SEC: One-Share Everywhere-Short Secret Sharing from Lossless Expanders

Muhammad Arzaki[1*], Muhammed F. Esgin[1], Amin Sakzad[1], Ron Steinfeld[1]

[1*]Faculty of Information Technology, Monash University, Melbourne, Australia.

*Corresponding author(s). E-mail(s): muhammad.arzaki@monash.edu;

Abstract

We introduce **SEC**, an expander-code-based ramp secret sharing scheme over $\mathbb{Z}$-modules, particularly suitable for lattice-based distributed cryptography. It is everywhere-short: bounded secret coefficients yield bounded share coefficients, and the integer reconstruction vectors have bounded entries. Moreover, each participant receives exactly one share, regardless of the number of parties and the privacy and reconstruction thresholds. This improves on prior schemes with bounded share coefficients and bounded entries in the integer reconstruction vectors, in which the number of sub-shares or shares per participant grows either with the number of parties and the threshold or with the security parameter. The construction rests on a coding-theoretic observation: in the bi-adjacency matrix of a lossless bipartite expander, unique-neighbor peeling completes every sufficiently small set of columns to a square unimodular submatrix. We prove correctness, everywhere-shortness, and linearity, and give a publicly computable entrywise bound on the reconstruction vector our algorithm returns for each authorized set. For privacy, we introduce hard-core local leakage resilience, which strengthens the leakage-resilience notion of Nielsen and Simkin (EUROCRYPT 2020) to protect a hard-core function of the secret, and we prove that **SEC** satisfies it against correlated multi-query local leakage under the **MatrixHint-MLWE** problem, which admits a reduction from the standard **MLWE** problem. For $n = 60$ participants and security parameter $\kappa = 128$, **SEC** achieves privacy threshold **6** against 2^{64} leakage queries and empirical reconstruction threshold **46**. The ideal entropy-coded secret and share sizes are about **3.5** KiB and **6.4** KiB, respectively.

Keywords: Secret sharing, Expander codes, Lossless expanders, Everywhere-short secret sharing, Lattice-based cryptography, Leakage resilience

MSC Classification: 94A60 , 94A62 , 94B05

1 Introduction

1.1 Background and Related Work

Secret sharing, introduced independently by Shamir and Blakley [1, 2], distributes a secret among n parties so that authorized sets can reconstruct it while unauthorized sets learn nothing about it. In a t-out-of-n threshold scheme, every set of at least t parties is authorized and every smaller set is unauthorized. Ramp schemes permit a gap between a privacy threshold and a larger reconstruction threshold, enabling a better information rate than perfect threshold sharing [3, 4].

Secret sharing is closely connected to coding theory. McEliece and Sarwate observed that Shamir's scheme is based on Reed-Solomon codes [5], and Massey characterized the access structure of a code-based scheme through the minimal codewords of the dual code [6]. Subsequent work constructs secret sharing schemes from random and algebraic codes [7, 8] and from span programs over rings and modules [9].

Black-box secret sharing, introduced by Desmedt and Frankel [10], works uniformly over finite abelian groups and reconstructs through group-independent integer linear combinations. Threshold black-box schemes with asymptotically optimal $O(\log n)$ expansion follow from number-theoretic and coding-theoretic constructions [11–13], while Cramer and Xing obtain near-threshold schemes with constant expansion [13]. Applebaum et al. construct near-threshold schemes with constant share size and $O(n)$-addition sharing and reconstruction by combining graph-based erasure codes with an extraction vector having small integer coefficients [14]. The near-threshold schemes of [13, 14] will be the closest relatives of SEC, and the construction of Applebaum et al. [14] is the closest structurally because it also uses sparse bipartite-graph codes and reusable public parameters. The requirement of Applebaum et al. [14] that the extraction vector lie outside the row span of every unauthorized restriction modulo every prime is analogous to the integral column-module condition of SEC: both prevent an unauthorized coalition from isolating the secret through exact linear algebra over $\mathbb{Z}$ or modulo an integer. Black-box schemes use uniform masks and obtain perfect privacy, whereas SEC uses the bounded-coefficient masks required in the lattice setting discussed next.

These coding-theoretic and black-box constructions provide the algebraic background for SEC, while lattice-based threshold cryptography supplies its principal application setting. Large-scale quantum computers would break the factoring and discrete-logarithm assumptions underlying efficient classical threshold cryptography, and the NIST call for multi-party threshold schemes has made practical post-quantum threshold signatures a concrete target [15]. Classical schemes such as FROST provide partially non-interactive signing and native identifiable abort by verifying partial signatures directly against participant-specific public keys [16]. In a Fiat-Shamir lattice signature such as Raccoon, however, the response $\boldsymbol{z} = c\boldsymbol{s} + \boldsymbol{r}$ must remain short [17]. Shamir sharing is unsuitable for directly generating such responses because its shares are uniform modulo q and its Lagrange reconstruction coefficients may be large. Consequently, the resulting partial responses are neither short nor safe to reveal, as a bounded-coefficient mask cannot hide the large share-dependent term [18, 19]. Threshold Raccoon therefore combines Shamir sharing with one-time additive masks [18], an

approach adopted by subsequent schemes [20–23]. Because masked partial responses cannot be verified directly, handling malicious signers requires robust, interactive, or proof-based mechanisms [24–26].

Small reconstruction coefficients have also been studied independently. Ball et al. prove share-size lower bounds for linear threshold schemes with binary reconstruction, motivated by lattice-based cryptography [27], and $\{0,1\}$-linear sharing has been used in threshold cryptosystems based on fully homomorphic encryption [28]. Benaloh-Leichter sharing [29], instantiated with Boppana's monotone formulas [30], gives reconstruction coefficients in $\{0,1\}$ or $\{0,\pm 1\}$, depending on the presentation, with $O(\min\{k_c, n-k_c\}^{4.3}\log n)$ shares per participant on average [27, 31], where k_c is the reconstruction threshold. Small reconstruction coefficients alone do not keep partial responses short, because the shares themselves must also have bounded coefficients. This motivates *everywhere-short* sharing: bounded secret coefficients yield bounded share coefficients, while the integer reconstruction vectors have bounded entries. These properties keep partial responses short and permit direct verification against participant-specific public material [19, 32].

Existing everywhere-short constructions assign multiple atomic sub-shares with bounded coefficients to each participant (Table 1). The sDKG framework of [32] uses replicated sharing, with combinatorial share multiplicity, and random-assignment ramp sharing, which uses several parallel sharings. Hermine uses Vandermonde sharing, which adapts the multiplicative construction of [33] to the additive group $(\mathcal{R}^\ell, +)$ and has a share count that grows with the threshold and $\log n$. For general thresholds, its best known per-participant upper bound is slightly superpolynomial, although it is claimed to remain concretely practical over Hermine's intended range [19, 34]. This multiplicity can substantially increase long-term secret-share storage and, in applications with participant-specific verification, the corresponding public material. This issue is particularly relevant on constrained platforms, where post-quantum cryptography must operate under limited memory, computation, and bandwidth [35–38]. We therefore study whether everywhere-short secret sharing over $\mathbb{Z}$-modules can be achieved with exactly one share per participant.

1.2 Our Contributions

We introduce SEC, an expander-code-based ramp secret sharing scheme over $\mathbb{Z}$-modules for lattice-based threshold cryptography. Unlike prior everywhere-short constructions, SEC assigns exactly one bounded-coefficient share to each participant, regardless of the number of parties and the privacy and reconstruction thresholds. As summarized in Table 1, it is the only scheme in the comparison that combines one share per participant, bounded share coefficients, and bounded entries in the integer reconstruction vectors. Privacy holds against every coalition of at most k_p parties. The empirical reconstruction threshold k_c is the smallest tested coalition size such that, for every tested size $t \geq k_c$, a uniformly random coalition of size t reconstructs with probability at least 0.99, where this probability is computed exactly or estimated by sampling. A larger deterministic reconstruction threshold k_c^{det} guarantees reconstruction for every coalition of at least that size. All 17 evaluated instances, with

Table 1: Comparison of SEC with secret sharing schemes with integer or short reconstruction used in lattice-based distributed cryptography

	Shamir [1]	Benaloh-Leichter + Boppana [29, 30]	Replicated	Random-assignment ramp	Vandermonde [33, 34]	SEC (this work)
Used in	[18, 24–26]	[22, 31]	sDKG [32]	sDKG [32]	Hermine [19]	–
Access	threshold	threshold	threshold	ramp[c]	threshold	ramp[d]
Shares per participant	1	$O(\mu^{4.3}\log n)$[a]	$\binom{n-1}{k_c-1}$	typically κ	$O\left(\binom{k_c-1+\lceil \lg n\rceil}{k_c-1}\right)$	1
Bounded share coefficients	×	×[b]	✓	✓	✓	✓
Reconstruction coefficients	large (Lagrange)	$\{0,\pm 1\}$	binary	binary	binary	small
Privacy	perfect	perfect[b]	comput.	comput.	comput.	comput.
Secret size[e]	–	–	6.02 KiB	6.28 KiB	3.76 KiB	3.54 KiB
Per-participant share storage[e]	–	–	259 GiB	964 KiB	699 KiB	6.42 KiB

Note: k_p and k_c denote the privacy and reconstruction thresholds, and $k_c = k_p + 1$ for a threshold scheme. "Binary" means reconstruction coefficients in $\{0,1\}$, while "small" permits bounded signed integers; "comput." abbreviates computational. "Shares per participant" counts elements of $\mathcal{R}^\ell$ for a module rank ℓ determined by the security parameter and independent of n; the size of each element, which may depend mildly on the parameters, is accounted for in the per-participant share storage.

[a] Here $\mu := \min\{k_c, n - k_c\}$ [27, 31].

[b] As instantiated in [22, 31], the sharing randomness is uniform over $\mathcal{R}_q$, so the share coefficients are not bounded and privacy is perfect; sampling bounded-coefficient pieces instead would bound the shares but make privacy computational.

[c] The privacy and reconstruction thresholds hold with overwhelming probability over a public seed [32].

[d] The privacy threshold k_p holds for every coalition of at most k_p parties, given a lossless expander, while the reconstruction threshold k_c is selected empirically.

[e] Ideal entropy-coded estimates of the secret and of all shares held by one participant at $\kappa = 128$ and $n = 60$ (Appendix J); SEC uses $Q = 2^{64}$ and privacy threshold $k_p = 6$; the replicated and Vandermonde schemes use the privacy-matched threshold $k_c = 7$, so that they tolerate the same 6 corrupted parties as SEC, while the reported ramp-sDKG storage is independent of the selected comparison threshold. The comparison schemes are extrapolated beyond their reported parameter ranges. The sDKG values include the error component $\boldsymbol{e}$, as does the systematic SEC secret. The Vandermonde values replace $\lg \sigma_{\mathbf{s}}$ by $\lceil \lg \sigma_{\mathbf{s}} \rceil$, adding less than one bit per coefficient, and report the participant with the fewest sub-shares.

$10 \le n \le 60$, satisfy $\tau_p := k_p/n \ge 0.1$ and $k/n < \tau_c := k_c/n \le (k+6)/n$, where k/n is the rate of the underlying code.

We prove correctness, everywhere-shortness, and $\mathbb{Z}$-linearity, and give a publicly computable entrywise bound on the reconstruction vector returned by our algorithm

for each authorized set. Because bounded-coefficient shares can reveal partial information about a bounded secret, SEC does not provide perfect information-theoretic privacy. Instead, we introduce a hard-core local leakage-resilience framework that strengthens [39, Definition 5] for the lattice setting. Under suitable parameters, we prove that a linear hard-core function of the secret, corresponding to an MLWE instance, remains computationally indistinguishable from uniform given the shares of a private coalition and prescribed noisy linear leakage from all shares, assuming the hardness of MatrixHint-MLWE, which in turn follows from the hardness of standard MLWE. This leakage is correlated across shares and queries and includes the partial responses arising in threshold signing, making the resulting privacy guarantee suitable for lattice-based threshold cryptography. Informally, our main result is as follows.

Theorem 1 (Main result, informal) *Let* SEC *be instantiated for n participants from a lossless bipartite expander, with privacy threshold k_p. Then the following hold.*

1. *Each participant holds exactly one share, whose coefficients are bounded whenever the coefficients of the secret and of the sharing randomness are bounded.*
2. *Every set for which the public reconstruction algorithm finds an integer reconstruction vector, called a certified set, recovers the secret exactly as an integer linear combination of its shares. The entries of the reconstruction vector satisfy a publicly computable bound. In particular, every coalition of at least $k_c^{\mathsf{det}} = n - b + 1$ parties is certified, where $b \geq \lfloor n/k \rfloor$ is the number of pairwise disjoint certified supports planted by the construction and k is the dimension of the underlying code.*
3. *For every coalition of at most k_p parties and every admissible choice of at most Q prescribed noisy linear leakage queries on all shares, a hard-core function of the secret remains pseudorandom under* MLWE *with secret width σ, provided that the widths of the sharing randomness and the leakage masks exceed σ by factors proportional to $\sqrt{k_p}$ and $\sqrt{Q}$, respectively, up to factors depending on n and the expander.*

Item 1 is formalized in Lemma 12, Item 2 in Lemmas 8, 11 and 16, and Item 3 in Theorem 4.

The main coding-theoretic ingredient may be of independent interest: in the bi-adjacency matrix of a lossless bipartite expander whose expansion exceeds half the left degree, unique-neighbor peeling completes every sufficiently small set of columns to a square unimodular submatrix (Lemma 4); we are not aware of a previous formulation of this unimodularity conclusion.

1.3 Technical Overview

Following the coding-theoretic approach of McEliece and Sarwate [5], SEC shares a secret through a public integer distribution matrix $\mathbf{G} := \begin{bmatrix}\mathbf{g}\ \hat{\mathbf{G}}\end{bmatrix} \in \mathbb{Z}^{n\times k}$, where $\hat{\mathbf{G}}$ is a signed and permuted transpose of the bi-adjacency matrix of a lossless bipartite expander and $\mathbf{g}$ is a secret-bearing column. The dealer multiplies $\mathbf{G}$ by a matrix whose first row is the secret and whose remaining rows are short Gaussian masks, giving one row of the product to each participant. A set $T \subseteq [n]$ is certified, and reconstructs exactly, when $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$, that is, when some $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ satisfies $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$. Applying the same integer combination to the shares recovers the secret;

our reconstruction algorithm searches for such a vector with small entries, and these entries satisfy the publicly computable bound of Lemma 16.

Because the masks are short rather than uniform, privacy is not perfect. Unlike sufficiently small sets of Shamir shares over a finite field, which are information-theoretically independent of the secret, shares generated from a sparse integer matrix offer no analogous guarantee: a coalition may attempt to cancel the masking contribution, possibly modulo an adversarially chosen integer $q' \geq 2$, while retaining a noiseless linear function of the secret. We rule out such attacks through an integral column-module condition, which the lossless-expander construction satisfies deterministically for every coalition of at most k_p parties, owing to the unimodular submatrices obtained by unique-neighbor peeling. This integral guarantee is stronger than sparsity or full rank over a fixed finite field and explains why an arbitrary sparse code does not suffice.

The structural condition excludes noiseless leakage but does not hide the noisy linear information carried by short shares. We therefore formalize privacy through the hard-core leakage experiment HC-Leak, in which the adversary must distinguish a designated hard-core function of the secret from uniform, given the shares of a private coalition and prescribed local leakage from every share. In our application, this function is the systematic MLWE target $f_{\boldsymbol{A}}(\boldsymbol{s}) := \boldsymbol{A}\boldsymbol{s} \bmod q$, where $\boldsymbol{A} := \left[\boldsymbol{I}\ \boldsymbol{A}'\right]$, and the leakage consists of noisy linear functions of the shares, such as the partial responses revealed by a threshold signing protocol.

The main challenge is that every share is a linear function of the same masking matrix, so the corrupted shares and the leakage released across multiple queries are jointly correlated. Treating them as independent hints would discard this correlation. We instead express the complete adversarial view as one stacked linear hint with a jointly distributed perturbation and bound it uniformly over admissible coalitions and leakage patterns. The integral image and kernel properties supplied by the unimodular submatrices allow us to apply the linear-transformation theorem of [40], showing that the joint perturbation is statistically close to the elliptic discrete Gaussian required by MatrixHint-MLWE. This enables a reduction of the multi-query HC-Leak experiment to MatrixHint-MLWE, whose hardness in turn follows from that of standard MLWE.

1.4 Organization

Section 2 recalls the required background and introduces the unimodular-submatrix lemma for lossless expanders, the variant of MatrixHint-MLWE used in our reduction, and the notions of everywhere-short secret sharing and hard-core local leakage resilience. Section 3 defines SEC and constructs its distribution matrix. Section 4 proves correctness, everywhere-shortness, and linearity, while Section 5 establishes multi-query privacy. Section 6 reports the concrete parameters and evaluation, and Section 7 discusses limitations and open problems. The appendices provide supporting definitions, proofs, and analyses, as well as a size comparison with other short-share secret sharing schemes (Appendix J).

2 Preliminaries

2.1 General Notation

For integers $a \le b$, let $[a,b] := \{a,\dots,b\}$ and $[n] := [1,n]$. We write $x \leftarrow \mathcal{D}$ for sampling from $\mathcal{D}$ and $x \xleftarrow{\$} S$ for uniform sampling from a finite set S. The notation $T \xleftarrow{\$} \binom{[n]}{t}$ means that T is sampled uniformly from the t-element subsets of $[n]$. All algorithms are probabilistic polynomial time unless stated otherwise, and $\mathsf{negl}(\kappa)$ denotes a negligible function of the security parameter κ.

Let $\mathcal{R} := \mathbb{Z}[X]/(X^d+1)$ and $\mathcal{R}_q := \mathcal{R}/(q\mathcal{R})$, where d is a power of two. Bold italic letters denote objects over $\mathcal{R}$ or $\mathcal{R}_q$, and bold upright letters denote coefficient embeddings or objects over $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$, or $\mathbb{Z}_q$. Thus, $\boldsymbol{a} \in \mathcal{R}^n$ is identified with $\mathbf{a} \in \mathbb{Z}^{dn}$, and $\boldsymbol{A} \in \mathcal{R}^{m\times n}$ with its induced coefficient-embedded map $\mathbf{A} \in \mathbb{Z}^{dm\times dn}$; analogously modulo q. Vectors are columns unless stated otherwise.

For a vector $\mathbf{v}$, its i-th entry is $\mathbf{v}[i]$, $(\mathbf{v})_T$ denotes its restriction to T, and $\mathrm{supp}(\mathbf{v}) := \{i : \mathbf{v}[i] \neq 0\}$. For a matrix $\mathbf{A}$, we use $\mathbf{A}[i,j]$, $\mathbf{A}[i,\cdot]$, and $\mathbf{A}[\cdot,j]$ for its entries, rows, and columns, and $\mathbf{A}[T,J]$ for a submatrix. We abbreviate $\mathbf{A}_T := \mathbf{A}[T,\cdot]$ and $\mathbf{A}_{(J)} := \mathbf{A}[\cdot,J]$. For a commutative ring or field Σ, we write $\mathrm{Col}_\Sigma(\mathbf{A})$, $\mathrm{Row}_\Sigma(\mathbf{A})$, and $\ker_\Sigma(\mathbf{A})$ for the column module, row module, and kernel of $\mathbf{A}$. When Σ is a field, $\mathrm{rank}_\Sigma(\mathbf{A})$ denotes its rank over Σ. We retain the subscript when the scalar domain matters.

For a tuple of matrices $(\mathbf{A}_i)_{i\in I}$ indexed by a finite ordered set I, $\left[(\mathbf{A}_i)_{i\in I}\right]$ denotes its ordered vertical stack; the same convention applies to vectors. Additional notation and linear-algebra conventions appear in Appendix A.

2.2 Lattices and Discrete Gaussians

A rank-r lattice generated by linearly independent vectors $\mathbf{b}_1,\dots,\mathbf{b}_r \in \mathbb{R}^m$ is $\Lambda := \{\sum_{i=1}^r z_i\mathbf{b}_i : z_i \in \mathbb{Z}\} = \mathcal{L}(\mathbf{B}) = \mathbf{B}\mathbb{Z}^r$, where $\mathbf{B} := \left[\mathbf{b}_1 \ \cdots \ \mathbf{b}_r\right]$. Its dual within $\mathrm{span}_{\mathbb{R}}(\Lambda)$ is $\Lambda^* := \{\mathbf{y} \in \mathrm{span}_{\mathbb{R}}(\Lambda) : \langle \mathbf{x},\mathbf{y}\rangle \in \mathbb{Z},\ \forall \mathbf{x} \in \Lambda\} = \mathbf{B}(\mathbf{B}^\top\mathbf{B})^{-1}\mathbb{Z}^r$. If Λ is full rank in $\mathbb{R}^m$, then $\Lambda^* = \mathbf{B}^{-\top}\mathbb{Z}^m$.

For $s > 0$, define $\rho_s(\mathbf{x}) := \exp\left(-\pi\frac{\|\mathbf{x}\|_2^2}{s^2}\right)$ and $\rho_s(S) := \sum_{\mathbf{x}\in S}\rho_s(\mathbf{x})$ for every countable set S for which the sum converges. For a countable set $S \subseteq \mathbb{R}^m$ for which the normalizing sum below is finite, a center $\mathbf{c} \in \mathbb{R}^m$, and a positive-definite parameter matrix $\mathbf{\Sigma} \in \mathbb{R}^{m\times m}$, let $\mathcal{D}_{S,\mathbf{c},\sqrt{\mathbf{\Sigma}}}$ denote the discrete Gaussian over S with probability mass function

$$\mathcal{D}_{S,\mathbf{c},\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) := \frac{\exp\left(-\frac{1}{2}(\mathbf{x}-\mathbf{c})^\top\mathbf{\Sigma}^{-1}(\mathbf{x}-\mathbf{c})\right)}{\displaystyle\sum_{\mathbf{z}\in S}\exp\left(-\frac{1}{2}(\mathbf{z}-\mathbf{c})^\top\mathbf{\Sigma}^{-1}(\mathbf{z}-\mathbf{c})\right)}, \qquad \mathbf{x} \in S. \tag{1}$$

We omit the center when $\mathbf{c} = \mathbf{0}$, write $\mathcal{D}_{S,\mathbf{c},\sigma}$ when $\mathbf{\Sigma} = \sigma^2\mathbf{I}_m$, and abbreviate this further to $\mathcal{D}_\sigma$ when $S = \mathbb{Z}^m$, $\mathbf{c} = \mathbf{0}$, and the dimension is clear. The matrix $\mathbf{\Sigma}$ is a Gaussian parameter matrix, not necessarily the exact covariance matrix of the resulting discrete distribution.

Using the standard-deviation convention of our Gaussians and the Matrix Hint-MLWE conditions in [24], we define the smoothing parameter of a nonzero lattice Λ for $\epsilon > 0$ by

$$\eta_\epsilon(\Lambda) := \inf\left\{\sigma > 0 : \sum_{\mathbf{y}\in\Lambda^*\setminus\{\mathbf{0}\}} \exp\left(-2\pi^2\sigma^2\|\mathbf{y}\|_2^2\right) \le \epsilon\right\}. \tag{2}$$

For numerical estimates, we use $\eta_\epsilon(\mathbb{Z}^N) \approx \sqrt{\frac{\ln(2N/\epsilon)}{2\pi^2}}$, while the formal statements retain $\eta_\epsilon(\mathbb{Z}^N)$ explicitly. The relation to the ρ_s-normalization of [40, 41] appears in Appendix B.

2.3 Linear Codes and Lossless Expanders

2.3.1 Linear Codes and Distance Transfer

A linear code C of length n over a commutative ring or field Σ is a Σ-submodule of Σ^n, whose codewords are written as columns. Its generator and parity-check matrices satisfy $C = \mathrm{Col}_\Sigma(\mathbf{G}) = \ker_\Sigma(\mathbf{H})$, and its dual is $C^\perp = \ker_\Sigma(\mathbf{G}^\top)$. For $\mathbf{u} \in \Sigma^n$, let $\mathrm{supp}(\mathbf{u})$ and $wt_H(\mathbf{u}) := |\mathrm{supp}(\mathbf{u})|$ denote its support and Hamming weight, and let $d(C)$ denote the minimum weight of a nonzero codeword. Over a field, $d(C)$ equals the smallest number of linearly dependent columns of a parity-check matrix.

In our construction, a parity-check matrix over a finite field is also viewed over $\mathbb{Z}$ and $\mathbb{Q}$ through fixed integer representatives. The following lemma, whose short proof we include for completeness, shows that lifting the matrix to $\mathbb{Q}$ does not decrease the associated code distance.

Lemma 1 (Distance under lifting to $\mathbb{Q}$) *Let p be prime and let $\mathbf{H} \in \mathbb{F}_p^{m\times n}$. Fix an integer lift of $\mathbf{H}$, still denoted by $\mathbf{H}$, whose reduction modulo p is the original matrix. Define $C_p := \ker_{\mathbb{F}_p}(\mathbf{H}) \subseteq \mathbb{F}_p^n$ and $C_\mathbb{Q} := \ker_\mathbb{Q}(\mathbf{H}) \subseteq \mathbb{Q}^n$. Then $d(C_\mathbb{Q}) \ge d(C_p)$.*

The proof clears denominators and removes the largest common power of p; it appears in Appendix D. Moreover, $\ker_\mathbb{Q}(\mathbf{UH}) = \ker_\mathbb{Q}(\mathbf{H})$ for every $\mathbf{U} \in GL_m(\mathbb{Z})$. Hence, unimodular row transformations, including signed permutations, preserve the rational code and its distance.

2.3.2 Lossless Bipartite Expanders and Expander Codes

Let $G = (L, R, E)$ be a simple bipartite graph with $|L| = n$ and $|R| = m$, whose vertices we identify with $[n]$ and $[m]$, respectively. The graph is $c_{\exp}$-left-regular if every left vertex has degree $c_{\exp}$. For $S \subseteq L$, let $N(S) \subseteq R$ be its neighbor set; a vertex of R adjacent to exactly one vertex in S is called a *unique neighbor* of S.

Definition 1 (Lossless bipartite expander and expander code, cf. [42], [43, Chapter 11]) Let $0 < \epsilon_{\exp} < 1/2$ and $\gamma \in (0, 1]$. An $(n, m, c_{\exp}, \gamma, (1-\epsilon_{\exp})c_{\exp})$ *lossless bipartite expander* is a

$c_{\exp}$-left-regular bipartite graph $G = (L, R, E)$ with $|L| = n$ and $|R| = m$ such that $|N(S)| \geq (1 - \epsilon_{\exp})c_{\exp}|S|$ for every $S \subseteq L$ with $|S| \leq \gamma n$. Its bi-adjacency matrix $\mathbf{H} \in \{0,1\}^{m \times n}$ has rows indexed by R, columns indexed by L, and $\mathbf{H}[i,j] = 1$ exactly when i and j are adjacent. The associated *binary expander code* is $C(G) := \ker_{\mathbb{F}_2}(\mathbf{H}) \subseteq \mathbb{F}_2^n$.

Since $\mathbf{H}$ has entries in $\{0,1\}$, it can be read over any ring. The binary code $C(G)$ is the standard expander code of [42] and is used only for the distance bound below, whereas our construction uses $\mathbf{H}$ over $\mathbb{Z}$, and Lemma 1 transfers distance bounds from $\mathbb{F}_p$ to the rational kernel.

The following theorem restates known existence results for convenience: its first part is [44, Lemma 1.9] and [43, Theorem 11.2.3], and its second part, which bounds the failure probability of a random construction, is [45, Lemma 2]. We use it only to show that constant-rate families of lossless bipartite expanders with the required asymptotic parameters exist. Our concrete instances do not rely on it and are certified directly in Section 6.1.

Theorem 2 (Existence of lossless bipartite expanders [44, Lemma 1.9], [43, Theorem 11.2.3]) *Fix $0 < \epsilon_{\exp} < 1/2$. For sufficiently large integers $n \geq m$, there exists an $(n, m, c_{\exp}, \gamma, (1 - \epsilon_{\exp})c_{\exp})$ lossless bipartite expander with $c_{\exp} = \Theta\left(\frac{\log(n/m)+\log(1/\epsilon_{\exp})}{\epsilon_{\exp}}\right)$ and $\gamma = \Theta\left(\frac{\epsilon_{\exp} m}{c_{\exp} n}\right)$. Moreover, if $m = \Theta(n)$, then by [45, Lemma 2], for every fixed integer $a \geq 1$, one can choose $c_{\exp} = \Theta\left(\frac{\log(n/m)+\log(1/\epsilon_{\exp})+a}{\epsilon_{\exp}}\right)$ such that a random $c_{\exp}$-left-regular simple bipartite graph, obtained by letting each left vertex independently choose a uniform $c_{\exp}$-subset of R, has the required expansion with probability at least $1 - n^{-a}$. The hidden constants may depend on the fixed rate, expansion gap, and failure exponent.*

Unique-Neighbor Properties

For $S \subseteq L$, let $U(S) := \{v \in R : v \text{ is adjacent to exactly one vertex in } S\}$ denote its unique-neighbor set. Lossless expansion implies that every sufficiently small set of left vertices has many unique neighbors.

Lemma 2 (Unique-neighbor bound [43, Lemma 11.3.4]) *Let $G = (L, R, E)$ be an $(n, m, c_{\exp}, \gamma, (1 - \epsilon_{\exp})c_{\exp})$ lossless bipartite expander with $0 < \epsilon_{\exp} < 1/2$. Then, for every $S \subseteq L$ satisfying $|S| \leq \gamma n$, $|U(S)| \geq c_{\exp}(1 - 2\epsilon_{\exp})|S|$.*

We extend Lemma 2 to a larger range, with the weaker conclusion that every such set has a unique neighbor.

Lemma 3 (Extended unique-neighbor property) *Let $G = (L, R, E)$ be an $\big(n, m, c_{\exp}, \gamma, (1 - \epsilon_{\exp})c_{\exp}\big)$ lossless bipartite expander with $0 < \epsilon_{\exp} < 1/2$ and $\lfloor \gamma n \rfloor \geq 1$. Then every nonempty set $S \subseteq L$ satisfying $|S| < 2(1 - \epsilon_{\exp})\lfloor \gamma n \rfloor$ has a unique neighbor.*

The proof applies Lemma 2 to a subset of size $\lfloor \gamma n \rfloor$ and counts the neighbors of the remaining vertices; see Appendix D. The same unique-neighbor argument gives $d(C(G)) \geq 2(1-\epsilon_{\exp})\gamma n$ for every $(n, m, c_{\exp}, \gamma, (1-\epsilon_{\exp})c_{\exp})$ lossless bipartite expander G [43, Theorem 11.3.8].

Unimodular Submatrices

An integer matrix contains a $t \times t$ *unimodular submatrix* if some choice of t rows and t columns yields a submatrix with determinant ± 1. Unique-neighbor peeling repeatedly removes a column that is the only remaining neighbor of some row. As in iterative erasure decoding, it stops before removing all columns exactly when the remaining columns form a stopping set, a nonempty set without unique neighbors [46], which Lemma 3 rules out for sufficiently small sets. Applying the domain-independent peeling method of [14] to a column restriction $\mathbf{H}_{(T)}$ and reversing the peeling order produces a triangular submatrix with unit diagonal. It is therefore invertible over $\mathbb{Z}$ and modulo every prime, properties used in both the structural and computational privacy analyses. Although unique-neighbor peeling and stopping-set arguments are standard, we are not aware of a previous formulation that extracts this unimodular-submatrix conclusion over $\mathbb{Z}$. We later use the following lemma to show that the necessary condition for privacy holds with probability 1 for every coalition T of size at most k_p.

Lemma 4 (Unimodular submatrices from lossless expansion) *Let $\mathbf{H} \in \{0,1\}^{m \times n}$ be the bi-adjacency matrix of an $(n, m, c_{\exp}, \gamma, (1-\epsilon_{\exp})c_{\exp})$ lossless bipartite expander, where $0 < \epsilon_{\exp} < 1/2$ and $\lfloor \gamma n \rfloor \geq 1$. Then, for every nonempty set $T \subseteq [n]$ satisfying $|T| \leq \lceil 2(1-\epsilon_{\exp})\lfloor \gamma n \rfloor \rceil - 1$, the column restriction $\mathbf{H}_{(T)}$ contains a $|T| \times |T|$ unimodular submatrix.*

Proof Fix a nonempty $T \subseteq [n]$, and let $\tau := |T| \leq \lceil 2(1-\epsilon_{\exp})\lfloor \gamma n \rfloor \rceil - 1$. Since $\tau \in \mathbb{Z}$, this implies $\tau < 2(1-\epsilon_{\exp})\lfloor \gamma n \rfloor$. Set $T_0 := T$. For each $i \in [\tau]$, Lemma 3 gives a unique neighbor r_i of the nonempty set T_{i-1}. Let t_i be the unique vertex of T_{i-1} adjacent to r_i, and set $T_i := T_{i-1} \setminus \{t_i\}$. The indices $t_1, \ldots, t_\tau$ are distinct by construction. The indices $r_1, \ldots, r_\tau$ are also distinct because, after t_i is removed, r_i has no neighbors in T_i. Let $\mathbf{M}_0$ be the submatrix whose rows and columns are ordered as $r_\tau, \ldots, r_1$ and $t_\tau, \ldots, t_1$, respectively. For each i, the row r_i is adjacent to t_i and to no other vertex in $T_{i-1} = \{t_i, t_{i+1}, \ldots, t_\tau\}$. Hence $\mathbf{M}_0$ is triangular with unit diagonal and $\det(\mathbf{M}_0) = 1$. Reordering the rows and columns changes the determinant only by a sign, so $\mathbf{H}_{(T)}$ contains a $|T| \times |T|$ unimodular submatrix. □

Deterministic Peeling

The proof of Lemma 4 is constructive. The scheme SEC defined in Section 3 uses $\hat{\mathbf{G}} := (\mathbf{PH})^\top$, where $\mathbf{P}$ is a signed permutation matrix, so every nonzero entry of $\hat{\mathbf{G}}$ lies in $\{-1, 1\}$. For $U \subseteq [n]$, we peel the row restriction $\hat{\mathbf{G}}_U$ deterministically as follows. Initially, all rows of $\hat{\mathbf{G}}_U$ are active. Each step selects the smallest-index column, among those not yet selected, that has exactly one nonzero entry in the active rows, records that column together with the row containing this entry, and then deactivates the row. When $|U| \leq \lceil 2(1-\epsilon_{\exp})\lfloor \gamma n \rfloor \rceil - 1$, such a column exists at every step by Lemma 3, exactly as in the proof of Lemma 4. The selected columns then form a

matrix $\mathbf{M}_U \in GL_{|U|}(\mathbb{Z})$; after ordering its rows according to the peeling order, it is triangular with diagonal entries in $\{-1, 1\}$. The smallest-index rule makes $\mathbf{M}_U$ a deterministic function of $\hat{\mathbf{G}}$ and U; Algorithm 3 in Appendix D states the procedure as pseudocode. We call $\mathbf{M}_U$ the *peeled unimodular submatrix*; its conditioning enters the kernel-smoothing bound of Section 5.4.3.

2.4 Hardness Assumptions

Our privacy analysis relies on decisional Module-LWE in systematic form and Matrix Hint-MLWE, which augments an MLWE secret with noisy linear leakage.

2.4.1 Module Learning with Errors

Definition 2 (Module-LWE in systematic form, cf. [47]) Let $n, m, q, d \in \mathbb{Z}^+$, and let ψ be a distribution over $\mathbb{Z}$, also denoting its coefficient-wise extension to $\mathcal{R}_q$. Sample $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_q^{n\times m}$ and define $\overline{\boldsymbol{A}} := \big[\boldsymbol{I}_n \ \boldsymbol{A}\big]$. The decisional $\mathsf{MLWE}_{n,m,q,d,\psi}$ problem is to distinguish $(\boldsymbol{A}, \overline{\boldsymbol{A}}\boldsymbol{s})$ from $(\boldsymbol{A}, \boldsymbol{u})$, where $\boldsymbol{s} \leftarrow \psi^{n+m}$ and $\boldsymbol{u} \xleftarrow{\$} \mathcal{R}_q^n$. The advantage of an adversary $\mathcal{A}$ is $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{MLWE}}(\kappa) := \big|\Pr\big[1 \leftarrow \mathcal{A}(\boldsymbol{A}, \overline{\boldsymbol{A}}\boldsymbol{s})\big] - \Pr[1 \leftarrow \mathcal{A}(\boldsymbol{A}, \boldsymbol{u})]\big|$.

Writing $\boldsymbol{s} = \begin{bmatrix}\boldsymbol{e}\\ \boldsymbol{s}'\end{bmatrix}$ with $\boldsymbol{e} \leftarrow \psi^n$ and $\boldsymbol{s}' \leftarrow \psi^m$ gives $\overline{\boldsymbol{A}}\boldsymbol{s} = \boldsymbol{A}\boldsymbol{s}' + \boldsymbol{e}$; thus, the systematic secret combines the conventional MLWE error and secret. When $\psi = \mathcal{D}_{\mathbb{Z},\sigma}$, we abbreviate $\mathsf{MLWE}_{n,m,q,d,\psi}$ as $\mathsf{MLWE}_{n,m,q,d,\sigma}$, where σ is the source width.

2.4.2 Matrix Hint-MLWE

Problem Definition

Hint-MLWE augments an MLWE sample with noisy scalar multiples of its secret [41], while Matrix Hint-MLWE permits noisy linear images [24]. Following the rectangular single-hint formulation of [48], we allow multiple rectangular hints with possibly different output dimensions. As in Definition 2, let $\overline{\boldsymbol{A}} := \big[\boldsymbol{I}_n \ \boldsymbol{A}\big]$.

Definition 3 (Matrix Hint-MLWE with varying hint dimensions) Let $n, m, q, d, h \in \mathbb{Z}^+$ and $(n_i)_{i\in[h]} \in (\mathbb{Z}^+)^h$. Let $\chi_{\boldsymbol{s}}$ be a distribution over $\mathcal{R}^{n+m}$. For every $i \in [h]$, let $\chi_{\boldsymbol{y}_i}$ be a distribution over $\mathcal{R}^{n_i}$, and let $\mathcal{M}_i$ be a distribution over $\mathcal{R}^{n_i\times(n+m)}$. The problem $\mathsf{MatrixHint\text{-}MLWE}^{h,(\chi_{\boldsymbol{y}_i})_{i\in[h]},(\mathcal{M}_i)_{i\in[h]}}_{(n_i)_{i\in[h]},n,m,q,d,\chi_{\boldsymbol{s}}}$ is to distinguish: (1) (Real) $\Big(\boldsymbol{A}, \overline{\boldsymbol{A}}\boldsymbol{s}, (\boldsymbol{M}_i)_{i\in[h]}, (\boldsymbol{z}_i)_{i\in[h]}\Big)$, and (2) (Random) $\Big(\boldsymbol{A}, \boldsymbol{u}, (\boldsymbol{M}_i)_{i\in[h]}, (\boldsymbol{z}_i)_{i\in[h]}\Big)$, where $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_q^{n\times m}$, $\boldsymbol{u} \xleftarrow{\$} \mathcal{R}_q^n$, $\boldsymbol{s} \leftarrow \chi_{\boldsymbol{s}}$, and, independently for every $i \in [h]$, $\boldsymbol{M}_i \leftarrow \mathcal{M}_i$, $\boldsymbol{y}_i \leftarrow \chi_{\boldsymbol{y}_i}$, and $\boldsymbol{z}_i := \boldsymbol{M}_i\boldsymbol{s} + \boldsymbol{y}_i$. The secret is reduced coefficient-wise modulo q when forming $\overline{\boldsymbol{A}}\boldsymbol{s} \in \mathcal{R}_q^n$, while the hints are computed over $\mathcal{R}$, or equivalently over their integer coefficient embeddings. The distinguishing advantage of an adversary $\mathcal{A}$ is $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{MHMLWE}}(\kappa) := |\Pr[\mathcal{A}(\mathsf{Real}) = 1] - \Pr[\mathcal{A}(\mathsf{Random}) = 1]|$, where the probabilities are taken over the stated sampling and the internal randomness of $\mathcal{A}$.

Joint-Precision Hardness Bound

We use the elliptic posterior-resampling argument of [24, 41, 48] and prove the following variant for multiple hints. Rather than separately bounding $\|\mathbf{M}_i^\top \mathbf{M}_i\|_2 \|\mathbf{\Sigma}_{\mathbf{y}_i}^{-1}\|_2$, it retains the tighter joint term $\|\mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1} \mathbf{M}_i\|_2$. Under coefficient embedding, let $N := d(n+m)$ and $N_i := dn_i$ for $i \in [h]$.

Theorem 3 (Joint-precision hardness of Matrix Hint-MLWE) *Consider the setting of Definition 3. Suppose that, under coefficient embedding,* $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \mathbf{c}_\mathbf{s}, \sqrt{\mathbf{\Sigma}_\mathbf{s}}}$*, and, for every* $i \in [h]$*, the* i*-th perturbation is distributed as* $\mathbf{y}_i \leftarrow \mathcal{D}_{\mathbb{Z}^{N_i}, \mathbf{c}_{\mathbf{y}_i}, \sqrt{\mathbf{\Sigma}_{\mathbf{y}_i}}}$*, where* $\mathbf{c}_\mathbf{s} \in \mathbb{R}^N$*,* $\mathbf{\Sigma}_\mathbf{s} \succ 0$*, and, for every* $i \in [h]$*,* $\mathbf{c}_{\mathbf{y}_i} \in \mathbb{R}^{N_i}$ *and* $\mathbf{\Sigma}_{\mathbf{y}_i} \succ 0$*. Fix a tuple of hint matrices in the support of* $(\mathcal{M}_i)_{i\in[h]}$*, and let* $\mathbf{M}_i \in \mathbb{Z}^{N_i \times N}$ *be the coefficient embedding of its* i*-th component. Define the corresponding effective source width* $\sigma > 0$ *by* $1/\sigma^2 := 2\big(\|\mathbf{\Sigma}_\mathbf{s}^{-1}\|_2 + \sum_{i=1}^h \|\mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1} \mathbf{M}_i\|_2\big)$*. If* $\sigma \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^N)$ *for some* $0 < \epsilon \leq 1/2$*, then there exists an efficient reduction from* $\mathsf{MLWE}_{n,m,q,d,\sigma}$ *to the corresponding fixed-matrix Matrix Hint-MLWE problem, with additive distinguishing loss at most* 2ϵ*.*

The posterior-resampling proof uses $\mathbf{\Sigma}_\mathbf{s}^{-1} + \sum_i \mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1} \mathbf{M}_i$ as the precision of the secret conditioned on the hints and requires only an upper bound on its norm; see Appendix B.

Remark 1 (Interpreting the effective source width) For an isotropic secret, $\mathbf{\Sigma}_\mathbf{s} = \sigma_\mathbf{s}^2 \mathbf{I}_N$, the effective source width of Theorem 3 satisfies $1/\sigma^2 = 2\left(1/\sigma_\mathbf{s}^2 + \Xi\right)$, where $\Xi := \sum_{i=1}^h \|\mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1} \mathbf{M}_i\|_2$ is the leakage term. Hence $\sigma \leq \sigma_\mathbf{s}/\sqrt{2}$, with equality in the absence of hints, and the hints reduce σ further. If the hint perturbations are wide enough that $\Xi \leq c/\sigma_\mathbf{s}^2$ for a constant $c \geq 0$, then $\sigma \geq \sigma_\mathbf{s}/\sqrt{2(1+c)}$, so the reduction loses only a constant factor in the MLWE source width; for example, $\Xi \leq 1/\sigma_\mathbf{s}^2$ gives $\sigma \geq \sigma_\mathbf{s}/2$. In SEC, the hint perturbations are the sharing and leakage masks, so Ξ decreases as $\sigma_\mathbf{R}$ and $\sigma_{\boldsymbol{r}}$ grow; the concrete parameters of Section 6.4 directly require the resulting width σ_{MQ} to meet the MLWE security target.

The reduction remains valid from a smaller source width that still satisfies the smoothing condition.

Corollary 1 (Reduction from a smaller source width) *Fix a tuple of hint matrices as in Theorem 3, and let* σ *be its effective source width. If, for some* $0 < \epsilon \leq 1/2$*,* $\sqrt{2}\,\eta_\epsilon(\mathbb{Z}^{d(n+m)}) \leq \sigma' \leq \sigma$*, then there is an efficient reduction from* $\mathsf{MLWE}_{n,m,q,d,\sigma'}$ *to the corresponding Matrix Hint-MLWE problem, with additive distinguishing loss at most* 2ϵ*.*

The proof verifies positive definiteness of the correction parameter matrix and smoothing of the convolution step; see Appendix B.

2.5 Secret Sharing

We recall the access structures, syntax, and correctness of secret sharing, and then formalize the two additional properties used in this work: everywhere-shortness and hard-core local leakage resilience.

2.5.1 Basic Definitions

We recall the following definition of partial and ramp access structures.

Definition 4 (Partial and ramp access structures, cf. [14, Definition 1]) A partial access structure over n parties is a pair $\Gamma = (\Gamma_0, \Gamma_1)$ of nonempty families $\Gamma_0, \Gamma_1 \subseteq 2^{[n]}$ such that $B \not\subseteq A$ for every $A \in \Gamma_0$ and $B \in \Gamma_1$. Sets in Γ_0 are called *private*, whereas sets in Γ_1 are called *authorized*. For $0 < \tau_p < \tau_c \leq 1$, define $k_p := \lceil \tau_p n \rceil$ and $k_c := \lfloor \tau_c n \rfloor$, and assume $k_p < k_c$. The (τ_p, τ_c)-ramp access structure over n parties is the partial access structure $\Gamma = (\Gamma_0, \Gamma_1)$ given by $\Gamma_0 := \{U \subseteq [n] : |U| \leq k_p\}$ and $\Gamma_1 := \{T \subseteq [n] : |T| \geq k_c\}$.

Remark 2 (Privacy vs. reconstruction) In a ramp scheme, the privacy and reconstruction thresholds need not coincide, and no guarantee is prescribed for sets in neither Γ_0 nor Γ_1: for a (k_p, k_c)-ramp access structure, sets with $k_p < |U| < k_c$ may learn partial information, and some may even reconstruct the secret. Accordingly, our privacy statements are always quantified over $U \in \Gamma_0$, whereas our correctness statements are quantified over $T \in \Gamma_1$.

We use the following secret-sharing syntax, which allows the access structure to depend on the public parameters.

Definition 5 (Secret sharing syntax, cf. [14, Section 2.1]) Let $n \in \mathbb{Z}^+$, let S be a secret space, and let $(S_i)_{i\in[n]}$ be the share spaces. A secret sharing scheme $\mathsf{Sharing} = (\mathsf{Setup}, \mathsf{Deal}, \mathsf{Rec})$ consists of the following efficient algorithms:

1. $\mathsf{Setup}(1^\kappa)$ is a randomized algorithm that outputs public parameters pp and determines a partial access structure $\Gamma(\mathsf{pp}) = (\Gamma_0(\mathsf{pp}), \Gamma_1(\mathsf{pp}))$ over n parties.
2. $\mathsf{Deal}_{\mathsf{pp}}(s)$, on input $s \in S$, is a randomized algorithm that outputs shares $(s_i)_{i\in[n]} \in \prod_{i=1}^n S_i$.
3. $\mathsf{Rec}_{\mathsf{pp}}(T, (s_i)_{i\in T})$, on input $T \subseteq [n]$ and the corresponding shares, is a deterministic algorithm that outputs an element of $S \cup \{\perp\}$.

Definition 6 (Correctness, cf. [14, Definition 2]) Let $\mathsf{Sharing} = (\mathsf{Setup}, \mathsf{Deal}, \mathsf{Rec})$ be a secret sharing scheme. The scheme is ϵ_c-correct if, for every pp in the support of $\mathsf{Setup}(1^\kappa)$, every authorized set $T \in \Gamma_1(\mathsf{pp})$, and every secret $s \in S$, $\Pr\left[\mathsf{Rec}_{\mathsf{pp}}(T, (s_i)_{i\in T}) \neq s \;\middle|\; (s_i)_{i\in[n]} \leftarrow \mathsf{Deal}_{\mathsf{pp}}(s)\right] \leq \epsilon_c$, where the probability is over the randomness of Deal. The scheme is *perfectly correct* if $\epsilon_c = 0$.

The conventional pairwise-secret privacy notion is recalled in Appendix C. The privacy property used in this work is instead the distributional leakage-resilience notion defined in Section 2.5.3.

2.5.2 Everywhere-Short Secret Sharing

Everywhere-short secret sharing strengthens the syntax and correctness of Definitions 5 and 6 by requiring that, for every secret of bounded infinity norm, the shares have bounded infinity norm, up to a permitted shortness error, and the reconstruction coefficients have bounded infinity norm. Following prior work, we call such objects *short*. The term was introduced informally in [19] for schemes whose secrets, shares, and reconstruction coefficients are simultaneously short. Related coefficient bounds appear in [31], without short shares, and in [32], as properties of distributed key generation. We formalize the notion below, building on the integer-linear reconstruction of black-box secret sharing [11, 13].

Definition 7 (Everywhere-short secret sharing) Let $\mathsf{Sharing} = (\mathsf{Setup}, \mathsf{Deal}, \mathsf{Rec})$ be a secret sharing scheme over n parties, and for pp in the support of $\mathsf{Setup}(1^\kappa)$ let $\Gamma(\mathsf{pp}) = (\Gamma_0(\mathsf{pp}), \Gamma_1(\mathsf{pp}))$ be the associated partial access structure. Let L denote the total number of atomic sub-shares, indexed by $[L]$. Suppose that the secret space S and the atomic share spaces $(\widehat{S}_j)_{j\in[L]}$ are contained in a common $\mathbb{Z}$-module with a fixed integer coefficient embedding, and let $\|\cdot\|_\infty$ denote the infinity norm, namely the largest absolute value of a coefficient. Let $(I_i)_{i\in[n]}$ be a partition of $[L]$, where party i holds the aggregate share $(\boldsymbol{s}_j)_{j\in I_i} \in S_i := \prod_{j\in I_i} \widehat{S}_j$. We say that $\mathsf{Sharing}$ is *everywhere-short* with secret bound β_{sec}, share bound β_{sh}, reconstruction bound $\beta_{\boldsymbol{\lambda}}$, and shortness error ϵ_{sh} if the following hold.

1. **Short shares:** for every pp in the support of $\mathsf{Setup}(1^\kappa)$ and every $\boldsymbol{s} \in S$ with $\|\boldsymbol{s}\|_\infty \le \beta_{\mathsf{sec}}$, if $\big((\boldsymbol{s}_j)_{j\in I_i}\big)_{i\in[n]} \leftarrow \mathsf{Deal}_{\mathsf{pp}}(\boldsymbol{s})$, then $\Pr\big[\max_{j\in[L]} \|\boldsymbol{s}_j\|_\infty > \beta_{\mathsf{sh}}\big] \le \epsilon_{\mathsf{sh}}$, where the probability is over the randomness of Deal.
2. **Short linear reconstruction:** for every pp in the support of $\mathsf{Setup}(1^\kappa)$ and every $T \in \Gamma_1(\mathsf{pp})$, $\mathsf{Rec}_{\mathsf{pp}}(T,\cdot)$ is defined by $\big((\boldsymbol{s}_j)_{j\in I_i}\big)_{i\in T} \mapsto \sum_{j\in\bigcup_{i\in T} I_i} \lambda_{T,j}\boldsymbol{s}_j$ for a public reconstruction vector $\boldsymbol{\lambda}_T \in \mathbb{Z}^L$ satisfying $\mathrm{supp}(\boldsymbol{\lambda}_T) \subseteq \bigcup_{i\in T} I_i$, depending only on pp and T, hence independent of the shares, with $\|\boldsymbol{\lambda}_T\|_\infty \le \beta_{\boldsymbol{\lambda}}$. Moreover, this map takes the value $\boldsymbol{s}$ on every sharing in the support of $\mathsf{Deal}_{\mathsf{pp}}(\boldsymbol{s})$, for every $\boldsymbol{s} \in S$ with $\|\boldsymbol{s}\|_\infty \le \beta_{\mathsf{sec}}$.

The *expansion factor* of the scheme is L/n.

All coefficient-magnitude bounds in Definition 7 use the infinity norm, as in [19, 31, 32]. Euclidean bounds follow from $\|\mathbf{x}\|_2 \le \sqrt{N}\|\mathbf{x}\|_\infty$ for $\mathbf{x} \in \mathbb{Z}^N$. Moreover, whenever $\max_{j\in[L]} \|\boldsymbol{s}_j\|_\infty \le \beta_{\mathsf{sh}}$, every reconstructed combination satisfies $\left\|\sum_{j\in\bigcup_{i\in T} I_i} \lambda_{T,j}\boldsymbol{s}_j\right\|_\infty \le wt_H(\boldsymbol{\lambda}_T)\,\beta_{\boldsymbol{\lambda}}\,\beta_{\mathsf{sh}}$ for every $T \in \Gamma_1(\mathsf{pp})$.

Remark 3 Since $\boldsymbol{\lambda}_T$ may depend only on pp and T, it may be published without affecting privacy. The reconstruction vector need not be unique: if $\boldsymbol{\lambda}_T$ and $\boldsymbol{\lambda}'_T$ both satisfy Item 2 in Definition 7 for the same (pp, T), then their induced linear combinations agree on every sharing in the support of $\mathsf{Deal}_{\mathsf{pp}}(\boldsymbol{s})$, for every $\boldsymbol{s} \in S$ with $\|\boldsymbol{s}\|_\infty \le \beta_{\mathsf{sec}}$, since both recover the shared secret.

Binary reconstruction [27, 28] is recovered by restricting the reconstruction vector in Item 2 of Definition 7 to $\boldsymbol{\lambda}_T \in \{0,1\}^L$, in which case $\beta_{\boldsymbol{\lambda}} = 1$. The preceding definition more generally permits short signed integer reconstruction coefficients.

2.5.3 Hard-Core Local Leakage Resilience

Everywhere-short shares may reveal partial information about a short secret, so conventional privacy is generally too strong for our setting. We instead protect a prescribed hard-core function of a secret sampled from a designated distribution, even given the shares of a private set and prescribed local leakage from all shares. Our model follows the local-leakage setting, in which each share is processed by its own leakage function, studied for non-malleable secret sharing in [49] and for linear secret sharing in [50].

Definition 8 (Local leakage function) Let $\mathsf{Sharing}$ have share spaces $(S_i)_{i\in[n]}$. A *local leakage function* is a tuple $L = (L_i)_{i\in[n]}$, where $L_i : S_i \to D_i$, acting component-wise as $L((s_i)_{i\in[n]}) := (L_i(s_i))_{i\in[n]}$. We write $\mathcal{L}$ for a family of such functions and $\mathcal{D}_{\mathcal{L}}$ for a distribution over $\mathcal{L}$.

The public description of L may include linear maps and fresh randomness. In our analysis, $L_i(\mathbf{s}_i) := \mathbf{L}_i\mathbf{s}_i + \mathbf{r}_i$, where $\mathbf{L}_i$ is public and $\mathbf{r}_i$ is fresh noise, possibly across several queries, as when a threshold signing protocol releases noisy linear functions of the shares. We omit the publicly known L from the adversary's displayed input below. We now define the experiment underlying our privacy notion.

Definition 9 (HC-Leak experiment) Let $\mathsf{Sharing}$ be a secret sharing scheme over n parties with secret space S. Let $\mathcal{D}_S$ be a distribution over S, $\mathcal{D}_{\mathcal{L}}$ a distribution over a family $\mathcal{L}$ of local leakage functions, and $\mathcal{D}_{\mathcal{F}}$ a distribution over a family $\mathcal{F}$ of functions $f : S \to Z$, where Z is finite. Fix pp in the support of $\mathsf{Setup}(1^\kappa)$ and a private set $U \subseteq [n]$ with $|U| \le k_p$. The experiment $\mathsf{HC\text{-}Leak}^{\mathsf{pp},U}_{\mathcal{A},\mathsf{Sharing}}(\kappa)$ proceeds as follows:

1. Sample $f \leftarrow \mathcal{D}_{\mathcal{F}}$, $L \leftarrow \mathcal{D}_{\mathcal{L}}$, and $s \leftarrow \mathcal{D}_S$, and generate $(s_i)_{i\in[n]} \leftarrow \mathsf{Deal}_{\mathsf{pp}}(s)$.
2. Sample $z_0 \xleftarrow{\$} Z$ and $b \xleftarrow{\$} \{0,1\}$, and set $z_1 := f(s)$.
3. Give $\big(\mathsf{pp}, f, (s_i)_{i\in U}, L((s_i)_{i\in[n]}), z_b\big)$ to $\mathcal{A}$, which outputs $b' \in \{0,1\}$.

The advantage of $\mathcal{A}$ is $\mathsf{Adv}^{\mathsf{pp},U}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) := 2\left|\Pr[b' = b] - 1/2\right|$.

In our application, the hard-core function is the systematic MLWE target $f_{\boldsymbol{A}}(\boldsymbol{s}) := \boldsymbol{A}\boldsymbol{s} \bmod q$, where $\boldsymbol{A} := \begin{bmatrix}\boldsymbol{I}_p & \boldsymbol{A}'\end{bmatrix} \in \mathcal{R}_q^{p\times\ell}$ and $\boldsymbol{A}'$ is sampled as part of the function description. Thus, $Z = \mathcal{R}_q^p$, and the experiment asks whether $f_{\boldsymbol{A}}(\boldsymbol{s})$ is indistinguishable from an independent uniform element of $\mathcal{R}_q^p$. The following definition strengthens the leakage-resilience notion of [39, Definition 5].

Definition 10 (Hard-core function for local leakage resilience) Let $\mathsf{Sharing} = (\mathsf{Setup}, \mathsf{Deal}, \mathsf{Rec})$ be an ϵ_c-correct secret sharing scheme with secret space S and privacy

threshold k_p. Let $\mathcal{D}_S$, $\mathcal{D}_{\mathcal{L}}$, and $\mathcal{D}_{\mathcal{F}}$ be as in Definition 9, let $\delta_p \in [0,1]$, and let ϵ_p assign a bound $\epsilon_p(\kappa, \mathcal{A}, U) \in [0,1]$ to every security parameter κ, PPT adversary $\mathcal{A}$, and coalition U. We say that $\mathcal{F}$ is a (δ_p, ϵ_p)-private hard-core function family for local leakage resilience with respect to $\mathsf{Sharing}$, $\mathcal{D}_S$, $\mathcal{D}_{\mathcal{L}}$, and $\mathcal{D}_{\mathcal{F}}$ if, for every κ,

$$\Pr_{\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)} \left[\forall \text{ PPT } \mathcal{A},\ \forall U \subseteq [n], |U| \le k_p : \mathsf{Adv}^{\mathsf{pp},U}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \epsilon_p(\kappa, \mathcal{A}, U) \right] \ge 1 - \delta_p.$$

In words, except with probability at most δ_p over $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)$, every PPT adversary has advantage at most $\epsilon_p(\kappa, \mathcal{A}, U)$ in the $\mathsf{HC\text{-}Leak}$ experiment for every coalition $U \subseteq [n]$ of size at most k_p. If, for every PPT adversary $\mathcal{A}$ and every such coalition U, the function $\epsilon_p(\kappa, \mathcal{A}, U)$ is negligible in κ, we say that $\mathcal{F}$ is δ_p-private.

Thus, δ_p bounds the probability of exceptional setup outputs, which are determined before the adversary and the coalition are chosen; outside them, every PPT adversary has advantage at most ϵ_p against every coalition of at most k_p parties. Unlike conventional privacy, this notion protects a prescribed hard-core function of a secret sampled from $\mathcal{D}_S$, rather than requiring indistinguishability of the complete views induced by arbitrary fixed secrets.

3 The SEC Construction

This section explains our ramp secret sharing scheme SEC along with the construction of its distribution matrix. Correctness, everywhere-shortness, and linearity are established in Section 4, and privacy in Section 5.

3.1 Overview

Our construction follows the coding-theoretic approach to secret sharing initiated by McEliece and Sarwate [5] and developed through linear secret sharing, span programs, and black-box secret sharing [6, 8, 11–14, 51, 52]. Unlike conventional finite-field schemes, SEC operates over a $\mathbb{Z}$-module and is targeted primarily for lattice-based distributed cryptography, where a common requirement is for both share coefficients and reconstruction coefficients to be bounded.

A designated dealer samples an $(n, m, c_{\exp}, \gamma, (1-\epsilon_{\exp})c_{\exp})$ lossless bipartite expander (Definition 1) with bi-adjacency matrix $\mathbf{H} \in \{0,1\}^{m \times n}$, viewed over $\mathbb{Z}$. It samples a signed permutation matrix $\mathbf{P} \in GL_m(\mathbb{Z})$, sets $\hat{\mathbf{G}} := (\mathbf{PH})^\top$, and constructs a secret-bearing column $\mathbf{g} \in \mathbb{Z}^n$ that supports reconstruction by sufficiently large sets. The distribution matrix is $\mathbf{G} := \begin{bmatrix} \mathbf{g} & \hat{\mathbf{G}} \end{bmatrix} \in \mathbb{Z}^{n \times k}$, where $k := m + 1$. The column $\mathbf{g}$ determines how the secret enters the shares, while the sparse columns of $\hat{\mathbf{G}}$ distribute short masking randomness and keep the shares short.

Let the secret be $\boldsymbol{s} \in \mathcal{R}^\ell$, where ℓ is the module rank and $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ has degree d, and let $\mathbf{s} \in \mathbb{Z}^D$, with $D := d\ell$, denote its coefficient embedding. To share $\mathbf{s}$, the dealer samples a short matrix $\mathbf{R} \in \mathbb{Z}^{(k-1)\times D}$ and computes $\mathbf{S} := \mathbf{G} \begin{bmatrix} \mathbf{s}^\top \\ \mathbf{R} \end{bmatrix} \in \mathbb{Z}^{n \times D}$. Party i receives the row $\mathbf{s}_i^\top := \mathbf{S}[i, \cdot]$. For a participant index list $T \subseteq [n]$, any $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ satisfying $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$ yields $\boldsymbol{\lambda}_T^\top \mathbf{S}_T = \mathbf{s}^\top$, so reconstruction by $T \subseteq [n]$ reduces to

finding such a vector with small ℓ_∞-norm. For every private set U, equivalently every $U \subseteq [n]$ with $|U| \leq k_p$, we additionally require $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$, which rules out exact cancellation of the masking term over the integers and modulo any integer. This structural condition does not by itself give privacy, which we prove computationally in Section 5.

3.2 Setup, Sharing, and Reconstruction

We first characterize reconstruction and give a deterministic procedure for finding a short reconstruction vector. We then formally define the setup, sharing, and reconstruction algorithms.

The following standard linear-reconstruction criterion underlies span-program, linear, and black-box secret sharing [11, 14, 51, 53].

Lemma 5 (Authorized reconstruction) *Let $\mathbf{G} \in \mathbb{Z}^{n\times k}$, $\mathbf{s} \in \mathbb{Z}^D$, $\mathbf{R} \in \mathbb{Z}^{(k-1)\times D}$, and define $\mathbf{S} := \mathbf{G}\begin{bmatrix}\mathbf{s}^\top \\ \mathbf{R}\end{bmatrix}$. For any $T \subseteq [n]$, if $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$, then there exists $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ such that $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$, and every such vector satisfies $(\boldsymbol{\lambda}_T^\top \mathbf{S}_T)^\top = \mathbf{s}$.*

Proof Suppose that $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$. By the definition of the integer row module, there exists $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ such that $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$. For every such vector, $\boldsymbol{\lambda}_T^\top \mathbf{S}_T = \boldsymbol{\lambda}_T^\top \mathbf{G}_T \begin{bmatrix}\mathbf{s}^\top \\ \mathbf{R}\end{bmatrix} = \mathbf{e}_1^\top \begin{bmatrix}\mathbf{s}^\top \\ \mathbf{R}\end{bmatrix} = \mathbf{s}^\top$. Transposing both sides gives $(\boldsymbol{\lambda}_T^\top \mathbf{S}_T)^\top = \mathbf{s}$. □

Conversely, if a share-independent $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ satisfies $(\boldsymbol{\lambda}_T^\top \mathbf{S}_T)^\top = \mathbf{s}$ for all $\mathbf{s}$ and $\mathbf{R}$, then $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$, since $\begin{bmatrix}\mathbf{s}^\top \\ \mathbf{R}\end{bmatrix}$ ranges over all of $\mathbb{Z}^{k\times D}$. Hence $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$ characterizes the sets that admit integer-linear reconstruction. The row Hermite normal form both tests this condition and extracts a reconstruction vector.

Lemma 6 (Row-HNF characterization of authorization) *Let $\mathbf{A} \in \mathbb{Z}^{a\times b}$ where $a, b \in \mathbb{Z}^+$, and let $\mathbf{U}_\mathbf{A} \in GL_a(\mathbb{Z})$ be such that $\mathbf{R}_\mathbf{A} := \mathbf{U}_\mathbf{A}\mathbf{A}$ is the row-HNF of $\mathbf{A}$. That is, for some $r \geq 0$, the nonzero rows of $\mathbf{R}_\mathbf{A}$ are its first r rows, the first nonzero entry $p_i > 0$ of row $i \in [r]$ lies in column j_i with $j_1 < \cdots < j_r$, and $0 \leq \mathbf{R}_\mathbf{A}[\iota, j_i] < p_i$ for all $\iota < i$. Then $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{A})$ if and only if $\mathbf{R}_\mathbf{A}[1,\cdot] = \mathbf{e}_1^\top$. In this case, $\boldsymbol{\lambda}_0 := (\mathbf{U}_\mathbf{A}[1,\cdot])^\top$ satisfies $\boldsymbol{\lambda}_0^\top \mathbf{A} = \mathbf{e}_1^\top$.*

We prove Lemma 6 in Appendix D; the proof uses only that unimodular row operations preserve $\mathrm{Row}_{\mathbb{Z}}$ and that the row-HNF pivots reduce the entries above them.

Although $\mathsf{Rec}_{\mathsf{pp}}$ ultimately computes only the vector-matrix product $(\boldsymbol{\lambda}_T^\top \mathbf{S}_T)^\top$, one must first select an integer vector $\boldsymbol{\lambda}_T$ satisfying $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$. By Lemma 6, the row-HNF determines whether such a vector exists and, when it does, yields the particular solution $\boldsymbol{\lambda}_{T,0} := (\mathbf{U}_{\mathbf{G}_T}[1,\cdot])^\top$. The transformation matrix $\mathbf{U}_{\mathbf{G}_T}$ need not be unique, and the resulting vector $\boldsymbol{\lambda}_{T,0}$ need not be short. Since adding any $\boldsymbol{\nu} \in K_T :=$

$\ker_{\mathbb{Z}}(\mathbf{G}_T^\top)$ preserves the reconstruction equation, all integer solutions form the affine coset $\boldsymbol{\lambda}_{T,0} + K_T$, and Algorithm 1 fixes a deterministic procedure for selecting a short representative of it.

Algorithm 1 Constructing a small reconstruction vector $\boldsymbol{\lambda}_T$ for $\mathbf{G}_T$.

Require: A submatrix $\mathbf{G}_T \in \mathbb{Z}^{|T|\times k}$; a fixed Lovász parameter $\delta_{\mathsf{LLL}} \in (1/4, 1)$; and fixed small greedy parameters $passes, window \in \mathbb{Z}^+$.
Ensure: A vector $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ satisfying $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$, or $\perp$ if $\mathbf{e}_1^\top \notin \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$.

Compute the row-HNF $\mathbf{R}_{\mathbf{G}_T}$ of $\mathbf{G}_T$ together with a $\mathbf{U}_{\mathbf{G}_T} \in GL_{|T|}(\mathbb{Z})$ such that $\mathbf{U}_{\mathbf{G}_T}\mathbf{G}_T = \mathbf{R}_{\mathbf{G}_T}$.
if $\mathbf{R}_{\mathbf{G}_T}[1,\cdot] \neq \mathbf{e}_1^\top$ **then**
  return $\perp$.
end if
Set $\boldsymbol{\lambda}_{T,0} := (\mathbf{U}_{\mathbf{G}_T}[1,\cdot])^\top$.
Let $\mathbf{B}_K \in \mathbb{Z}^{r\times|T|}$ be the row-HNF basis of $K_T := \ker_{\mathbb{Z}}(\mathbf{G}_T^\top)$, where $r := \mathrm{rank}(K_T)$.
if $r = 0$ **then**
  return $\boldsymbol{\lambda}_{T,0}$.
end if
Compute $\mathbf{B}'_K := \mathsf{LLL}_{\delta_{\mathsf{LLL}}}(\mathbf{B}_K)$ in exact rational arithmetic using the fixed reduction and tie-breaking conventions.
Compute $\boldsymbol{\lambda}_T \leftarrow \mathsf{BabaiGreedy}(\boldsymbol{\lambda}_{T,0}, \mathbf{B}'_K, passes, window)$ according to Algorithm 4.
if $\boldsymbol{\lambda}_T^\top \mathbf{G}_T \neq \mathbf{e}_1^\top$ **then**
  return $\perp$. ▷ exact certification
end if
return $\boldsymbol{\lambda}_T$.

Finding a representative of minimum ℓ_∞-norm is a closest-vector problem and is NP-hard in general [54, 55]; we therefore use exact-arithmetic LLL reduction and Babai decoding [56, 57], followed by a deterministic greedy search that directly targets the ℓ_∞-norm. Since $\mathsf{BabaiGreedy}$ (Algorithm 4) modifies $\boldsymbol{\lambda}_{T,0}$ only by vectors in K_T, its output remains in the affine coset $\boldsymbol{\lambda}_{T,0}+K_T$ and therefore preserves the reconstruction equation. The final equality test certifies this equation exactly. The resulting coefficient size is analyzed in Section 4.2.2 and evaluated empirically in Section 6.3. The detailed Babai and greedy procedures are given in Algorithm 4 in Appendix E.

We formally define SEC, an expander-code-based secret sharing scheme, as follows.

Definition 11 (Expander-code-based secret sharing) Fix $d, \ell, n, m, c_{\exp} \in \mathbb{Z}^+$ with $n > m$, let $D := d\ell$ and $k := m + 1$, and fix $\gamma \in (0, 1]$, $\epsilon_{\exp} \in (0, 1/2)$, $\sigma_{\mathbf{s}}, \sigma_{\mathbf{R}} \in \mathbb{R}^+$, a privacy threshold $k_p \in [n-1]$ satisfying $k_p \leq \lceil 2(1-\epsilon_{\exp})\lfloor \gamma n \rfloor \rceil - 1$, a target reconstruction-coverage error $\epsilon_c \in [0, 0.01]$, and the fixed public parameters used by Algorithm 1. Let the secret space be $S := \mathbb{Z}^D$, let each party $i \in [n]$ have share space $S_i := \mathbb{Z}^D$, and let $\mathcal{D}_S := \mathcal{D}_{\sigma_{\mathbf{s}}}^D$ be the designated secret distribution for the everywhere-shortness and privacy analyses. The expander-code-based secret sharing scheme $\mathsf{SEC} = (\mathsf{Setup}, \mathsf{Deal}, \mathsf{Rec})$ consists of the following algorithms.

1. $\mathsf{Setup}(1^\kappa)$. Sample a lossless bipartite expander graph $(n, m, c_{\exp}, \gamma, (1-\epsilon_{\exp})c_{\exp})$, and let $\mathbf{H} \in \{0,1\}^{m\times n}$ be its bi-adjacency matrix, viewed over $\mathbb{Z}$. Sample a signed permutation matrix $\mathbf{P} \in GL_m(\mathbb{Z})$, set $\hat{\mathbf{G}} := (\mathbf{PH})^\top \in \mathbb{Z}^{n\times m}$, construct $\mathbf{g} \in \mathbb{Z}^n$ according to Section 3.3, and define $\mathbf{G} := \begin{bmatrix}\mathbf{g} & \hat{\mathbf{G}}\end{bmatrix} \in \mathbb{Z}^{n\times k}$. By the restrictions on $\epsilon_{\exp}$ and k_p above, $\hat{\mathbf{G}}$ satisfies $\mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U) = \mathbb{Z}^{|U|}$, and hence $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$, for every $U \subseteq [n]$ with $|U| \le k_p$. Let k_c be the smallest integer in $[k_p+1, n]$ such that, for every $t \in [k_c, n]$, $\Pr_{T \xleftarrow{\$} \binom{[n]}{t}}\left[\mathbf{e}_1^\top \notin \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)\right] \le \epsilon_c$. Define $\Gamma_0(\mathsf{pp}) := \{U \subseteq [n] : |U| \le k_p\}$ and $\Gamma_1(\mathsf{pp}) := \{T \subseteq [n] : |T| \ge k_c \wedge \mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)\}$, and output the public parameters pp, including $\mathbf{G}$, k_c, and the public sampling and reconstruction parameters.
2. $\mathsf{Deal}_{\mathsf{pp}}(\mathbf{s})$. On input $\mathbf{s} \in \mathbb{Z}^D$, sample a masking matrix $\mathbf{R} \in \mathbb{Z}^{(k-1)\times D}$ with independent entries from $\mathcal{D}_{\sigma_{\mathbf{R}}}$, compute
$$\mathbf{S} := \mathbf{G}\begin{bmatrix}\mathbf{s}^\top \\ \mathbf{R}\end{bmatrix}, \tag{3}$$
and output the shares $(\mathbf{s}_i)_{i\in[n]}$, where $\mathbf{s}_i := (\mathbf{S}[i,\cdot])^\top \in \mathbb{Z}^D$.
3. $\mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{s}_i)_{i\in T})$. If $|T| < k_c$, output $\perp$. Otherwise, let $\mathbf{S}_T \in \mathbb{Z}^{|T|\times D}$ be the matrix whose rows are the transposes of the shares indexed by T, ordered according to the indices in T, and compute $\boldsymbol{\lambda}_T$ by running Algorithm 1 on input $\mathbf{G}_T$. If the algorithm returns $\perp$, output $\perp$; otherwise, output $(\boldsymbol{\lambda}_T^\top \mathbf{S}_T)^\top$.

By Lemmas 5 and 6, $\mathsf{Rec}_{\mathsf{pp}}$ exactly recovers the secret for every $T \in \Gamma_1(\mathsf{pp})$ (Lemma 11). The private-set condition asserted in Definition 11 follows from the expander structure under the restrictions on $\epsilon_{\exp}$ and k_p (Lemmas 4 and 7).

Remark 4 (Deterministic selection of reconstruction vectors) Although the reconstruction equation may admit multiple integer solutions, Algorithm 1 fixes a deterministic output for each $\mathbf{G}_T$ and its stated parameters. This follows from the unique row-HNF basis of K_T, the fixed exact-arithmetic LLL and Babai conventions, and the deterministic greedy refinement; in particular, the final output is independent of the transformation matrix used to obtain $\boldsymbol{\lambda}_{T,0}$. The decoding and tie-breaking rules, together with the determinism justification, are given in Algorithm 4 in Appendix E. Once the resulting public vector $\boldsymbol{\lambda}_T$ is fixed, reconstruction is the vector-matrix multiplication $\mathbf{S}_T \mapsto (\boldsymbol{\lambda}_T^\top \mathbf{S}_T)^\top$, and is therefore efficient and deterministic.

3.3 Constructing the Distribution Matrix

We describe how Setup constructs the distribution matrix $\mathbf{G} = \begin{bmatrix}\mathbf{g} & \hat{\mathbf{G}}\end{bmatrix}$. The expander-derived masking block $\hat{\mathbf{G}}$ alone guarantees the private-set condition, allowing the secret-bearing column $\mathbf{g}$ to be chosen to promote reconstruction. We construct $\mathbf{g}$ by planting pairwise disjoint certified supports, thereby obtaining a deterministic reconstruction threshold while keeping its entries small.

The expander structure ensures that every restriction of $\hat{\mathbf{G}}$ to at most k_p rows is integrally surjective.

Lemma 7 (Integral surjectivity of row restrictions to private sets) *Let* $\hat{\mathbf{G}} = (\mathbf{PH})^\top \in \{-1,0,1\}^{n\times m}$, *where* $\mathbf{P} \in GL_m(\mathbb{Z})$ *is a signed permutation matrix and* $\mathbf{H} \in \{0,1\}^{m\times n}$ *is*

the bi-adjacency matrix of an $(n, m, c_{\exp}, \gamma, (1-\epsilon_{\exp})c_{\exp})$ *lossless bipartite expander with* $0 < \epsilon_{\exp} < 1/2$. *If* $k_p \le \lceil 2(1-\epsilon_{\exp})\lfloor \gamma n \rfloor \rceil - 1$, *then* $\mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U) = \mathbb{Z}^{|U|}$ *for every* $U \subseteq [n]$ *with* $|U| \le k_p$.

Proof The claim is immediate if $U = \emptyset$. Now let $\emptyset \neq U \subseteq [n]$ with $|U| \le k_p$. By Lemma 4, the restricted matrix $\mathbf{H}_{(U)} \in \{0,1\}^{m \times |U|}$ contains a $|U| \times |U|$ unimodular submatrix, and hence $\mathrm{Row}_{\mathbb{Z}}(\mathbf{H}_{(U)}) = \mathbb{Z}^{|U|}$. Since $\mathbf{P} \in GL_m(\mathbb{Z})$, left multiplication by $\mathbf{P}$ preserves the integer row module, so $\mathrm{Row}_{\mathbb{Z}}((\mathbf{PH})_{(U)}) = \mathbb{Z}^{|U|}$. Finally, $\hat{\mathbf{G}}_U = ((\mathbf{PH})_{(U)})^{\top}$, and therefore $\mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U) = \mathrm{Row}_{\mathbb{Z}}((\mathbf{PH})_{(U)}) = \mathbb{Z}^{|U|}$. □

Consequently, $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$ for every $\mathbf{g} \in \mathbb{Z}^n$ and every $U \subseteq [n]$ with $|U| \le k_p$. Thus, after fixing $\hat{\mathbf{G}}$, the construction of $\mathbf{g}$ can focus on promoting reconstruction without affecting the private-set condition required in Definition 11.

To motivate the construction of $\mathbf{g}$, write $\mathbf{H}' := \mathbf{PH}$, so that $\hat{\mathbf{G}}_T = (\mathbf{H}'_{(T)})^{\top}$. By Lemma 5, a set $T \subseteq [n]$ reconstructs whenever $\mathbf{e}_1^{\top} \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$. Since $\mathbf{G}_T = \left[(\mathbf{g})_T \ \hat{\mathbf{G}}_T\right]$, this condition can be rewritten as

$$\mathbf{e}_1^{\top} \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T) \quad \Longleftrightarrow \quad \exists \boldsymbol{\lambda}_T \in \ker_{\mathbb{Z}}(\mathbf{H}'_{(T)}) : \boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = 1. \tag{4}$$

If $\ker_{\mathbb{Z}}(\mathbf{H}'_{(T)}) \neq \{\mathbf{0}\}$, then dividing any nonzero vector in this kernel by the greatest common divisor of its entries gives a primitive vector $\boldsymbol{\lambda}_T$. The generalized Bézout identity then gives $(\mathbf{g})_T \in \mathbb{Z}^{|T|}$ satisfying $\boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = 1$. The remaining challenge is to construct a single global vector $\mathbf{g} \in \mathbb{Z}^n$ that satisfies this condition for many sets T, while keeping its entries small.

We address this challenge by planting pairwise disjoint supports that satisfy (4). This guarantees that every set of size at least a deterministic threshold k_c^{det} reconstructs. Smaller sets are not guaranteed to reconstruct, but Section 6.2.2 shows empirically that most sets of size at least the smaller threshold k_c of Definition 11 do.

We call $T \subseteq [n]$ a *certified support* if $\mathbf{e}_1^{\top} \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$. Certification is monotone: if T is certified, then every $S \supseteq T$ is also certified. Our construction plants pairwise disjoint certified supports so that every sufficiently large set contains at least one of them.

Lemma 8 (Deterministic reconstruction threshold) *If* $T_1, \ldots, T_b \subseteq [n]$ *are pairwise disjoint certified supports, then every* $S \subseteq [n]$ *with* $|S| \ge n - b + 1$ *is certified.*

Proof Let $R := [n] \setminus \bigcup_{i=1}^{b} T_i$. If $S \subseteq [n]$ contains none of $T_1, \ldots, T_b$, then it omits at least one element from each T_i, so $|S| \le |R| + \sum_{i=1}^{b}(|T_i| - 1) = n - b$. Therefore, every $S \subseteq [n]$ with $|S| \ge n - b + 1$ contains some T_i. Since certification is monotone, such an S is certified. □

As a result, planting b pairwise disjoint certified supports gives the deterministic reconstruction threshold $k_c^{\mathsf{det}} := n - b + 1$. The construction therefore seeks many

small pairwise disjoint sets T_i for which $\ker_{\mathbb{Z}}(\mathbf{H}'_{(T_i)}) \neq \{\mathbf{0}\}$, chooses a short primitive vector in each kernel, and selects $(\mathbf{g})_{T_i}$ using the generalized Bézout identity so that (4) holds. The following rank argument gives a deterministic baseline for the number of such supports.

Lemma 9 (Rank-based disjoint supports) *Let $\mathbf{H}' \in \mathbb{Z}^{m\times n}$, and let $r := \mathrm{rank}_{\mathbb{Q}}(\mathbf{H}')$. Then there exist pairwise disjoint sets $T_1, \ldots, T_b \subseteq [n]$ with $|T_i| = r+1$ and $\ker_{\mathbb{Z}}(\mathbf{H}'_{(T_i)}) \neq \{\mathbf{0}\}$ for every $i \in [b]$, such that the leftover set $R := [n] \setminus \bigcup_{i=1}^{b} T_i$ satisfies $|R| \leq r$.*

Indeed, any $r+1$ columns of $\mathbf{H}'$ are linearly dependent over $\mathbb{Q}$, so partitioning $[n]$ into $\lfloor n/(r+1)\rfloor$ blocks of size $r+1$, together with a leftover set of size at most r, proves the claim (Appendix D). Making each block certified as above gives $b = \lfloor n/(r+1)\rfloor$, and hence $k_c^{\mathsf{det}} \leq n - \lfloor n/(r+1)\rfloor + 1$. In practice, the construction searches for smaller kernel supports, thereby increasing b and lowering the deterministic reconstruction threshold. The complete construction is given in Algorithm 2; it prefers Bézout solutions in $\{-1, 0, 1\}^{|T_i|}$ when possible and assigns small values to the remaining coordinates of $\mathbf{g}$.

For every support $T \in \mathcal{T}$, the algorithm produces $\boldsymbol{\lambda}_T \in \ker_{\mathbb{Z}}(\mathbf{H}'_{(T)})$ and $(\mathbf{g})_T$ satisfying $\boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = 1$, so T is certified by (4). Every selected support is removed from Ω before the next iteration, so the supports are pairwise disjoint, and Lemma 8 guarantees reconstruction for every set of size at least $k_c^{\mathsf{det}} = n - b + 1$, where $b := |\mathcal{T}|$.

The branch $t = t_{\max}$ always admits a nonzero integer kernel vector whenever $|\Omega| \geq t_{\max}$. Indeed, $t_{\max} = \mathrm{rank}_{\mathbb{Q}}(\mathbf{H}') + 1$, so any $t_{\max}$ columns of $\mathbf{H}'$ are linearly dependent over $\mathbb{Q}$, and clearing denominators yields a nonzero integer kernel vector. Thus, the deterministic baseline of Lemma 9 is recovered even when the searches for smaller supports are unsuccessful.

The choices that keep $\mathbf{g}$ short do not affect certification, which depends only on the exact relations $\boldsymbol{\lambda}_T \in \ker_{\mathbb{Z}}(\mathbf{H}'_{(T)})$ and $\boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = 1$; the resulting value of $\|\mathbf{g}\|_\infty$ enters the individual-share bound of Section 4.2.1.

For fixed search parameters, the construction has polynomial bit complexity; the operation count is given in Appendix D. The existence of a certified support also ensures that the secret-bearing column is not contained in the integer column module generated by the masking columns.

Lemma 10 *If $\mathbf{e}_1^{\top} \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$ for some $T \subseteq [n]$, then $\mathbf{g} \notin \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}})$.*

Proof Suppose, toward a contradiction, that $\mathbf{g} \in \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}})$. Then there exists $\mathbf{a} \in \mathbb{Z}^{k-1}$ such that $\mathbf{g} = \hat{\mathbf{G}}\mathbf{a} = (\mathbf{H}')^{\top}\mathbf{a}$. Let $T \subseteq [n]$ satisfy $\mathbf{e}_1^{\top} \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$. By (4), there exists $\boldsymbol{\lambda}_T \in \ker_{\mathbb{Z}}(\mathbf{H}'_{(T)})$ satisfying $\boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = 1$. Restricting $\mathbf{g} = (\mathbf{H}')^{\top}\mathbf{a}$ to T gives $(\mathbf{g})_T = (\mathbf{H}'_{(T)})^{\top}\mathbf{a}$. Therefore, $\boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = (\mathbf{H}'_{(T)}\boldsymbol{\lambda}_T)^{\top}\mathbf{a} = 0$, contradicting $\boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = 1$. Hence $\mathbf{g} \notin \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}})$. □

Accordingly, whenever the construction produces at least one certified support, the resulting distribution matrix admits nontrivial secret reconstruction while retaining the private-set condition established in Lemma 7.

Algorithm 2 Constructing $\mathbf{g}$ from disjoint certified supports.

Require: A matrix $\mathbf{H}' = \mathbf{PH} \in \mathbb{Z}^{m\times n}$ with $n > m$; a target starting size $t_0 \in [t_{\max}]$, where $t_{\max} := \operatorname{rank}_{\mathbb{Q}}(\mathbf{H}') + 1$; and a fixed bound $maxTrials \in \mathbb{Z}^+$ on the number of sampled subsets per size.

Ensure: A vector $\mathbf{g} \in \mathbb{Z}^n$ and a family $\mathcal{T}$ of pairwise disjoint certified supports.

Initialize $\Omega := [n]$, $\mathbf{g} := \mathbf{0}$, and $\mathcal{T} := \emptyset$.
while $|\Omega| \geq t_0$ **do**
    Set *found* := **False**.
    for $t = t_0$ **to** $\min\{t_{\max}, |\Omega|\}$ **do**
        if $t < t_{\max}$ **then**
            Sample and examine up to *maxTrials* subsets $T \subseteq \Omega$ with $|T| = t$, using the fixed sampling and tie-breaking conventions.
            if some T satisfies $\ker_{\mathbb{Z}}(\mathbf{H}'_{(T)}) \neq \{\mathbf{0}\}$ **then**
                Record T, set *found* := **True**, and **break**.
            end if
        else
            Choose the lexicographically first $T \subseteq \Omega$ with $|T| = t$.
            Set *found* := **True** and **break**.
        end if
    end for
    if not *found* **then**
        break. ▷ no further certified support is found
    end if
    Compute a short primitive vector $\boldsymbol{\lambda}_T \in \ker_{\mathbb{Z}}(\mathbf{H}'_{(T)})$.
    Compute a small solution $(\mathbf{g})_T \in \mathbb{Z}^{|T|}$ to $\boldsymbol{\lambda}_T^{\top}(\mathbf{g})_T = 1$ using the generalized Bézout identity.
    Write $T = \{t_1, \ldots, t_{|T|}\}$ with $t_1 < \cdots < t_{|T|}$, and set $\mathbf{g}[t_a] = (\mathbf{g})_T[a]$ for every $a \in [|T|]$.
    Set $\mathcal{T} := \mathcal{T} \cup \{T\}$ and $\Omega := \Omega \setminus T$.
end while
Assign each remaining coordinate $\mathbf{g}[j]$, for $j \in \Omega$, a fixed value in $\{-1, 0, 1\}$, taking $\mathbf{g}[j] = 0$ by default.
return $(\mathbf{g}, \mathcal{T})$.

4 Correctness, Everywhere-Shortness, and Linearity

This section establishes the functional properties of SEC. Section 4.1 shows that SEC is perfectly correct in the sense of Definition 6 and analyzes the reconstruction coverage measured by ϵ_c. Section 4.2 shows that SEC is everywhere-short in the sense of Definition 7 by bounding the shares and the reconstruction coefficients separately; since each party holds a single share, the expansion factor is 1. Finally, Section 4.3 proves the linearity of reconstruction and shows that public linear operations may be applied locally to the shares before reconstruction.

4.1 Correctness

4.1.1 Exact Reconstruction

The correctness analysis of SEC separates exact reconstruction from reconstruction coverage: every authorized set reconstructs exactly, while the reconstruction error ϵ_c measures only how often a sampled set fails the integer row-span criterion.

Lemma 11 (Exact reconstruction for authorized sets) *Fix public parameters* pp *generated by* $\mathsf{Setup}(1^\kappa)$. *For every* $T \subseteq [n]$ *satisfying* $|T| \geq k_c$ *and* $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$, *every secret* $\mathbf{s} \in \mathbb{Z}^D$, *and every* $(\mathbf{s}_i)_{i\in[n]}$ *in the support of* $\mathsf{Deal}_{\mathsf{pp}}(\mathbf{s})$, *we have* $\mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{s}_i)_{i\in T}) = \mathbf{s}$.

Proof Since $|T| \geq k_c$, $\mathsf{Rec}_{\mathsf{pp}}$ does not reject T based on its cardinality, and Algorithm 1 returns a vector $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ satisfying $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$. By Lemma 5, we have $(\boldsymbol{\lambda}_T^\top \mathbf{S}_T)^\top = \mathbf{s}$, and hence $\mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{s}_i)_{i\in T}) = \mathbf{s}$. □

For a fixed instantiation of $\mathbf{G}$, the reconstruction threshold k_c is selected so that, for every $t \geq k_c$, a uniformly sampled set $T \subseteq [n]$ of size t satisfies $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$ except with probability at most ϵ_c, and we require $\epsilon_c \leq 0.01$; for every certified set of size at least k_c, reconstruction is exact by Lemma 11. The planted certified supports of Algorithm 2 additionally guarantee reconstruction, with zero failure probability, for every set of size at least $k_c^{\mathsf{det}} = n - b + 1$ (Lemma 8). Accordingly, the correctness error of Definition 6 is zero for authorized sets, while throughout the remainder of the paper ϵ_c denotes the reconstruction-coverage error defined above.

Both k_c and k_c^{det} depend on the instantiated generator matrix $\mathbf{G}$. We derive analytical lower bounds on the certification probability induced by the planted certified supports in Section 4.1.2, and report the instantiated reconstruction thresholds in Section 6.2.2.

4.1.2 Reconstruction-Coverage Analysis

Reconstruction coverage measures the probability that a set sampled uniformly from all subsets of $[n]$ of a fixed size satisfies the integer row-span criterion.

Fix $\mathbf{G}$ and pairwise disjoint certified supports $T_1, \ldots, T_b \subseteq [n]$, and let $S \xleftarrow{\$} \binom{[n]}{s}$. Since every superset of a certified support is certified, $\Pr[\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_S)] \geq \Pr[\bigcup_{i=1}^{b}(T_i \subseteq S)]$. The bound improves if each T_i is replaced by the support $W_i := \mathrm{supp}(\boldsymbol{\lambda}_{T_i}) \subseteq T_i$ of its planted reconstruction vector $\boldsymbol{\lambda}_{T_i} \in \ker_{\mathbb{Z}}(\mathbf{H}'_{(T_i)})$ satisfying $\boldsymbol{\lambda}_{T_i}^\top (\mathbf{g})_{T_i} = 1$. Indeed, restricting $\boldsymbol{\lambda}_{T_i}$ to its support shows that W_i is itself certified, so $W_i \subseteq S$ implies certification. Here and below, $\binom{a}{c} := 0$ whenever $c < 0$ or $c > a$.

Corollary 2 (Coverage from planted supports) *Let* $t_i := |T_i|$. *For* $S \xleftarrow{\$} \binom{[n]}{s}$,

$$\Pr\left[\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_S)\right] \geq \frac{1}{\binom{n}{s}} \sum_{\emptyset \neq J \subseteq [b]} (-1)^{|J|+1} \binom{n - \sum_{i\in J} t_i}{s - \sum_{i\in J} t_i} \tag{5}$$

$$\geq \sum_{i=1}^{b} \frac{\binom{n-t_i}{s-t_i}}{\binom{n}{s}} - \sum_{1\leq i<j\leq b} \frac{\binom{n-t_i-t_j}{s-t_i-t_j}}{\binom{n}{s}}. \tag{6}$$

Corollary 3 (Coverage from witness supports) *The bounds (5) and (6) remain valid with t_i replaced by $w_i := |W_i| \leq t_i$.*

The right-hand side of (5) is the exact probability of containing at least one planted support, by inclusion-exclusion over disjoint sets, and (6) is the second-order Bonferroni bound; the witness version and the heuristic design rules $k_c \approx (\ln(1/\epsilon_c)/b)^{1/t} n$ and $k_c \approx (\ln(1/\epsilon_c)/b)^{1/w} n$ derived from them are given in Appendix D. Containing a planted or witness support is sufficient but not necessary for certification, so the certification probability of an instantiated matrix may be substantially larger than these bounds. The empirical evaluation in Section 6.2.2 therefore computes k_c exactly when feasible and otherwise estimates it from the observed certification rate.

4.2 Everywhere-Shortness

4.2.1 Share Shortness

We first bound the individual shares produced by $\mathsf{Deal}_{\mathsf{pp}}$. By (3), each share $\mathbf{s}_i$ is a linear function of the secret $\mathbf{s}$ and the masking matrix $\mathbf{R}$:

$$\begin{aligned} \mathbf{s}_i^\top &= \mathbf{G}[i,\cdot] \begin{bmatrix} \mathbf{s}^\top \\ \mathbf{R} \end{bmatrix} = \begin{bmatrix} \mathbf{g}[i] & \hat{\mathbf{G}}[i,\cdot] \end{bmatrix} \begin{bmatrix} \mathbf{s}^\top \\ \mathbf{R} \end{bmatrix} \\ &= \mathbf{g}[i]\mathbf{s}^\top + \hat{\mathbf{G}}[i,\cdot]\mathbf{R}. \end{aligned} \tag{7}$$

Lemma 12 (Boundedness of the shares) *Let $\beta_{\mathbf{s}}, \beta_{\mathbf{R}} \in \mathbb{R}^+$, and suppose that $\|\mathbf{s}\|_\infty \leq \beta_{\mathbf{s}}$ and $\|\mathbf{R}\|_{\max} \leq \beta_{\mathbf{R}}$. Then, for every $i \in [n]$, $\|\mathbf{s}_i\|_\infty \leq |\mathbf{g}[i]|\beta_{\mathbf{s}} + \|\hat{\mathbf{G}}[i,\cdot]\|_1 \beta_{\mathbf{R}}$. In particular, if $\hat{\mathbf{G}} = (\mathbf{PH})^\top$, where $\mathbf{P}$ is a signed permutation matrix and $\mathbf{H}$ is the bi-adjacency matrix of a $c_{\exp}$-left-regular bipartite lossless expander, then $\|\mathbf{s}_i\|_\infty \leq |\mathbf{g}[i]|\beta_{\mathbf{s}} + c_{\exp}\beta_{\mathbf{R}}$. If, moreover, $\|\mathbf{g}\|_\infty \leq \beta_{\mathbf{g}}$ for some $\beta_{\mathbf{g}} \in \mathbb{R}^+$, then $\|\mathbf{s}_i\|_\infty \leq \beta_{\mathbf{g}}\beta_{\mathbf{s}} + c_{\exp}\beta_{\mathbf{R}}$.*

Proof Coordinate-wise, (7) gives $|\mathbf{s}_i[j]| \leq |\mathbf{g}[i]||\mathbf{s}[j]| + \sum_t |\hat{\mathbf{G}}[i,t]||\mathbf{R}[t,j]| \leq |\mathbf{g}[i]|\beta_{\mathbf{s}} + \|\hat{\mathbf{G}}[i,\cdot]\|_1\beta_{\mathbf{R}}$. For $\hat{\mathbf{G}} = (\mathbf{PH})^\top$, $\|\hat{\mathbf{G}}[i,\cdot]\|_1 = \|\mathbf{H}[\cdot,i]\|_1 = c_{\exp}$ by left-regularity, and $|\mathbf{g}[i]| \leq \beta_{\mathbf{g}}$ gives the last claim. □

Under the designated secret distribution $\mathcal{D}_S = \mathcal{D}_{\sigma_{\mathbf{s}}}^D$ and the entrywise discrete Gaussian sampling of $\mathbf{R}$, we choose $\beta_{\mathbf{s}}$ and $\beta_{\mathbf{R}}$ so that the event $\|\mathbf{s}\|_\infty \leq \beta_{\mathbf{s}}$ and $\|\mathbf{R}\|_{\max} \leq \beta_{\mathbf{R}}$ holds except with probability at most ϵ_{sh}. Consequently, if $\|\mathbf{g}\|_\infty \leq \beta_{\mathbf{g}}$, then $\max_{i\in[n]} \|\mathbf{s}_i\|_\infty \leq \beta_{\mathbf{g}}\beta_{\mathbf{s}} + c_{\exp}\beta_{\mathbf{R}}$ except with probability at most ϵ_{sh}. Thus, SEC satisfies the individual-share shortness requirement of Definition 7 with share bound $\beta_{\mathsf{sh}} := \beta_{\mathbf{g}}\beta_{\mathbf{s}} + c_{\exp}\beta_{\mathbf{R}}$ and failure probability ϵ_{sh}. The tail parameters in Section 6.4.1 choose $\beta_{\mathbf{s}}$ and $\beta_{\mathbf{R}}$ so that $\epsilon_{\mathsf{sh}} \leq 2 \cdot 2^{-\kappa}$.

4.2.2 Reconstruction-Coefficient Bounds

We next bound the reconstruction coefficients produced by Algorithm 1. Fix an authorized set $T \in \Gamma_1(\mathsf{pp})$, and write $\mathbf{A}_T := \mathbf{G}_T^\top$, a local abbreviation rather than a row restriction. The reconstruction equation $\boldsymbol{\lambda}^\top \mathbf{G}_T = \mathbf{e}_1^\top$ is equivalent to $\mathbf{A}_T \boldsymbol{\lambda} = \mathbf{e}_1$. We first decompose every real solution into the minimum-ℓ_2-norm solution and a real-kernel component, and then bound the kernel component of the integer vector returned by Algorithm 1.

Affine Reconstruction Cosets

Lemma 13 (Affine coset of reconstruction vectors) *Let $T \in \Gamma_1(\mathsf{pp})$, let $\mathbf{A}_T := \mathbf{G}_T^\top$, and let $\boldsymbol{\lambda}_{T,0} \in \mathbb{Z}^{|T|}$ satisfy $\mathbf{A}_T \boldsymbol{\lambda}_{T,0} = \mathbf{e}_1$. Then the set of integer reconstruction vectors for T is the affine lattice coset $\mathcal{L}_T := \boldsymbol{\lambda}_{T,0} + K_T$, where $K_T := \ker_{\mathbb{Z}}(\mathbf{A}_T)$.*

Proof For every $\boldsymbol{\nu} \in K_T$, $\mathbf{A}_T(\boldsymbol{\lambda}_{T,0} + \boldsymbol{\nu}) = \mathbf{e}_1$. Conversely, every $\boldsymbol{\lambda} \in \mathbb{Z}^{|T|}$ with $\mathbf{A}_T \boldsymbol{\lambda} = \mathbf{e}_1$ satisfies $\mathbf{A}_T(\boldsymbol{\lambda} - \boldsymbol{\lambda}_{T,0}) = \mathbf{0}$, so $\boldsymbol{\lambda} - \boldsymbol{\lambda}_{T,0} \in K_T$. □

Although the representative $\boldsymbol{\lambda}_{T,0}$ obtained from a transformation matrix in the row-HNF computation need not be unique, the coset $\mathcal{L}_T$ depends only on $\mathbf{G}_T$.

Orthogonal Decomposition of Reconstruction Vectors

We next identify the unavoidable real least-norm component of every reconstruction vector. Let $r_T := \operatorname{rank}_{\mathbb{R}}(\mathbf{A}_T) = \operatorname{rank}_{\mathbb{R}}(\mathbf{G}_T)$, and let $d_T := \dim(\ker_{\mathbb{R}}(\mathbf{A}_T)) = |T| - r_T$. We write $\sigma^+_{\min}(\mathbf{G}_T)$ for the smallest positive singular value of $\mathbf{G}_T$.

Lemma 14 (Minimum-norm real reconstruction vector) *Let $T \in \Gamma_1(\mathsf{pp})$, let $\mathbf{A}_T := \mathbf{G}_T^\top$, and define $\boldsymbol{\lambda}_{T,\mathbb{R}} := \mathbf{A}_T^+ \mathbf{e}_1$, where $(\cdot)^+$ denotes the Moore-Penrose pseudoinverse. Then $\boldsymbol{\lambda}_{T,\mathbb{R}}$ is the unique solution of $\mathbf{A}_T \boldsymbol{\lambda} = \mathbf{e}_1$ having minimum Euclidean norm. Moreover, $\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 = \mathbf{e}_1^\top (\mathbf{A}_T \mathbf{A}_T^\top)^+ \mathbf{e}_1 = \mathbf{e}_1^\top (\mathbf{G}_T^\top \mathbf{G}_T)^+ \mathbf{e}_1$, and $\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2 \le \frac{1}{\sigma^+_{\min}(\mathbf{G}_T)}$.*

Proof Since $T \in \Gamma_1(\mathsf{pp})$, we have $\mathbf{e}_1 \in \operatorname{Col}_{\mathbb{R}}(\mathbf{A}_T)$, so the system $\mathbf{A}_T \boldsymbol{\lambda} = \mathbf{e}_1$ is consistent. The Moore-Penrose pseudoinverse gives its unique minimum-Euclidean-norm solution, and hence $\boldsymbol{\lambda}_{T,\mathbb{R}} = \mathbf{A}_T^+ \mathbf{e}_1$. Using the pseudoinverse identity $(\mathbf{A}_T^+)^\top \mathbf{A}_T^+ = (\mathbf{A}_T \mathbf{A}_T^\top)^+$, we obtain

$$\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 = (\mathbf{A}_T^+ \mathbf{e}_1)^\top (\mathbf{A}_T^+ \mathbf{e}_1) = \mathbf{e}_1^\top (\mathbf{A}_T^+)^\top \mathbf{A}_T^+ \mathbf{e}_1 = \mathbf{e}_1^\top (\mathbf{A}_T \mathbf{A}_T^\top)^+ \mathbf{e}_1 = \mathbf{e}_1^\top (\mathbf{G}_T^\top \mathbf{G}_T)^+ \mathbf{e}_1.$$

Furthermore, $\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2 \le \|\mathbf{A}_T^+\|_2 \|\mathbf{e}_1\|_2 = \frac{1}{\sigma^+_{\min}(\mathbf{A}_T)} = \frac{1}{\sigma^+_{\min}(\mathbf{G}_T)}$, where we used $\|\mathbf{e}_1\|_2 = 1$ and the fact that a matrix and its transpose have the same positive singular values. □

The following decomposition separates the kernel component of a reconstruction vector from the unavoidable minimum-ℓ_2-norm real component.

Lemma 15 (Orthogonal decomposition of reconstruction vectors) *Let $T \in \Gamma_1(\mathsf{pp})$. For every $\boldsymbol{\lambda} \in \mathcal{L}_T$, $\|\boldsymbol{\lambda}\|_2^2 = \|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 + \|\boldsymbol{\lambda} - \boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2$. In particular, fix any $\boldsymbol{\lambda}_{T,0} \in \mathcal{L}_T$, define*

$\mathbf{t}_T := \boldsymbol{\lambda}_{T,\mathbb{R}} - \boldsymbol{\lambda}_{T,0}$, *and write* $\boldsymbol{\lambda} = \boldsymbol{\lambda}_{T,0} + \boldsymbol{\nu}$ *for some* $\boldsymbol{\nu} \in K_T$. *Then* $\|\boldsymbol{\lambda}_{T,0} + \boldsymbol{\nu}\|_2^2 = \|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 + \|\boldsymbol{\nu} - \mathbf{t}_T\|_2^2$.

Proof The minimum-Euclidean-norm solution satisfies $\boldsymbol{\lambda}_{T,\mathbb{R}} \in \mathrm{Row}_{\mathbb{R}}(\mathbf{A}_T) = \mathrm{Col}_{\mathbb{R}}(\mathbf{A}_T^\top) = (\ker_{\mathbb{R}}(\mathbf{A}_T))^\perp$. Since $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}_{T,\mathbb{R}}$ satisfy $\mathbf{A}_T\boldsymbol{\lambda} = \mathbf{e}_1$ and $\mathbf{A}_T\boldsymbol{\lambda}_{T,\mathbb{R}} = \mathbf{e}_1$, respectively, their difference lies in $\ker_{\mathbb{R}}(\mathbf{A}_T)$. Hence $\boldsymbol{\lambda}_{T,\mathbb{R}}$ and $\boldsymbol{\lambda} - \boldsymbol{\lambda}_{T,\mathbb{R}}$ are orthogonal, and the first claim follows from the Pythagorean identity. For the second claim, observe that $\boldsymbol{\lambda} - \boldsymbol{\lambda}_{T,\mathbb{R}} = \boldsymbol{\lambda}_{T,0} + \boldsymbol{\nu} - \boldsymbol{\lambda}_{T,\mathbb{R}} = \boldsymbol{\nu} - \mathbf{t}_T$, and substitute this equality into the first identity. □

Replacing $\boldsymbol{\lambda}_{T,0}$ by another element of $\mathcal{L}_T$ translates $\mathbf{t}_T$ by an element of K_T, so the lattice-decoding problem is independent of the HNF transformation matrix used to obtain the initial representative. By Lemma 15, minimizing $\|\boldsymbol{\lambda}_{T,0} + \boldsymbol{\nu}\|_2^2$ over $\boldsymbol{\nu} \in K_T$ is equivalent to the closest-vector problem $\min_{\boldsymbol{\nu} \in K_T} \|\boldsymbol{\nu} - \mathbf{t}_T\|_2^2$, since the remaining term $\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2$ is fixed.

Basis-Dependent Babai Certificate

We write $B^*_{\boldsymbol{\lambda}}(T) := \min_{\boldsymbol{\lambda} \in \mathcal{L}_T} \|\boldsymbol{\lambda}\|_\infty$ for the optimal reconstruction-coefficient size and $B_{\boldsymbol{\lambda}}(T) := \|\boldsymbol{\lambda}_T\|_\infty$ for the value achieved by Algorithm 1. Since $\boldsymbol{\lambda}_T \in \mathcal{L}_T$, we have $B^*_{\boldsymbol{\lambda}}(T) \le B_{\boldsymbol{\lambda}}(T)$. Computing $B^*_{\boldsymbol{\lambda}}(T)$ is an ℓ_∞-closest-vector problem in K_T, which is hard in general. The algorithm instead applies Babai's nearest-plane procedure to an LLL-reduced basis of K_T, followed by a greedy refinement that accepts only updates strictly decreasing the ℓ_∞-norm. The following lemma gives a Euclidean certificate for the resulting vector.

Lemma 16 (Certified coefficient bound) *Let* $T \in \Gamma_1(\mathsf{pp})$, *let* $d_T := \dim(\ker_{\mathbb{R}}(\mathbf{G}_T^\top))$, *and let* $\sigma^+_{\min}(\mathbf{G}_T)$ *denote the smallest positive singular value of* $\mathbf{G}_T$. *Let* $\mathbf{B}'_K \in \mathbb{Z}^{d_T \times |T|}$, *with rows* $\mathbf{b}_1, \ldots, \mathbf{b}_{d_T}$, *be the LLL-reduced basis of* $K_T := \ker_{\mathbb{Z}}(\mathbf{G}_T^\top)$ *used by Algorithm 1, and let* $\mathbf{b}^*_1, \ldots, \mathbf{b}^*_{d_T}$ *be its Gram-Schmidt vectors. If* $\boldsymbol{\lambda}_T$ *is the output of Algorithm 1, then*

$$B_{\boldsymbol{\lambda}}(T) := \|\boldsymbol{\lambda}_T\|_\infty \le \left(\frac{1}{\left(\sigma^+_{\min}(\mathbf{G}_T)\right)^2} + \frac{1}{4} \sum_{i=1}^{d_T} \|\mathbf{b}^*_i\|_2^2 \right)^{1/2}. \tag{8}$$

When $d_T = 0$, *the sum is interpreted as zero.*

Proof If $d_T = 0$, then $K_T = \{\mathbf{0}\}$, so $\mathcal{L}_T$ contains a unique vector. This vector equals $\boldsymbol{\lambda}_{T,\mathbb{R}}$, since $\ker_{\mathbb{R}}(\mathbf{A}_T) = \{\mathbf{0}\}$ makes the real solution unique, and hence $\|\boldsymbol{\lambda}_T\|_\infty \le \|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2 \le \frac{1}{\sigma^+_{\min}(\mathbf{G}_T)}$ by Lemma 14. This is (8) with an empty Gram-Schmidt sum.

Suppose now that $d_T > 0$. Let $\boldsymbol{\lambda}_{\mathrm{dec}}$ be the output of Stage 1 of Algorithm 4, and let $\boldsymbol{\nu}_{\mathsf{Babai}} := \boldsymbol{\lambda}_{\mathrm{dec}} - \boldsymbol{\lambda}_{T,0} \in K_T$. By Lemma 15 and $\|\cdot\|_\infty \le \|\cdot\|_2$,

$$\|\boldsymbol{\lambda}_{\mathrm{dec}}\|_\infty \le \|\boldsymbol{\lambda}_{\mathrm{dec}}\|_2 \le \left(\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 + \|\boldsymbol{\nu}_{\mathsf{Babai}} - \mathbf{t}_T\|_2^2 \right)^{1/2}. \tag{9}$$

Since $\mathbf{A}_T\mathbf{t}_T = \mathbf{A}_T\boldsymbol{\lambda}_{T,\mathbb{R}} - \mathbf{A}_T\boldsymbol{\lambda}_{T,0} = \mathbf{0}$, the target $\mathbf{t}_T$ lies in $\ker_{\mathbb{R}}(\mathbf{A}_T)$, which equals $\mathrm{span}_{\mathbb{R}}(K_T) = \mathrm{span}_{\mathbb{R}}(\mathbf{b}^*_1, \ldots, \mathbf{b}^*_{d_T})$ because $\mathbf{A}_T$ has integer entries. Stage 1 of Algorithm 4

starts from $\boldsymbol{\lambda} = \boldsymbol{\lambda}_{T,0}$ and processes $i = d_T, \dots, 1$ in turn, following Babai's nearest-plane procedure [56]. At step i, it subtracts the integer multiple of $\mathbf{b}_i$ that places the coefficient of the current vector $\boldsymbol{\lambda}$ along $\mathbf{b}_i^*$ in $[-1/2, 1/2)$. Later steps subtract only multiples of $\mathbf{b}_j$ with $j < i$, which are orthogonal to $\mathbf{b}_i^*$, so this coefficient is not changed afterward. Since $\boldsymbol{\lambda}_{T,\mathbb{R}} \in (\ker_{\mathbb{R}}(\mathbf{A}_T))^{\perp}$ is orthogonal to every $\mathbf{b}_i^*$, the coefficients of $\boldsymbol{\lambda}_{\mathrm{dec}} - \boldsymbol{\lambda}_{T,\mathbb{R}} = \boldsymbol{\nu}_{\mathsf{Babai}} - \mathbf{t}_T$ along the vectors $\mathbf{b}_i^*$ coincide with those of $\boldsymbol{\lambda}_{\mathrm{dec}}$. This difference lies in $\mathrm{span}_{\mathbb{R}}(\mathbf{b}_1^*, \dots, \mathbf{b}_{d_T}^*)$, so $\boldsymbol{\nu}_{\mathsf{Babai}} - \mathbf{t}_T = \sum_{i=1}^{d_T} c_i' \mathbf{b}_i^*$ with $c_i' \in [-1/2, 1/2)$, and since the Gram-Schmidt vectors are mutually orthogonal,

$$\|\boldsymbol{\nu}_{\mathsf{Babai}} - \mathbf{t}_T\|_2 \le R_{\mathsf{Babai}}(\mathbf{B}_K') := \frac{1}{2}\left(\sum_{i=1}^{d_T} \|\mathbf{b}_i^*\|_2^2\right)^{1/2}. \tag{10}$$

Combining (9) and (10) gives $\|\boldsymbol{\lambda}_{\mathrm{dec}}\|_\infty \le \left(\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 + \frac{1}{4}\sum_{i=1}^{d_T} \|\mathbf{b}_i^*\|_2^2\right)^{1/2}$. Stage 2 of Algorithm 4 accepts only updates that strictly decrease the ℓ_∞-norm, so $\|\boldsymbol{\lambda}_T\|_\infty \le \|\boldsymbol{\lambda}_{\mathrm{dec}}\|_\infty$. Finally, Lemma 14 gives $\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2 \le \frac{1}{\sigma_{\min}^+(\mathbf{G}_T)}$. Substituting this bound proves (8). □

The certificate in (8) is computable from $\mathbf{G}_T$ and the reduced kernel basis used by the reconstruction algorithm, and applies to the vector actually returned. Unlike the initial HNF representative $\boldsymbol{\lambda}_{T,0}$, it is independent of the transformation matrix used in the row-HNF computation. The certificate may nevertheless be conservative, since it does not capture the full improvement from the greedy refinement. The proof in fact bounds the Euclidean norm of the Stage-1 Babai vector $\boldsymbol{\lambda}_{\mathrm{dec}}$; it is stated for the ℓ_∞-norm because the greedy refinement guarantees only $\|\boldsymbol{\lambda}_T\|_\infty \le \|\boldsymbol{\lambda}_{\mathrm{dec}}\|_\infty$ and may increase the Euclidean norm. Heuristically, if the coordinates of $\boldsymbol{\lambda}_{\mathrm{dec}}$ have comparable magnitudes, its ℓ_∞-norm is about $\sqrt{|T|}$ times smaller than its Euclidean norm, so (8) may overestimate $B_{\boldsymbol{\lambda}}(T)$ by roughly this factor. Appendix F relates the terms of the certificate to the sparsity of $\mathbf{G}$ through the kernel-lattice determinant and gives heuristic estimates under the geometric-series assumption.

Corollary 4 (Coefficient bound from the LLL root-Hermite factor) *Let $T \in \Gamma_1(\mathsf{pp})$ with $d_T > 0$, and let $D_T := \det(K_T)^{1/d_T}$. Let $\delta_{\mathsf{LLL}} \in (1/4, 1)$ be the Lovász parameter used by Algorithm 1, and let $\delta_{\mathrm{RHF}}^{\mathrm{wc}} := (\delta_{\mathsf{LLL}} - 1/4)^{-1/4}$ be the corresponding worst-case root-Hermite factor of LLL. If the Gram-Schmidt vectors in Lemma 16 satisfy $\|\mathbf{b}_i^*\|_2 \le \|\mathbf{b}_1^*\|_2$ for every $i \in [d_T]$, then*

$$B_{\boldsymbol{\lambda}}(T) \le \left(\frac{1}{\left(\sigma_{\min}^+(\mathbf{G}_T)\right)^2} + \frac{d_T}{4}\left(\delta_{\mathrm{RHF}}^{\mathrm{wc}}\right)^{2(d_T-1)} D_T^2\right)^{1/2}. \tag{11}$$

Proof The Lovász condition and size reduction give $\|\mathbf{b}_{i+1}^*\|_2^2 \ge (\delta_{\mathsf{LLL}} - 1/4)\|\mathbf{b}_i^*\|_2^2$ for every $i \in [d_T - 1]$, and hence $\|\mathbf{b}_1^*\|_2^2 \le (\delta_{\mathsf{LLL}} - 1/4)^{-(i-1)}\|\mathbf{b}_i^*\|_2^2$ for every $i \in [d_T]$. Multiplying these d_T inequalities and using $\prod_{i=1}^{d_T} \|\mathbf{b}_i^*\|_2 = \det(K_T)$ gives $\|\mathbf{b}_1\|_2 = \|\mathbf{b}_1^*\|_2 \le (\delta_{\mathsf{LLL}} - 1/4)^{-(d_T-1)/4} \det(K_T)^{1/d_T} = \left(\delta_{\mathrm{RHF}}^{\mathrm{wc}}\right)^{d_T-1} D_T$ [58]. By the hypothesis, $\sum_{i=1}^{d_T} \|\mathbf{b}_i^*\|_2^2 \le d_T\|\mathbf{b}_1^*\|_2^2 \le d_T\left(\delta_{\mathrm{RHF}}^{\mathrm{wc}}\right)^{2(d_T-1)} D_T^2$. Substituting this bound into (8) proves the claim. □

Here the root-Hermite factor is normalized by $\|\mathbf{b}_1\|_2 = \delta^{d_T-1}\det(K_T)^{1/d_T}$, consistently with (F4). The hypothesis holds whenever the Gram-Schmidt norms are non-increasing, as in the geometric-series model of Appendix F; LLL does not guarantee it, but it can be checked directly from $\mathbf{B}'_K$. For $\delta_{\mathsf{LLL}} = 0.99$, as used in Section 6.3, $\delta^{\mathrm{wc}}_{\mathrm{RHF}} \approx 1.078$, whereas LLL typically achieves a root-Hermite factor of about 1.02 in practice [59, 60]; substituting the latter for $\delta^{\mathrm{wc}}_{\mathrm{RHF}}$ gives a heuristic estimate. Bounds on D_T in terms of the sparsity of $\mathbf{G}$ follow from (F2) and (F3).

Since $\Gamma_1(\mathsf{pp})$ is finite, the maximum of the right-hand side of (8) over all $T \in \Gamma_1(\mathsf{pp})$ is a valid choice of the reconstruction bound in Item 2 of Definition 7, so SEC is everywhere-short with this bound. Evaluating this maximum exactly requires the certificate for every authorized set, which is infeasible for the instances considered in Section 6.

For the concrete instantiations, we evaluate $B_{\boldsymbol{\lambda}}(T) = \|\boldsymbol{\lambda}_T\|_\infty$ over the authorized sets examined during the parameter search. We choose $\beta_{\boldsymbol{\lambda}}$ slightly above the largest observed value $\hat{\beta}_{\boldsymbol{\lambda}}$, and similarly choose $\omega_{\boldsymbol{\lambda}}$ above the largest observed weight $\hat{\omega}_{\boldsymbol{\lambda}}$. Thus, $(\beta_{\boldsymbol{\lambda}}, \omega_{\boldsymbol{\lambda}})$ are empirically selected parameters rather than worst-case consequences of (8) over all $T \in \Gamma_1(\mathsf{pp})$. An application may impose both as additional admissibility checks; the selected margins then affect how many candidate sets are rejected, rather than the guarantees conditioned on satisfying the bounds. The sampling methodology, observed values, and selected margins are reported in Section 6.3.

If $\|\boldsymbol{\lambda}_T\|_\infty \le \beta_{\boldsymbol{\lambda}}$ and $wt_H(\boldsymbol{\lambda}_T) \le \omega_{\boldsymbol{\lambda}}$, then every family $(\mathbf{x}_i)_{i\in T}$ with $\|\mathbf{x}_i\|_\infty \le \beta_{\mathsf{sh}}$, such as shares bounded by Lemma 12, satisfies $\|\sum_{i\in T}\lambda_{T,i}\mathbf{x}_i\|_\infty \le \omega_{\boldsymbol{\lambda}}\beta_{\boldsymbol{\lambda}}\beta_{\mathsf{sh}} \le |T|\beta_{\boldsymbol{\lambda}}\beta_{\mathsf{sh}}$. This bounds every linear combination computed with a reconstruction vector satisfying the selected coefficient and weight bounds, as required by applications such as threshold signing.

4.3 Linearity and Public-Map Compatibility

For fixed public parameters pp and an authorized set $T \in \Gamma_1(\mathsf{pp})$, let $\boldsymbol{\lambda}_T \in \mathbb{Z}^{|T|}$ be the deterministic reconstruction vector produced by Algorithm 1. For vectors $(\mathbf{x}_i)_{i\in T}$ of any common dimension, let $\mathbf{X}_T$ denote the matrix whose rows are their transposes, ordered according to the indices in T. We use the same notation $\mathsf{Rec}_{\mathsf{pp}}$ for the dimension-independent reconstruction rule $\mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{x}_i)_{i\in T}) := (\boldsymbol{\lambda}_T^\top \mathbf{X}_T)^\top$. We record two consequences that make SEC compatible with public linear computations on shares.

Lemma 17 (Linearity of reconstruction) *Fix public parameters* pp *and an authorized set* $T \in \Gamma_1(\mathsf{pp})$. *For every* $\alpha \in \mathbb{Z}$ *and every* $\mathbf{x}_i, \mathbf{y}_i \in \mathbb{Z}^D$, $i \in T$, *we have* $\mathsf{Rec}_{\mathsf{pp}}(T, (\alpha\mathbf{x}_i + \mathbf{y}_i)_{i\in T}) = \alpha\,\mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{x}_i)_{i\in T}) + \mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{y}_i)_{i\in T})$.

Lemma 18 (Compatibility with public matrix multiplication) *Fix public parameters* pp *and an authorized set* $T \in \Gamma_1(\mathsf{pp})$. *Let* $\mathbf{M} \in \mathbb{Z}^{p\times D}$ *and let* $\mathbf{x}_i \in \mathbb{Z}^D$ *for every* $i \in T$. *Then* $\mathbf{M}\,\mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{x}_i)_{i\in T}) = \mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{M}\mathbf{x}_i)_{i\in T})$.

Both lemmas follow directly from $\mathsf{Rec}_{\mathsf{pp}}(T, (\mathbf{x}_i)_{i\in T}) = (\boldsymbol{\lambda}_T^\top \mathbf{X}_T)^\top$; the short computations are given in Appendix D. The first property permits local linear combinations of shares to be reconstructed as the corresponding linear combinations of their reconstructed values, and the second permits a public matrix to be applied locally to each share before reconstruction.

5 Privacy of SEC

The privacy analysis follows the five-step roadmap in Table 2: we exclude exact cancellation attacks; express the correlated multi-query view as a stacked linear hint; bound its joint precision and perturbation parameters; derive uniform bounds and the required Gaussian approximation; and reduce the resulting HC-Leak experiment to Matrix Hint-MLWE and standard MLWE. The notation is summarized in Table 3, and lengthy proofs are deferred to Appendix H.

Table 2: Roadmap of the privacy analysis.

Step	Section	Main statement	Output
Structural condition	5.1	Proposition 1	no exact cancellation for $\lvert U\rvert \le k_p$
Stacked representation	5.2	(20), (21)	$\mathbf{z} = \mathbf{M}\mathbf{s} + \boldsymbol{\Phi}\mathbf{r}$
Joint precision	5.3	Lemma 26	pointwise precision bound
Uniform parameters	5.4	Lemmas 28 and 32	Ξ_{MQ}, σ_{MQ}, loss $2\epsilon'$
Reduction	5.5	Theorem 4	HC-LLR from MLWE

5.1 Structural Setup Condition

Fix a private coalition $U \in \Gamma_0$, whose share matrix is $\mathbf{S}_U = (\mathbf{g})_U \mathbf{s}^\top + \hat{\mathbf{G}}_U \mathbf{R}$. Before reducing privacy to MLWE, we must exclude exact linear combinations that annihilate the masking term while retaining a nonzero function of the secret. Because the shares are integral, an adversary may also reduce them modulo any $q' \ge 2$ and attempt cancellation over $\mathbb{Z}_{q'}$. By Lemma 42 in Appendix G, if $(\mathbf{g})_U \notin \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$, then some modular linear combination cancels the masking term and reveals a noiseless nonzero linear function of $\mathbf{s} \bmod q'$. We therefore impose the structural setup condition $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$ for every $U \in \Gamma_0$.

Proposition 1 (Exclusion of structural noise cancellation) *Fix public parameters* pp *and a private coalition $U \in \Gamma_0$. If $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$, then: (1) $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Q}}(\hat{\mathbf{G}}_U)$, and hence rational noise cancellation is impossible; (2) $\mathrm{Col}_{\mathbb{Z}}(\mathbf{G}_U) = \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$ and $\ker_{\mathbb{Z}}(\mathbf{G}_U^\top) = \ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U^\top)$, and hence integer noise cancellation is impossible; (3) for every $q' \ge 2$, $\overline{(\mathbf{g})_U} \in \mathrm{Col}_{\mathbb{Z}_{q'}}(\overline{\hat{\mathbf{G}}_U})$, and hence modular noise cancellation is impossible.*

Table 3: Notation for the privacy analysis.

Symbol	Meaning	Symbol	Meaning
$D = d\ell$	coefficient dimension of a secret or share	Q	number of leakage queries
U_{raw}	parties whose raw shares are exposed	V_s	parties whose responses leak in query s
$c_s,\ \mathbf{L}_s$	leakage description and its local map	ω_c	uniform bound on $\|\mathbf{L}_s\|_2$
$\mathcal{S}$	leakage-block indices (s,i), $i \in V_s$	$\mathcal{S}_{\mathrm{in}}, \mathcal{S}_{\mathrm{out}}$	blocks with $i \in U_{\mathrm{raw}}$, blocks with $i \notin U_{\mathrm{raw}}$
$\mathbf{r}_{\mathbf{R}}$	$\mathrm{vec}(\mathbf{R}^\top)$, sharing mask, width $\sigma_{\mathbf{R}}$	$\mathbf{r}_{i,s}$	fresh response mask, width $\sigma_{\mathbf{r}}$
$\hat{\mathbf{G}}_{\mathrm{raw}}$	$\hat{\mathbf{G}}_{U_{\mathrm{raw}}} \otimes \mathbf{I}_D$	$\mathbf{T}_{\mathcal{S}}$	$\big(\hat{\mathbf{G}}[i,\cdot] \otimes \mathbf{L}_s\big)_{(s,i)\in\mathcal{S}}$ (vertical stack of $\hat{\mathbf{G}}[i,\cdot] \otimes \mathbf{L}_s$ for $(s,i) \in \mathcal{S}$)
$\mathbf{M}_{\mathcal{U},\mathbf{c}},\ \mathbf{y}_{\mathcal{U},\mathbf{c}}$	stacked hint matrix and perturbation	$\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}$	target parameter matrix of $\mathbf{y}_{\mathcal{U},\mathbf{c}}$
$\mathbf{W}$	public elimination transformation	$\mathbf{\Sigma}_{\mathrm{cor}}, \mathbf{B}, \mathbf{C}$	correlated block, transfer matrix, Schur complement
α_{k_p}	$\min_{\|V\|=k_p} \sigma_{\min}(\hat{\mathbf{G}}_V)$	β_{k_p}	$\max_{\|U\|\le k_p} \|\hat{\mathbf{G}}_U\|_2$
γ_{k_p}	$\min_{\|U\|\le k_p} \sigma_{\min}$ of peeled submatrix	$\gamma^{\mathrm{ker}}_{k_p}$	uniform kernel Gram-Schmidt bound
σ_{eff}	effective perturbation width	$\Xi_{\mathrm{MQ}},\ \sigma_{\mathrm{MQ}}$	uniform precision and MLWE source width

Writing $(\mathbf{g})_U = \hat{\mathbf{G}}_U\mathbf{a}$ for some integral $\mathbf{a}$ proves all three claims; the details are given in Appendix G. The rational condition also has a coding-theoretic interpretation through the dual distance of the masking code.

Lemma 19 *Let $\hat{C}_{\mathbb{Q}} := \mathrm{Col}_{\mathbb{Q}}(\hat{\mathbf{G}}) \subseteq \mathbb{Q}^n$, so that $\hat{C}^{\perp}_{\mathbb{Q}} = \ker_{\mathbb{Q}}(\hat{\mathbf{G}}^\top)$, and let $\hat{d}^{\perp}_{\mathbb{Q}}$ be the minimum Hamming distance of $\hat{C}^{\perp}_{\mathbb{Q}}$. Then, for every $U \subseteq [n]$ satisfying $|U| \le \hat{d}^{\perp}_{\mathbb{Q}} - 1$, we have $\ker_{\mathbb{Q}}(\hat{\mathbf{G}}_U^\top) = \{\mathbf{0}\}$. Consequently, rational noise cancellation is impossible for every such U.*

Proof A nonzero $\boldsymbol{\alpha}_U \in \ker_{\mathbb{Q}}(\hat{\mathbf{G}}_U^\top)$, extended by zero outside U, would be a nonzero codeword of $\hat{C}^{\perp}_{\mathbb{Q}}$ of weight at most $|U| \le \hat{d}^{\perp}_{\mathbb{Q}} - 1$, a contradiction. The final claim follows from Lemma 37. □

When $\mathbf{H}$ is a lifted finite-field parity-check matrix, Lemma 1 transfers the finite-field distance bound to the rational kernel code, which is unchanged by multiplication with the signed permutation $\mathbf{P}$. Dual distance, however, neither implies the integral condition nor excludes cancellation modulo an adversarially chosen modulus. The lossless-expander construction gives the stronger guarantee $\mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U) = \mathbb{Z}^{|U|}$ for every

$|U| \le k_p$ by Lemma 7, so the structural condition holds for every private coalition independently of $\mathbf{g}$.

5.2 Multi-Query HC-Leak Representation

Fix $d, \ell \in \mathbb{Z}^+$ and write $D := d\ell$. Let $U_{\mathrm{raw}} \subseteq [n]$ with $|U_{\mathrm{raw}}| \le k_p$ be the set of parties whose raw shares are exposed. The raw-share view is

$$\mathbf{S}_{U_{\mathrm{raw}}} = (\mathbf{g})_{U_{\mathrm{raw}}} \mathbf{s}^\top + \hat{\mathbf{G}}_{U_{\mathrm{raw}}} \mathbf{R}, \tag{12}$$

where $\mathbf{s} \in \mathbb{Z}^D$ is the coefficient embedding of the shared secret and $\mathbf{R} \in \mathbb{Z}^{(k-1)\times D}$ is the sharing mask. By Section 5.1, $\mathbf{S}_{U_{\mathrm{raw}}} = \hat{\mathbf{G}}_{U_{\mathrm{raw}}}(\mathbf{R} + \mathbf{a}\mathbf{s}^\top)$ for some $\mathbf{a} \in \mathbb{Z}^{k-1}$, so the secret enters the raw-share view only through a shift of the common mask.

The adversary also observes noisy linear local leakage across at most Q queries. This model is motivated by lattice-based threshold signatures, in which party i answers query s with a partial response $\boldsymbol{z}_{i,s} := c_s \boldsymbol{s}_i + \boldsymbol{r}_{i,s} \in \mathcal{R}^\ell$, where c_s is the challenge and $\boldsymbol{r}_{i,s}$ is a fresh mask. Each query $s \in [Q]$ is therefore specified by a public $c_s \in \mathcal{C}$, where $\mathcal{C} \subseteq \mathcal{R}$ is a fixed set of bounded-norm ring elements. Under the coefficient embedding, multiplication by c_s on $\mathcal{R}^\ell$ is represented by $\mathbf{L}_s := \mathbf{I}_\ell \otimes \mathrm{Rot}(c_s) \in \mathbb{Z}^{D\times D}$. We assume $\|\mathbf{L}_s\|_2 \le \omega_c$ for every $c_s \in \mathcal{C}$. For example, if every element of $\mathcal{C}$ has exactly ω_c nonzero coefficients in $\{-1, 1\}$, then $\|\mathbf{L}_s\|_2 \le \|c_s\|_1 = \omega_c$.

Let $V_s \subseteq [n]$ index the response blocks exposed in query s. For every $i \in V_s$, the exposed coefficient-embedded block is

$$\mathbf{z}_{i,s} := \mathbf{L}_s \mathbf{s}_i + \mathbf{r}_{i,s}, \tag{13}$$

where $\mathbf{r}_{i,s} \leftarrow \mathcal{D}_{\mathbb{Z}^D,\sigma_{\mathbf{r}}}$ are mutually independent and independent of $\mathbf{R}$ and $\mathbf{s}$. We write $\mathcal{U} := \big(U_{\mathrm{raw}}, (V_s)_{s\in[Q]}\big)$ for the multi-query leakage pattern and $\mathcal{S} := \{(s,i) : s \in [Q],\, i \in V_s\}$ for its leakage-block index set.

5.2.1 Fixed-Pattern HC-Leak View

Fix an admissible pattern $\mathcal{U}$, a public leakage-description sequence $\mathbf{c} := (c_s)_{s\in[Q]}$, and a public ordering of $\mathcal{S}$. Let $\boldsymbol{A}' \xleftarrow{\$} \mathcal{R}_q^{p\times(\ell-p)}$, set $\boldsymbol{A} := \big[\boldsymbol{I}_p \ \boldsymbol{A}'\big]$, and define the hard-core map $f_{\boldsymbol{A}}(\boldsymbol{s}) := \boldsymbol{A}\boldsymbol{s} \bmod q \in \mathcal{R}_q^p$. We encode the auxiliary view $\mathsf{Leak}_{\mathcal{U},\mathbf{c}} := (\mathbf{S}_{U_{\mathrm{raw}}}, (\mathbf{z}_{i,s})_{(s,i)\in\mathcal{S}})$ by

$$\mathbf{z}_{\mathcal{U},\mathbf{c}} := \begin{bmatrix} \mathrm{vec}\left(\mathbf{S}_{U_{\mathrm{raw}}}^\top\right) \\ (\mathbf{z}_{i,s})_{(s,i)\in\mathcal{S}} \end{bmatrix} \in \mathbb{Z}^{M'}, \tag{14}$$

where $M' := (|U_{\mathrm{raw}}| + |\mathcal{S}|)D$. In the fixed-pattern $\mathsf{HC\text{-}Leak}$ experiment, the real branch supplies $f_{\boldsymbol{A}}(\boldsymbol{s})$ together with $\mathsf{Leak}_{\mathcal{U},\mathbf{c}}$, whereas the random branch replaces $f_{\boldsymbol{A}}(\boldsymbol{s})$ by an independently uniform element of $\mathcal{R}_q^p$; the auxiliary leakage has the same marginal distribution in both branches.

5.2.2 Stacked Matrix-Hint Representation

Write $\mathbf{r}_{\mathbf{R}} := \mathrm{vec}(\mathbf{R}^\top) \in \mathbb{Z}^{(k-1)D}$ and $\hat{\mathbf{g}}_i^\top := \hat{\mathbf{G}}[i,\cdot]$ for $i \in [n]$, and define

$$\hat{\mathbf{G}}_{\mathrm{raw}} := \hat{\mathbf{G}}_{U_{\mathrm{raw}}} \otimes \mathbf{I}_D \in \mathbb{Z}^{|U_{\mathrm{raw}}|D \times (k-1)D}, \tag{15}$$

$$\mathbf{T}_{\mathcal{S}} := \left(\hat{\mathbf{g}}_i^\top \otimes \mathbf{L}_s\right)_{(s,i)\in\mathcal{S}} \in \mathbb{Z}^{|\mathcal{S}|D \times (k-1)D}, \tag{16}$$

where the block rows of $\mathbf{T}_{\mathcal{S}}$ are stacked in the fixed ordering of $\mathcal{S}$. Applying $\mathrm{vec}(\mathbf{XYZ}) = (\mathbf{Z}^\top \otimes \mathbf{X})\mathrm{vec}(\mathbf{Y})$ to (12) and to $\mathbf{s}_i = \mathbf{g}[i]\mathbf{s} + \mathbf{R}^\top \hat{\mathbf{g}}_i$ gives

$$\mathrm{vec}\left(\mathbf{S}_{U_{\mathrm{raw}}}^\top\right) = \mathbf{M}_{\mathrm{raw}}\mathbf{s} + \mathbf{y}_{\mathrm{raw}}, \tag{17}$$

$$\mathbf{z}_{i,s} = \mathbf{M}_{i,s}\mathbf{s} + \mathbf{y}_{i,s}, \qquad (s,i) \in \mathcal{S}, \tag{18}$$

where $\mathbf{M}_{\mathrm{raw}} := (\mathbf{g})_{U_{\mathrm{raw}}} \otimes \mathbf{I}_D$, $\mathbf{y}_{\mathrm{raw}} := \hat{\mathbf{G}}_{\mathrm{raw}}\mathbf{r}_{\mathbf{R}}$, $\mathbf{M}_{i,s} := \mathbf{g}[i]\mathbf{L}_s$, and $\mathbf{y}_{i,s} := (\hat{\mathbf{g}}_i^\top \otimes \mathbf{L}_s)\mathbf{r}_{\mathbf{R}} + \mathbf{r}_{i,s}$. Stacking these blocks in the fixed ordering, define

$$\mathbf{M}_{\mathcal{U},\mathbf{c}} := \begin{bmatrix} \mathbf{M}_{\mathrm{raw}} \\ (\mathbf{M}_{i,s})_{(s,i)\in\mathcal{S}} \end{bmatrix} \in \mathbb{Z}^{M' \times D}, \qquad \mathbf{y}_{\mathcal{U},\mathbf{c}} := \begin{bmatrix} \mathbf{y}_{\mathrm{raw}} \\ (\mathbf{y}_{i,s})_{(s,i)\in\mathcal{S}} \end{bmatrix} \in \mathbb{Z}^{M'}. \tag{19}$$

The complete auxiliary view is then the single stacked linear hint

$$\mathbf{z}_{\mathcal{U},\mathbf{c}} = \mathbf{M}_{\mathcal{U},\mathbf{c}}\mathbf{s} + \mathbf{y}_{\mathcal{U},\mathbf{c}}. \tag{20}$$

Its perturbation is one integer linear image of the independent masks: with $\mathbf{r}_{\mathcal{S}} := (\mathbf{r}_{i,s})_{(s,i)\in\mathcal{S}}$,

$$\mathbf{y}_{\mathcal{U},\mathbf{c}} = \mathbf{\Phi}_{\mathcal{U},\mathbf{c}} \begin{bmatrix} \mathbf{r}_{\mathbf{R}} \\ \mathbf{r}_{\mathcal{S}} \end{bmatrix}, \qquad \mathbf{\Phi}_{\mathcal{U},\mathbf{c}} := \begin{bmatrix} \hat{\mathbf{G}}_{\mathrm{raw}} & \mathbf{O} \\ \mathbf{T}_{\mathcal{S}} & \mathbf{I}_{|\mathcal{S}|D} \end{bmatrix}. \tag{21}$$

Accordingly, we define the target parameter matrix of the joint perturbation by

$$\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}} := \mathbf{\Phi}_{\mathcal{U},\mathbf{c}}\mathbf{\Sigma}_0\mathbf{\Phi}_{\mathcal{U},\mathbf{c}}^\top, \tag{22}$$

$$\mathbf{\Sigma}_0 := \mathrm{diag}\left(\sigma_{\mathbf{R}}^2\mathbf{I}_{(k-1)D}, \sigma_{\mathbf{r}}^2\mathbf{I}_{|\mathcal{S}|D}\right), \tag{23}$$

where $\mathbf{\Sigma}_0$ is the parameter matrix of the independent masks $(\mathbf{r}_{\mathbf{R}}, \mathbf{r}_{\mathcal{S}})$. Because the blocks of $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ share the common mask $\mathbf{r}_{\mathbf{R}}$, the matrix $\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}$ need not be block diagonal, and the view must be treated as one Matrix-Hint instance rather than as independent per-block hints. Section 5.3 bounds the resulting precision term, and Section 5.4 shows that $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ is statistically close to $\mathcal{D}_{\mathbb{Z}^{M'},\sqrt{\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}}}$.

5.3 Bounding the Joint Precision Term

To apply Theorem 3, we bound $\|\mathbf{M}_{\mathcal{U},\mathbf{c}}^\top \mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}^{-1} \mathbf{M}_{\mathcal{U},\mathbf{c}}\|_2$. The response blocks are divided according to whether the same party's raw share is exposed:

$$\mathcal{S}_{\mathrm{in}} := \{(s,i) \in \mathcal{S} : i \in U_{\mathrm{raw}}\}, \qquad \mathcal{S}_{\mathrm{out}} := \{(s,i) \in \mathcal{S} : i \notin U_{\mathrm{raw}}\}. \tag{24}$$

For $(s,i) \in \mathcal{S}_{\text{in}}$, the adversary can compute $\mathbf{z}_{i,s} - \mathbf{L}_s \mathbf{s}_i = \mathbf{r}_{i,s}$, so a public transformation isolates these blocks as secret-independent fresh masks. After this elimination, a Schur-complement decomposition separates the raw-share precision from the conditional precision of the unpaired blocks, and bounding both components gives the pointwise estimate of Lemma 26.

5.3.1 Transformation and Block Structure

We reorder the block rows of $\mathbf{M}_{\mathcal{U},\mathbf{c}}$ and $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ as raw-share blocks, then $\mathcal{S}_{\text{in}}$, then $\mathcal{S}_{\text{out}}$, and write the corresponding blocks as $\mathbf{M}_{\text{raw}}, \mathbf{M}_{\text{in}}, \mathbf{M}_{\text{out}}$ and $\mathbf{y}_{\text{raw}}, \mathbf{y}_{\text{in}}, \mathbf{y}_{\text{out}}$. A common permutation of the hint and perturbation preserves the precision matrix, so we keep the notation; empty blocks are omitted throughout. Let $\mathbf{L}_{\mathcal{S}_{\text{in}}}$ be the block-sparse matrix whose block in row $(s,i) \in \mathcal{S}_{\text{in}}$ and column $i \in U_{\text{raw}}$ equals $\mathbf{L}_s$. Then

$$\mathbf{L}_{\mathcal{S}_{\text{in}}} \mathbf{M}_{\text{raw}} = \mathbf{M}_{\text{in}}, \tag{25}$$

$$\mathbf{y}_{\text{in}} - \mathbf{L}_{\mathcal{S}_{\text{in}}} \mathbf{y}_{\text{raw}} = (\mathbf{r}_{i,s})_{(s,i)\in\mathcal{S}_{\text{in}}} =: \mathbf{r}_{\text{in}}. \tag{26}$$

Define the public transformation

$$\mathbf{W} := \begin{bmatrix} \mathbf{I} & \mathbf{O} & \mathbf{O} \\ -\mathbf{L}_{\mathcal{S}_{\text{in}}} & \mathbf{I} & \mathbf{O} \\ \mathbf{O} & \mathbf{O} & \mathbf{I} \end{bmatrix} \in GL_{M'}(\mathbb{Z}), \tag{27}$$

which is unimodular because it is block lower triangular with identity diagonal blocks. Writing $\widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}} := \mathbf{W}\mathbf{M}_{\mathcal{U},\mathbf{c}}$ and $\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}} := \mathbf{W}\mathbf{y}_{\mathcal{U},\mathbf{c}}$, (25) and (26) yield the zero block:

$$\widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}} = \begin{bmatrix} \mathbf{M}_{\text{raw}} \\ \mathbf{O} \\ \mathbf{M}_{\text{out}} \end{bmatrix}, \qquad \widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}} = \begin{bmatrix} \mathbf{y}_{\text{raw}} \\ \mathbf{r}_{\text{in}} \\ \mathbf{y}_{\text{out}} \end{bmatrix}. \tag{28}$$

Since $\mathbf{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}} = \mathbf{W}\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}\mathbf{W}^\top$ and $\mathbf{W}$ is invertible,

$$\mathbf{M}_{\mathcal{U},\mathbf{c}}^\top \mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}^{-1} \mathbf{M}_{\mathcal{U},\mathbf{c}} = \widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}}^\top \mathbf{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}^{-1} \widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}}. \tag{29}$$

Let $\mathbf{\Sigma}_{\text{raw}}$, $\mathbf{\Sigma}_{\text{out}}$, and $\mathbf{\Sigma}_{\text{raw,out}}$ be the parameter blocks of $\mathbf{y}_{\text{raw}}$, of $\mathbf{y}_{\text{out}}$, and between them, and define

$$\mathbf{\Sigma}_{\text{cor}} := \begin{bmatrix} \mathbf{\Sigma}_{\text{raw}} & \mathbf{\Sigma}_{\text{raw,out}} \\ \mathbf{\Sigma}_{\text{raw,out}}^\top & \mathbf{\Sigma}_{\text{out}} \end{bmatrix}. \tag{30}$$

Lemma 20 (Block structure of the transformed perturbation) *Assume that $\sigma_{\mathbf{r}} > 0$ and $\mathbf{\Sigma}_{\text{cor}} \succ 0$. Write $\mathbf{\Sigma}_{\text{cor}}^{-1}[a,b]$, $a,b \in \{1,2\}$, for the blocks of $\mathbf{\Sigma}_{\text{cor}}^{-1}$ in the raw-share and $\mathcal{S}_{\text{out}}$ coordinates. Then, in the $(U_{\text{raw}}, \mathcal{S}_{\text{in}}, \mathcal{S}_{\text{out}})$ ordering,*

$$\mathbf{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}^{-1} = \begin{bmatrix} \mathbf{\Sigma}_{\text{cor}}^{-1}[1,1] & \mathbf{O} & \mathbf{\Sigma}_{\text{cor}}^{-1}[1,2] \\ \mathbf{O} & \sigma_{\mathbf{r}}^{-2}\mathbf{I} & \mathbf{O} \\ \mathbf{\Sigma}_{\text{cor}}^{-1}[2,1] & \mathbf{O} & \mathbf{\Sigma}_{\text{cor}}^{-1}[2,2] \end{bmatrix}, \tag{31}$$

and $\|\mathbf{\Sigma}^{-1}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}\|_2 = \max\{\|\mathbf{\Sigma}^{-1}_{\mathrm{cor}}\|_2, \sigma_{\mathbf{r}}^{-2}\}$.

The lemma holds because $\mathbf{r}_{\mathrm{in}}$ is independent of $\mathbf{y}_{\mathrm{raw}}$ and $\mathbf{y}_{\mathrm{out}}$; the proof is given in Appendix H.1. Thus, all remaining dependence induced by the common sharing mask is confined to the raw-share and $\mathcal{S}_{\mathrm{out}}$ coordinates.

5.3.2 Elimination and Schur-Complement Decomposition

The isolated fresh-mask block is multiplied on both sides by the zero block of $\widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}}$ and therefore contributes nothing to the precision. Define $\mathbf{M}_{\mathrm{cor}} := \left[\mathbf{M}_{\mathrm{raw}}^\top \ \mathbf{M}_{\mathrm{out}}^\top\right]^\top$.

Lemma 21 (Elimination of paired response blocks) *Under the conditions of Lemma 20,*

$$\mathbf{M}_{\mathcal{U},\mathbf{c}}^\top \mathbf{\Sigma}^{-1}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}} \mathbf{M}_{\mathcal{U},\mathbf{c}} = \mathbf{M}_{\mathrm{cor}}^\top \mathbf{\Sigma}_{\mathrm{cor}}^{-1} \mathbf{M}_{\mathrm{cor}}. \tag{32}$$

Assume that $\mathbf{\Sigma}_{\mathrm{raw}} \succ 0$, and define the parameter-transfer matrix and the Schur complement of $\mathbf{\Sigma}_{\mathrm{raw}}$ in $\mathbf{\Sigma}_{\mathrm{cor}}$ by

$$\mathbf{B} := \mathbf{\Sigma}_{\mathrm{raw}}^{-1} \mathbf{\Sigma}_{\mathrm{raw,out}}, \tag{33}$$

$$\mathbf{C} := \mathbf{\Sigma}_{\mathrm{out}} - \mathbf{\Sigma}_{\mathrm{raw,out}}^\top \mathbf{\Sigma}_{\mathrm{raw}}^{-1} \mathbf{\Sigma}_{\mathrm{raw,out}}. \tag{34}$$

Lemma 22 (Schur-complement precision decomposition) *Assume that* $\mathbf{\Sigma}_{\mathrm{raw}} \succ 0$ *and* $\mathbf{C} \succ 0$. *Then*

$$\mathbf{M}_{\mathrm{cor}}^\top \mathbf{\Sigma}_{\mathrm{cor}}^{-1} \mathbf{M}_{\mathrm{cor}} = \mathbf{M}_{\mathrm{raw}}^\top \mathbf{\Sigma}_{\mathrm{raw}}^{-1} \mathbf{M}_{\mathrm{raw}} + \left(\mathbf{M}_{\mathrm{out}} - \mathbf{B}^\top \mathbf{M}_{\mathrm{raw}}\right)^\top \mathbf{C}^{-1} \left(\mathbf{M}_{\mathrm{out}} - \mathbf{B}^\top \mathbf{M}_{\mathrm{raw}}\right), \tag{35}$$

where the second term is omitted if $\mathcal{S}_{\mathrm{out}} = \emptyset$.

Both lemmas follow from (29), (31), and the block $\mathbf{LDL}^\top$ factorization of $\mathbf{\Sigma}_{\mathrm{cor}}$; see Appendix H.1.

5.3.3 Bounds on the Parameter and Hint Blocks

For $1 \le |U_{\mathrm{raw}}| \le k_p$, Lemma 7 implies that $\hat{\mathbf{G}}_{U_{\mathrm{raw}}}$ has full row rank, so $\sigma_{\min}(\hat{\mathbf{G}}_{U_{\mathrm{raw}}}) > 0$. Recall $\hat{\mathbf{G}}_{\mathrm{raw}}$ and $\mathbf{T}_{\mathcal{S}}$ from (15), and let

$$\mathbf{T}_{\mathrm{out}} := \left(\hat{\mathbf{g}}_i^\top \otimes \mathbf{L}_s\right)_{(s,i)\in\mathcal{S}_{\mathrm{out}}} \tag{36}$$

be the rows of $\mathbf{T}_{\mathcal{S}}$ indexed by $\mathcal{S}_{\mathrm{out}}$, so that $\mathbf{y}_{\mathrm{out}} = \mathbf{T}_{\mathrm{out}}\mathbf{r}_{\mathbf{R}} + \mathbf{r}_{\mathrm{out}}$ with $\mathbf{r}_{\mathrm{out}} := (\mathbf{r}_{i,s})_{(s,i)\in\mathcal{S}_{\mathrm{out}}}$. By (22),

$$\mathbf{\Sigma}_{\mathrm{raw}} = \sigma_{\mathbf{R}}^2 \hat{\mathbf{G}}_{\mathrm{raw}} \hat{\mathbf{G}}_{\mathrm{raw}}^\top, \tag{37}$$

$$\mathbf{\Sigma}_{\mathrm{raw,out}} = \sigma_{\mathbf{R}}^2 \hat{\mathbf{G}}_{\mathrm{raw}} \mathbf{T}_{\mathrm{out}}^\top, \tag{38}$$

$$\mathbf{\Sigma}_{\mathrm{out}} = \sigma_{\mathbf{R}}^2 \mathbf{T}_{\mathrm{out}} \mathbf{T}_{\mathrm{out}}^\top + \sigma_{\mathbf{r}}^2 \mathbf{I}. \tag{39}$$

Lemma 23 (Raw parameter and parameter-transfer bounds) *Suppose that* $1 \le |U_{\mathrm{raw}}| \le k_p$ *and* $\sigma_{\mathbf{R}} > 0$. *Then* $\mathbf{\Sigma}_{\mathrm{raw}} \succ 0$ *and* $\|\mathbf{\Sigma}_{\mathrm{raw}}^{-1}\|_2 = 1/(\sigma_{\mathbf{R}}^2 \sigma_{\min}^2(\hat{\mathbf{G}}_{U_{\mathrm{raw}}}))$. *If* $\mathcal{S}_{\mathrm{out}} \neq \emptyset$, *then* $\|\mathbf{B}\|_2 \le \omega_c \sqrt{Q} \|\hat{\mathbf{G}}\|_2 / \sigma_{\min}(\hat{\mathbf{G}}_{U_{\mathrm{raw}}})$. *Moreover, for the expander-based construction,* $\|\hat{\mathbf{G}}\|_2 \le \sqrt{c_{\exp} n}$.

Lemma 24 (Conditional response-parameter bound) *Suppose that* $1 \le |U_{\mathrm{raw}}| \le k_p$, $\mathcal{S}_{\mathrm{out}} \neq \emptyset$, *and* $\sigma_{\mathbf{R}}, \sigma_{\mathbf{r}} > 0$. *Then* $\mathbf{C} \succeq \sigma_{\mathbf{r}}^2 \mathbf{I} \succ 0$, *and hence* $\|\mathbf{C}^{-1}\|_2 \le \sigma_{\mathbf{r}}^{-2}$.

Lemma 25 (Raw, outside, and residual hint bounds) *Suppose that* $1 \le |U_{\mathrm{raw}}| \le k_p$ *and* $\|\mathbf{g}\|_\infty \le \beta_{\mathbf{g}}$. *Then* $\|\mathbf{M}_{\mathrm{raw}}\|_2 \le \beta_{\mathbf{g}} \sqrt{k_p}$. *If* $\mathcal{S}_{\mathrm{out}} \neq \emptyset$, *then* $\|\mathbf{M}_{\mathrm{out}}\|_2 \le \beta_{\mathbf{g}} \omega_c \sqrt{nQ}$ *and*

$$\left\| \mathbf{M}_{\mathrm{out}} - \mathbf{B}^\top \mathbf{M}_{\mathrm{raw}} \right\|_2 \le \beta_{\mathbf{g}} \omega_c \sqrt{Q} \left(\sqrt{n} + \sqrt{k_p} \, \frac{\|\hat{\mathbf{G}}\|_2}{\sigma_{\min}(\hat{\mathbf{G}}_{U_{\mathrm{raw}}})} \right). \tag{40}$$

The first lemma uses the Kronecker structure of $\mathbf{\Sigma}_{\mathrm{raw}}$ and the factorization $\mathbf{T}_{\mathrm{out}} = \mathbf{L}_{\mathcal{S}_{\mathrm{out}}}(\hat{\mathbf{G}}_{\mathcal{S}_{\mathrm{out}}} \otimes \mathbf{I}_D)$, in which each party appears in at most Q block rows. The second follows because $\mathbf{C} - \sigma_{\mathbf{r}}^2 \mathbf{I}$ is a congruence of an orthogonal projector. The third follows from the block structure of the hint matrices, the operator-norm bounds on the local maps, and the triangle inequality. The proofs are given in Appendix H.1.

5.3.4 Final Joint Precision Bound

Lemma 26 (Joint precision bound) *Suppose that* $1 \le |U_{\mathrm{raw}}| \le k_p$, $\sigma_{\mathbf{R}}, \sigma_{\mathbf{r}} > 0$, *and* $\|\mathbf{g}\|_\infty \le \beta_{\mathbf{g}}$. *Then*

$$\left\| \mathbf{M}_{\mathcal{U},\mathbf{c}}^\top \mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U}},\mathbf{c}}^{-1} \mathbf{M}_{\mathcal{U},\mathbf{c}} \right\|_2 \le \frac{\beta_{\mathbf{g}}^2 k_p}{\sigma_{\mathbf{R}}^2 \sigma_{\min}^2(\hat{\mathbf{G}}_{U_{\mathrm{raw}}})} + \frac{\beta_{\mathbf{g}}^2 \omega_c^2 Q}{\sigma_{\mathbf{r}}^2} \left(\sqrt{n} + \sqrt{k_p} \, \frac{\|\hat{\mathbf{G}}\|_2}{\sigma_{\min}(\hat{\mathbf{G}}_{U_{\mathrm{raw}}})} \right)^2. \tag{41}$$

Proof sketch By Lemmas 21 and 22, the precision matrix is the sum of $\mathbf{M}_{\mathrm{raw}}^\top \mathbf{\Sigma}_{\mathrm{raw}}^{-1} \mathbf{M}_{\mathrm{raw}}$ and the term involving $\mathbf{C}^{-1}$. Bounding the first by Lemmas 23 and 25 and the second by Lemma 24 and (40) gives the two terms. The full argument, including the case $\mathcal{S}_{\mathrm{out}} = \emptyset$, is in Appendix H.1. □

The first term bounds the precision contributed by the exposed raw shares, and the second the conditional precision contributed by the response blocks whose raw shares are not exposed; paired response blocks contribute no separate term. By Remark 1, the reduction loses only a constant factor in the MLWE source width when the right-hand side of (41) is at most $c/\sigma_{\mathbf{s}}^2$ for a constant $c > 0$. It suffices that each of its two terms is at most $c/(2\sigma_{\mathbf{s}}^2)$, that is, $\sigma_{\mathbf{R}}/\sigma_{\mathbf{s}} \ge \sqrt{2/c}\, \beta_{\mathbf{g}} \sqrt{k_p} / \sigma_{\min}(\hat{\mathbf{G}}_{U_{\mathrm{raw}}})$ and $\sigma_{\mathbf{r}}/\sigma_{\mathbf{s}} \ge \sqrt{2/c}\, \beta_{\mathbf{g}} \omega_c \sqrt{Q} \left(\sqrt{n} + \sqrt{k_p}\, \|\hat{\mathbf{G}}\|_2 / \sigma_{\min}(\hat{\mathbf{G}}_{U_{\mathrm{raw}}}) \right)$. Under this sufficient allocation, the sharing-mask width exceeds the secret width by a factor proportional to $\sqrt{k_p}$, and the leakage-mask width exceeds it by a factor proportional to $\sqrt{Q}$. The concrete parameters of Section 6.4 do not fix $c = \sigma_{\mathbf{s}}^2 \, \Xi_{\mathrm{MQ}}$ in advance, where Ξ_{MQ} is the uniform precision parameter in (43). Instead, they impose the uniform version of this bound through σ_{MQ} in (58), and the selected widths give c between about 0.5 and 8.5, with $\sigma_{\mathbf{r}}/\sigma_{\mathbf{s}} \approx 2^{40}$ to $2^{42.5}$ for $Q = 2^{64}$.

Remark 5 (Boundary leakage patterns) If $U_{\mathrm{raw}} = \emptyset$ and $\mathcal{S} \neq \emptyset$, then $\|\mathbf{M}_{\mathcal{U},\mathbf{c}}^{\top}\boldsymbol{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}^{-1}\mathbf{M}_{\mathcal{U},\mathbf{c}}\|_2 \leq \beta_{\mathbf{g}}^2\omega_c^2 nQ/\sigma_{\mathbf{r}}^2$. If $U_{\mathrm{raw}} \neq \emptyset$ and $\mathcal{S}_{\mathrm{out}} = \emptyset$, only the first term of (41) remains. If $U_{\mathrm{raw}} = \emptyset$ and $\mathcal{S} = \emptyset$, no hint is present and the precision term is defined to be 0. These cases are verified in Appendix H.1.

The joint bound is tighter than the estimate obtained by bounding the two factors separately, namely $\|\mathbf{M}_{\mathcal{U},\mathbf{c}}^{\top}\mathbf{M}_{\mathcal{U},\mathbf{c}}\|_2\|\boldsymbol{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}^{-1}\|_2$, which still pays for the $\mathcal{S}_{\mathrm{in}}$ blocks whose hint-matrix block vanishes after the transformation $\mathbf{W}$. With the remaining parameters fixed, the separate estimate grows cubically in Q, whereas the leakage-dependent term in (41) grows linearly. We therefore apply Theorem 3 in the reduction.

5.4 Uniform Parameters for Multi-Query Privacy

We now replace the pattern-dependent quantities in Lemma 26 by parameters valid uniformly over every admissible leakage pattern and leakage-description sequence, that is, in the worst case over every set $U_{\mathrm{raw}} \subseteq [n]$ of exposed raw shares with $|U_{\mathrm{raw}}| \leq k_p$, all response sets $V_s \subseteq [n]$, and all $\mathbf{c} \in \mathcal{C}^Q$. We then establish the distributional condition required by Matrix Hint-MLWE.

5.4.1 Uniform Conditioning and Precision

Define

$$\alpha_{k_p} := \min\left\{\sigma_{\min}(\hat{\mathbf{G}}_V) : V \subseteq [n],\ |V| = k_p\right\}. \tag{42}$$

Lemma 27 (Uniform smallest-singular-value bound) *Suppose that $\hat{\mathbf{G}}_U$ has full row rank for every $U \subseteq [n]$ with $1 \leq |U| \leq k_p$. Then $\alpha_{k_p} > 0$, and $\sigma_{\min}(\hat{\mathbf{G}}_U) \geq \alpha_{k_p}$ for every nonempty $U \subseteq [n]$ with $|U| \leq k_p$.*

The second claim follows from eigenvalue interlacing, since $\hat{\mathbf{G}}_U\hat{\mathbf{G}}_U^{\top}$ is a principal submatrix of $\hat{\mathbf{G}}_V\hat{\mathbf{G}}_V^{\top}$ for every $V \supseteq U$ with $|V| = k_p$. For the expander-based instantiation, the full-row-rank hypothesis follows from Lemma 7. Assume the hypothesis of Lemma 27, so that $\alpha_{k_p} > 0$, and that $\sigma_{\mathbf{R}}, \sigma_{\mathbf{r}} > 0$ and $\|\mathbf{g}\|_\infty \leq \beta_{\mathbf{g}}$. Define the uniform multi-query precision parameter

$$\Xi_{\mathrm{MQ}} := \frac{\beta_{\mathbf{g}}^2 k_p}{\sigma_{\mathbf{R}}^2\alpha_{k_p}^2} + \frac{\beta_{\mathbf{g}}^2\omega_c^2 Q}{\sigma_{\mathbf{r}}^2}\left(\sqrt{n} + \sqrt{k_p}\,\frac{\|\hat{\mathbf{G}}\|_2}{\alpha_{k_p}}\right)^2. \tag{43}$$

Lemma 28 (Uniform multi-query precision bound) *For every admissible pattern $\mathcal{U}$ with $|U_{\mathrm{raw}}| \leq k_p$ and every $\mathbf{c} \in \mathcal{C}^Q$,*

$$\left\|\mathbf{M}_{\mathcal{U},\mathbf{c}}^{\top}\boldsymbol{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}^{-1}\mathbf{M}_{\mathcal{U},\mathbf{c}}\right\|_2 \leq \Xi_{\mathrm{MQ}}, \tag{44}$$

where the left-hand side is interpreted as 0 when the auxiliary view is empty.

The bound follows from Lemmas 26 and 27 and Remark 5; it is independent of $\mathbf{c}$ because each c_s enters only through $\|\mathbf{L}_s\|_2 \le \omega_c$. The proofs of both lemmas are given in Appendix H.2.

5.4.2 Linear-Image and Kernel Structure

The precision bound alone does not establish the distributional condition of Matrix Hint-MLWE: by (21), $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ is an integer linear image of independent discrete Gaussians and need not itself be an elliptic discrete Gaussian. This section, together with the subsequent Sections 5.4.3 and 5.4.4, derives kernel-smoothing conditions under which $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ is, uniformly over admissible realizations, within statistical distance ϵ' of the elliptic discrete Gaussian assumed in Theorem 3. Splitting $\mathbf{r}_{\mathcal{S}}$ into $\mathbf{r}_{\mathrm{in}}$ and $\mathbf{r}_{\mathrm{out}}$ and applying $\mathbf{W}$ to (21) gives

$$\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}} = \widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}\,\mathbf{r}_0, \qquad \widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}} := \mathbf{W}\mathbf{\Phi}_{\mathcal{U},\mathbf{c}} = \begin{bmatrix} \hat{\mathbf{G}}_{\mathrm{raw}} & \mathbf{O} & \mathbf{O} \\ \mathbf{O} & \mathbf{I} & \mathbf{O} \\ \mathbf{T}_{\mathrm{out}} & \mathbf{O} & \mathbf{I} \end{bmatrix}, \tag{45}$$

where $\mathbf{r}_0 := \begin{bmatrix}\mathbf{r}_{\mathbf{R}}^\top & \mathbf{r}_{\mathrm{in}}^\top & \mathbf{r}_{\mathrm{out}}^\top\end{bmatrix}^\top$ in agreement with (28). Let $M := (k-1+|\mathcal{S}|)D$, so that $\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}} \in \mathbb{Z}^{M'\times M}$ and $\mathbf{r}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^M,\sqrt{\mathbf{\Sigma}_0}}$, where $\mathbf{\Sigma}_0$ is as in (23), with the fresh-mask coordinates ordered as $(\mathcal{S}_{\mathrm{in}}, \mathcal{S}_{\mathrm{out}})$.

Lemma 29 (Integral surjectivity of the transformed leakage map) *For every admissible fixed public realization* $(\mathcal{U},\mathbf{c})$, $\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}\mathbb{Z}^M = \mathbb{Z}^{M'}$.

Surjectivity follows from $\hat{\mathbf{G}}_{U_{\mathrm{raw}}}\mathbb{Z}^{k-1} = \mathbb{Z}^{|U_{\mathrm{raw}}|}$ (Lemma 7) and the identity blocks in (45); see Appendix H.2. The block equations in (45) also give the integer kernel

$$\ker_{\mathbb{Z}}\left(\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}\right) = \left\{ \begin{bmatrix} \mathbf{v} \\ \mathbf{0} \\ -\mathbf{T}_{\mathrm{out}}\mathbf{v} \end{bmatrix} : \mathbf{v} \in \ker_{\mathbb{Z}}\left(\hat{\mathbf{G}}_{\mathrm{raw}}\right) \right\}, \tag{46}$$

so the $\mathcal{S}_{\mathrm{in}}$ blocks do not contribute to the kernel geometry. Finally, the target parameter matrix of the transformed perturbation is

$$\mathbf{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}} = \widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}\mathbf{\Sigma}_0\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}^\top. \tag{47}$$

For every nonempty auxiliary view, Lemma 29 implies that $\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}$ has full row rank; since $\mathbf{\Sigma}_0 \succ 0$ and $\mathbf{W}$ is invertible, both $\mathbf{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}$ and $\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}$ are positive definite.

5.4.3 Uniform Kernel-Smoothing Condition

For an integer lattice Λ, let $\widetilde{bl}(\Lambda)$ denote the minimum, over all ordered bases of Λ, of the largest Euclidean norm of a Gram-Schmidt vector, with $\widetilde{bl}(\{\mathbf{0}\}) := 0$. Define the

uniform kernel-geometry parameter and the effective perturbation width

$$\gamma_{k_p}^{\ker} := \max\left\{1,\ \max_{\substack{U\subseteq[n]\\ 1\le|U|\le k_p}} \widetilde{bl}\left(\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U)\right)\right\}, \tag{48}$$

$$\sigma_{\text{eff}} := \left(\frac{1}{\sigma_{\mathbf{R}}^2} + \frac{\omega_c^2 Q\|\hat{\mathbf{G}}\|_2^2}{\sigma_{\mathbf{r}}^2}\right)^{-1/2}. \tag{49}$$

For $0 < \epsilon' \le 1/2$, we impose the uniform kernel-smoothing condition

$$\sigma_{\text{eff}} \ge \gamma_{k_p}^{\ker}\sqrt{\frac{\ln\left(2(k-1)D\left(2+1/\epsilon'\right)\right)}{2\pi^2}}. \tag{50}$$

All quantities in (50) depend only on the public sharing matrix, the distribution parameters, and the global bounds k_p, ω_c, and Q.

Lemma 30 (Uniform whitened-kernel smoothing) *Assume* (50)*, and let* $\epsilon_0 := \epsilon'/(1+\epsilon')$*. Then, for every admissible fixed public realization* $(\mathcal{U}, \mathbf{c})$ *with a nonempty auxiliary view and* $\ker_{\mathbb{Z}}(\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}) \neq \{\mathbf{0}\}$*,*

$$\eta_{\epsilon_0}\left(\mathbf{\Sigma}_0^{-1/2}\ker_{\mathbb{Z}}\left(\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}\right)\right) \le 1. \tag{51}$$

The proof, in Appendix H.2, represents the whitened kernel as the image of $\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_{\text{raw}})$ under a linear map of norm at most $1/\sigma_{\text{eff}}$, bounds its Gram-Schmidt profile by $\gamma_{k_p}^{\ker}/\sigma_{\text{eff}}$, and applies the basis-dependent smoothing estimate of [40, Lemma 2.8]. To make $\gamma_{k_p}^{\ker}$ explicit, define

$$\beta_{k_p} := \max_{\substack{U\subseteq[n]\\ 1\le|U|\le k_p}} \|\hat{\mathbf{G}}_U\|_2, \tag{52}$$

$$\gamma_{k_p} := \min_{\substack{U\subseteq[n]\\ 1\le|U|\le k_p}} \sigma_{\min}(\mathbf{M}_U), \tag{53}$$

where $\mathbf{M}_U$ is the peeled unimodular submatrix of $\hat{\mathbf{G}}_U$ from Section 2.3.2. Since every $\mathbf{M}_U$ is invertible, $\gamma_{k_p} > 0$.

Lemma 31 (Explicit kernel basis and bound on $\gamma_{k_p}^{\ker}$) *Let* $U \subseteq [n]$ *satisfy* $1 \le |U| \le k_p$*, write* $r := k-1$*, and let* $\mathbf{M}_U = \hat{\mathbf{G}}_U[\cdot, J_U]$ *be the peeled unimodular submatrix, where* $|J_U| = |U|$*. Set* $t := r - |U|$*,* $\overline{J_U} := [r]\setminus J_U$*,* $\mathbf{N}_U := \hat{\mathbf{G}}_U[\cdot, \overline{J_U}]$*, and let* $\mathbf{P}_U$ *be the permutation matrix with* $\hat{\mathbf{G}}_U\mathbf{P}_U = \begin{bmatrix}\mathbf{M}_U & \mathbf{N}_U\end{bmatrix}$*. Then the vectors*

$$\mathbf{b}_a := \mathbf{P}_U\begin{bmatrix}-\mathbf{M}_U^{-1}\mathbf{N}_U\mathbf{e}_a\\ \mathbf{e}_a\end{bmatrix} \in \mathbb{Z}^r, \qquad a \in [t], \tag{54}$$

form a basis of $\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$. *Consequently,*

$$\gamma_{k_p}^{\ker} \le \sqrt{1 + \min\left\{k_p,\, \beta_{k_p}^2\right\}/\gamma_{k_p}^2}. \tag{55}$$

The basis is integral because $\mathbf{M}_U^{-1}$ is integral, and each $\|\mathbf{b}_a\|_2^2$ is at most $1 + \|\mathbf{N}_U \mathbf{e}_a\|_2^2/\sigma_{\min}^2(\mathbf{M}_U)$; the proof is given in Appendix H.2. Combining Lemma 31 with (50) gives the computable sufficient condition

$$\sigma_{\mathrm{eff}} \ge \sqrt{1 + \frac{\min\left\{k_p,\, \beta_{k_p}^2\right\}}{\gamma_{k_p}^2}} \sqrt{\frac{\ln\left(2(k-1)D\left(2 + 1/\epsilon'\right)\right)}{2\pi^2}}. \tag{56}$$

This is the only point of the privacy analysis at which the peeling procedure enters quantitatively; the values of β_{k_p} and γ_{k_p} for our instances are reported in Table 4.

5.4.4 Uniform Gaussian Approximation

Lemma 32 (Uniform Gaussian approximation of the joint perturbation) *Suppose that* $\sigma_{\mathbf{R}}, \sigma_{\mathbf{r}} > 0$ *and that* (50) *holds. Then, for every admissible fixed public realization* $(\mathcal{U}, \mathbf{c})$ *inducing a nonempty auxiliary view,* $\Delta(\mathbf{y}_{\mathcal{U},\mathbf{c}}, \mathcal{D}_{\mathbb{Z}^{M'}, \sqrt{\boldsymbol{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}}}) \le \epsilon'$.

Proof sketch If $\ker_{\mathbb{Z}}(\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}}) = \{\mathbf{0}\}$, then integral surjectivity makes $\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}}$ unimodular, so the pushforward is exact. Otherwise, Lemmas 29 and 30 and the linear-transformation theorem of [40, Theorem 3.1] give statistical distance at most $\epsilon_0/(1-\epsilon_0) = \epsilon'$. Applying the bijection $\mathbf{W}^{-1}$ preserves this distance. The details are given in Appendix H.2. □

In the fixed-pattern HC-Leak experiment, replacing the actual perturbation by its target elliptic discrete Gaussian therefore costs at most ϵ' in each branch and at most $2\epsilon'$ in total; the empty view incurs no loss.

5.4.5 Uniform Matrix Hint-MLWE Source Parameter

Suppose that the coefficient-embedded secret has parameter matrix $\sigma_{\mathbf{s}}^2 \mathbf{I}_D$. For a fixed nonempty realization $(\mathcal{U}, \mathbf{c})$, the effective source width of Theorem 3 and its uniform replacement are

$$\sigma_{\mathcal{U},\mathbf{c}} := \left(2\left(\frac{1}{\sigma_{\mathbf{s}}^2} + \left\|\mathbf{M}_{\mathcal{U},\mathbf{c}}^{\top} \boldsymbol{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}^{-1} \mathbf{M}_{\mathcal{U},\mathbf{c}}\right\|_2\right)\right)^{-1/2}, \tag{57}$$

$$\sigma_{\mathrm{MQ}} := \left(2\left(\frac{1}{\sigma_{\mathbf{s}}^2} + \Xi_{\mathrm{MQ}}\right)\right)^{-1/2}. \tag{58}$$

By Lemma 28, $\sigma_{\mathcal{U},\mathbf{c}} \ge \sigma_{\mathrm{MQ}}$ for every admissible nonempty realization. Hence, if $\sigma_{\mathrm{MQ}} \ge \sqrt{2}\eta_\epsilon(\mathbb{Z}^D)$ for some $0 < \epsilon \le 1/2$, Corollary 1 lets every such instance be reduced from $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathrm{MQ}}}$ with loss 2ϵ. If $U_{\mathrm{raw}} = \emptyset$ and $\mathcal{S} = \emptyset$, then the auxiliary view is empty and contains no hint; this case is handled directly under $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathbf{s}}}$.

5.5 Multi-Query HC-LLR Privacy

5.5.1 Fixed-Pattern Reduction to MLWE

Fix an admissible public realization $(\mathcal{U}, \mathbf{c})$ inducing a nonempty auxiliary view. Define the Gaussianized hint and perturbation

$$\check{\mathbf{z}}_{\mathcal{U},\mathbf{c}} := \mathbf{M}_{\mathcal{U},\mathbf{c}}\mathbf{s} + \check{\mathbf{y}}_{\mathcal{U},\mathbf{c}}, \qquad \check{\mathbf{y}}_{\mathcal{U},\mathbf{c}} \leftarrow \mathcal{D}_{\mathbb{Z}^{M'},\sqrt{\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}}} \tag{59}$$

where $\check{\mathbf{y}}_{\mathcal{U},\mathbf{c}}$ is sampled independently of the secret.

Lemma 33 (Fixed-pattern reduction to Matrix Hint-MLWE) *Suppose that the hypotheses of Lemma 32 hold for some $0 < \epsilon' \le 1/2$. Then, for every PPT adversary $\mathcal{A}$ against the fixed-pattern* HC-Leak *experiment for $(\mathcal{U}, \mathbf{c})$, there exists a PPT adversary $\mathcal{B}$ against the one-hint Matrix Hint-MLWE instance with hint matrix $\mathbf{M}_{\mathcal{U},\mathbf{c}}$ and hint $\check{\mathbf{z}}_{\mathcal{U},\mathbf{c}}$ in (59) such that $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U},\mathbf{c})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}^{\mathsf{MHMLWE}}_{\mathcal{B}}(\kappa) + 2\epsilon'$.*

Lemma 34 (Fixed-pattern reduction to MLWE) *Suppose that the hypotheses of Lemmas 28 and 32 hold and that $\sigma_{\mathrm{MQ}} \ge \sqrt{2}\eta_\epsilon(\mathbb{Z}^D)$ for some $0 < \epsilon \le 1/2$. Then, for every PPT adversary $\mathcal{A}$ against the fixed-pattern* HC-Leak *experiment for $(\mathcal{U}, \mathbf{c})$, there exists a PPT adversary $\mathcal{B}$ against $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathrm{MQ}}}$ such that $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U},\mathbf{c})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}_{\mathcal{B},\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathrm{MQ}}}}(\kappa) + 2\epsilon' + 2\epsilon$.*

The first lemma replaces the actual joint perturbation by its target elliptic discrete Gaussian, parses $\check{\mathbf{z}}_{\mathcal{U},\mathbf{c}}$ according to (14), and loses at most ϵ' in each branch by Lemma 32. The second composes this reduction with Corollary 1, using $\sigma_{\mathrm{MQ}} \le \sigma_{\mathcal{U},\mathbf{c}}$, and incurs the additional loss 2ϵ. The proofs are given in Appendix H.3.

5.5.2 Admissible Multi-Query Leakage Patterns

Define $\mathcal{U}_{k_p,Q} := \big\{\big(U_{\mathrm{raw}}, (V_s)_{s\in[Q]}\big) : U_{\mathrm{raw}} \in \Gamma_0,\ |U_{\mathrm{raw}}| \le k_p,\ V_s \subseteq [n]\ \forall s \in [Q]\big\}$. No disjointness is required between U_{raw} and the sets V_s, the same party may occur in several queries, and empty sets are permitted, including the completely empty auxiliary view. For every $\mathbf{c} \in \mathcal{C}^Q$, the induced local maps satisfy $\|\mathbf{L}_s\|_2 \le \omega_c$, so Lemma 34 applies pointwise to every $(\mathcal{U}, \mathbf{c})$ with a nonempty view, with the same source width σ_{MQ}.

5.5.3 Expander Setup Event

Let $\mathsf{Good}_{\mathrm{MQ}}(\mathsf{pp})$ be the event that, for every nonempty $U \subseteq [n]$ with $|U| \le k_p$, the matrix $\hat{\mathbf{G}}_U$ contains a unimodular $|U| \times |U|$ column submatrix, and let

$$\delta_{\mathrm{MQ}} := \Pr_{\mathsf{pp}\leftarrow\mathsf{Setup}(1^\kappa)}\left[\neg\mathsf{Good}_{\mathrm{MQ}}(\mathsf{pp})\right]. \tag{60}$$

On $\mathsf{Good}_{\mathrm{MQ}}(\mathsf{pp})$, every such $\hat{\mathbf{G}}_U$ is integrally surjective and has full row rank, so the uniform conditioning, integral-surjectivity, and kernel-geometry conclusions of

Section 5.4 hold simultaneously over $\mathcal{U}_{k_p,Q}$. The probability δ_{MQ} is taken only over the setup output; the conditions (50) and $\sigma_{\mathrm{MQ}} \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^D)$ are deterministic, and their statistical errors ϵ' and ϵ are accounted for separately. For the certified expander-based construction, Lemma 4 shows that $\mathsf{Good}_{\mathrm{MQ}}(\mathsf{pp})$ holds for every setup output, so $\delta_{\mathrm{MQ}} = 0$.

5.5.4 Multi-Query Privacy Theorem

For fixed pp and $\mathcal{U} \in \mathcal{U}_{k_p,Q}$, write $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa)$ for the advantage in the experiment in which the leakage-description sequence is sampled internally as $\mathbf{c} \leftarrow \mathcal{C}^Q$. The theorem instantiates Definition 10 with coalitions replaced by patterns $\mathcal{U} \in \mathcal{U}_{k_p,Q}$; its exceptional setup outputs are those violating $\mathsf{Good}_{\mathrm{MQ}}$, which depends neither on the adversary nor on the pattern.

Let $\mathcal{F}_{\mathrm{MQ}} := \{f_{\boldsymbol{A}} : \boldsymbol{s} \mapsto \boldsymbol{A}\boldsymbol{s} \bmod q \mid \boldsymbol{A} = \left[\boldsymbol{I}_p \ \boldsymbol{A}'\right],\ \boldsymbol{A}' \in \mathcal{R}_q^{p\times(\ell-p)}\}$, with $\boldsymbol{A}'$ uniform. For each κ, let Good_κ be the set of setup outputs pp in the support of $\mathsf{Setup}(1^\kappa)$ satisfying $\mathsf{Good}_{\mathrm{MQ}}(\mathsf{pp})$. For a PPT adversary $\mathcal{A}$, a pattern $\mathcal{U} \in \mathcal{U}_{k_p,Q}$, and $\mathsf{pp} \in \mathsf{Good}_\kappa$, let $\mathcal{B}_{\mathsf{pp}}$ be obtained from $\mathcal{A}$, $\mathcal{U}$, and pp as follows: if the auxiliary view of $\mathcal{U}$ is nonempty, $\mathcal{B}_{\mathsf{pp}}$ samples $\mathbf{c} \leftarrow \mathcal{C}^Q$ and runs the reduction of Lemma 34 against $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathrm{MQ}}}$; otherwise, $\mathcal{B}_{\mathsf{pp}}$ is the direct reduction to $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathbf{s}}}$ given in the proof in Appendix H.3. Define

$$\epsilon_{\mathrm{MQ}}(\kappa,\mathcal{A},\mathcal{U}) := \max_{\mathsf{pp}\in\mathsf{Good}_\kappa} \mathsf{Adv}_{\mathcal{B}_{\mathsf{pp}},\mathsf{MLWE}}(\kappa) + 2\epsilon' + 2\epsilon, \tag{61}$$

where the maximum is over a finite set and is taken to be 0 if $\mathsf{Good}_\kappa = \emptyset$. Since $\mathcal{B}_{\mathsf{pp}}$ uses pp only as polynomial-size advice, ϵ_{MQ} is negligible whenever the relevant MLWE problems are hard against non-uniform PPT adversaries and $2\epsilon' + 2\epsilon$ is negligible.

Theorem 4 (Multi-query HC-LLR privacy from MLWE) *Suppose that the following hold:*

1. *the secret, sharing masks, and response masks satisfy the distributional and independence conditions of Section 5.2;*
2. *(50) holds for some $0 < \epsilon' \leq 1/2$, and $\sigma_{\mathrm{MQ}} \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^D)$ for some $0 < \epsilon \leq 1/2$, where σ_{MQ} is defined in (58);*
3. *$\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathrm{MQ}}}$ and $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathbf{s}}}$ are hard against non-uniform PPT adversaries, and $2\epsilon' + 2\epsilon$ is negligible in κ.*

Then, for every κ, with δ_{MQ} and ϵ_{MQ} as in (60) and (61), respectively,

$$\Pr_{\mathsf{pp}\leftarrow\mathsf{Setup}(1^\kappa)}\left[\forall \text{PPT } \mathcal{A},\ \forall \mathcal{U} \in \mathcal{U}_{k_p,Q} : \mathsf{Adv}^{(\mathsf{pp},\mathcal{U})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \leq \epsilon_{\mathrm{MQ}}(\kappa,\mathcal{A},\mathcal{U})\right] \geq 1 - \delta_{\mathrm{MQ}}.$$

That is, except with probability at most δ_{MQ} over $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)$, every PPT adversary $\mathcal{A}$ has advantage at most $\epsilon_{\mathrm{MQ}}(\kappa,\mathcal{A},\mathcal{U})$ for every admissible pattern $\mathcal{U}$. Consequently, $\mathcal{F}_{\mathrm{MQ}}$ is $(\delta_{\mathrm{MQ}},\epsilon_{\mathrm{MQ}})$-private for SEC in the multi-query $\mathsf{HC\text{-}LLR}$ sense of Definition 10, and hence δ_{MQ}-private because $\epsilon_{\mathrm{MQ}}(\kappa,\mathcal{A},\mathcal{U})$ is negligible in κ for every PPT adversary $\mathcal{A}$ and every admissible pattern $\mathcal{U}$. For the certified expander-based construction, $\delta_{\mathrm{MQ}} = 0$.

Proof sketch Fix $\mathsf{pp} \in \mathsf{Good}_\kappa$, a PPT adversary $\mathcal{A}$, and a pattern $\mathcal{U} \in \mathcal{U}_{k_p,Q}$. If the auxiliary view of $\mathcal{U}$ is empty, $\mathcal{B}_{\mathsf{pp}}$ directly distinguishes $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathbf{s}}}$, so $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}_{\mathcal{B}_{\mathsf{pp}},\mathsf{MLWE}}(\kappa)$. Otherwise, $\mathsf{Good}_{\mathrm{MQ}}(\mathsf{pp})$ ensures that Lemma 34 applies uniformly over $\mathbf{c} \in \mathcal{C}^Q$, with source width σ_{MQ} and loss $2\epsilon' + 2\epsilon$, so $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}_{\mathcal{B}_{\mathsf{pp}},\mathsf{MLWE}}(\kappa) + 2\epsilon' + 2\epsilon$. In both cases, the advantage is at most $\epsilon_{\mathrm{MQ}}(\kappa,\mathcal{A},\mathcal{U})$ by its definition as a maximum over Good_κ. The result follows because Good_κ has probability $1-\delta_{\mathrm{MQ}}$ and depends neither on $\mathcal{A}$ nor on $\mathcal{U}$. The complete proof is given in Appendix H.3. □

The two MLWE assumptions correspond to different cases: σ_{MQ} applies uniformly to every nonempty auxiliary view, whereas the empty view uses the original secret width $\sigma_{\mathbf{s}}$. By (58), $\sigma_{\mathrm{MQ}} \le \sigma_{\mathbf{s}}/\sqrt{2}$. The theorem states both assumptions separately, but for our concrete instantiations it suffices in practice to evaluate hardness at σ_{MQ}; see Section 6.4.3.

6 Parameters and Evaluation

We empirically evaluate concrete SEC distribution matrices for party counts $n \in \{10, 15, \ldots, 60\}$. Our purpose is not to characterize random expander-code instances, but to identify concrete matrices satisfying the structural, privacy, reconstruction, and shortness properties established in Sections 3 to 5. We report all 17 selected matrices, including multiple selected matrices for some party counts. Among the selected matrices with $n \ge 25$, the $n = 40$ instance achieves the largest observed ratio $\tau_p/\tau_c \approx 0.192$, while the $n = 60$ instance is the largest configuration considered. We use these two instances for the concrete privacy instantiations. One should be able to obtain results for larger n values by expending more computation.

6.1 Search and Certification Methodology

6.1.1 Search Space

For each party count, we consider masking dimensions in the regimes $m \approx n/2$ and $m \approx 2n/3$, whenever applicable. For each pair (n, m), we consider candidate parameter tuples $(n, m, c_{\mathrm{exp}}, \gamma, \epsilon_{\mathrm{exp}})$, where $\epsilon_{\mathrm{exp}} \in (0, 1/2)$ and $\gamma \ge 0.1$. We search increasing values of c_{exp} and retain the smallest left degree for which the search finds a suitable certified lossless-expander instance. The selected instances use $c_{\mathrm{exp}} \in \{2, 3\}$.

The principal constraint is obtaining a sufficiently large certified value of γ, since the certified expansion range determines the private-coalition threshold $k_p := \lceil 2(1-\epsilon_{\mathrm{exp}})\lfloor \gamma n \rfloor \rceil - 1$. For every reported instance, γ is chosen so that $\gamma n \in \mathbb{Z}$, and the underlying full-precision parameters satisfy $\lceil 2(1-\epsilon_{\mathrm{exp}})\lfloor \gamma n \rfloor \rceil - 1 = \lfloor \gamma n \rfloor$. Hence $k_p = \gamma n$ and $\tau_p := k_p/n = \gamma$.

The values of γ, ϵ_{exp}, α_{k_p}, $\|\hat{\mathbf{G}}\|_2$, and $\|\mathbf{G}\|_2$ are reported to three decimal places. Expansion verification and the computation of k_p use the underlying full-precision parameters rather than the rounded table entries.

6.1.2 Expander Sampling and Certification

For each candidate tuple, we sample a $c_{\exp}$-left-regular bipartite graph with left vertices $L = [n]$ and right vertices $R = [m]$ by letting each left vertex independently choose a uniformly random $c_{\exp}$-subset of R, so that $\mathbf{H} \in \{0,1\}^{m\times n}$ has exactly $c_{\exp}$ nonzero entries in each column. Every sampled candidate is then certified by exhaustive verification, which the regime $n \le 60$ permits: for every $S \subseteq L$ with $|S| \le \lfloor \gamma n \rfloor$, we verify $|N(S)| \ge (1-\epsilon_{\exp})c_{\exp}|S|$, and only candidates passing every check are selected. Thus, the concrete matrices rely on exhaustive certification rather than Theorem 2, and Lemma 4 converts the certified expansion property into the unimodular submatrices required by the structural privacy condition.

6.1.3 Construction of the Distribution Matrix

For each selected $\mathbf{H}$, we construct $\mathbf{G} = \begin{bmatrix}\mathbf{g}\ \hat{\mathbf{G}}\end{bmatrix}$ as in Definition 11, with $\mathbf{g}$ from Algorithm 2; for every selected instance, $\|\mathbf{g}\|_\infty = 1$, and hence $\beta_{\mathbf{g}} = 1$. Among the candidates satisfying the certified expansion and generator-matrix construction requirements, we select representative instances based primarily on their empirical reconstruction thresholds and observed reconstruction-vector magnitudes, while also accounting for their privacy-side matrix statistics. For some party counts, more than one instance is selected to compare different choices of m or $c_{\exp}$.

6.2 Structural and Reconstruction Parameters

6.2.1 Selected Instances and Privacy-Side Parameters

The privacy-side statistics of all selected instances are reported in Table 4. The threshold $\hat{d}^{\perp} - 1$ is the certified rational no-cancellation threshold associated with the code having parity-check matrix $\mathbf{PH}$ (Lemma 19). The quantities α_{k_p} and $\|\hat{\mathbf{G}}\|_2$ enter the uniform precision parameter Ξ_{MQ} in (43), while β_{k_p} and γ_{k_p} enter the explicit kernel-smoothing condition (56). The quantities α_{k_p}, β_{k_p}, and γ_{k_p} are evaluated exhaustively over their subset domains.

To compare the integral and rational no-cancellation thresholds independently of $\mathbf{g}$, let $\hat{d}^{\perp}_{\mathbb{Z}} := \min_p \hat{d}^{\perp}_p$, where $\hat{d}^{\perp}_p$ is the minimum distance of $\hat{C}^{\perp}_p := \ker_{\mathbb{F}_p}(\overline{\mathbf{H}})$, and the minimum ranges over all primes p. By the Smith normal form, $\mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U) = \mathbb{Z}^{|U|}$ for every $|U| \le \hat{d}^{\perp}_{\mathbb{Z}} - 1$. Moreover, $k_p \le \hat{d}^{\perp}_{\mathbb{Z}} - 1 \le \hat{d}^{\perp}_{\mathbb{Q}} - 1$, where the first inequality follows from Lemma 7 and the second from Lemma 1. Thus, a strict gap between k_p and $\hat{d}^{\perp} - 1$ may reflect both slack in the expander-based certificate and a genuine gap between the integral and rational no-cancellation thresholds. For coalitions satisfying $k_p < |U| \le \hat{d}^{\perp} - 1$, rational cancellation remains impossible, but the expander certificate no longer guarantees resistance to modular cancellation. If $(\mathbf{g})_U \notin \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$, such a coalition admits a modular noise-cancellation attack by Lemma 42. Coalitions beyond k_p lie outside the stated privacy guarantee.

Table 4: Selected SEC instances and their privacy-side matrix statistics

No.	n	$k = m + 1$	$c_{\exp}$	$\gamma\ (= \tau_p)$	$\epsilon_{\exp}$	k_p	$\hat{d}^{\perp} - 1$	α_{k_p}	$\|\hat{\mathbf{G}}\|_2$	β_{k_p}	γ_{k_p}
1	10	6	2	0.200	0.333	2	3	1.000	2.828	1.732	0.618
2	10	6	3	0.200	0.333	2	3	1.000	4.243	2.236	0.618
3	15	8	3	0.200	0.450	3	3	0.792	4.440	2.646	0.445
4	20	10	2	0.150	0.333	3	3	0.765	3.000	2.000	0.445
5	20	10	3	0.150	0.490	3	3	0.792	4.634	2.646	0.445
6	25	13	2	0.120	0.333	3	3	0.765	3.143	2.000	0.445
7	25	13	3	0.120	0.450	3	3	0.792	4.433	2.646	0.445
8	30	15	2	0.100	0.333	3	3	0.765	3.298	2.000	0.445
9	30	15	3	0.100	0.450	3	3	0.792	4.628	2.646	0.445
10	30	20	2	0.133	0.450	4	5	0.618	2.769	2.236	0.347
11	30	20	3	0.133	0.450	4	5	0.618	4.053	2.665	0.295
12	35	18	3	0.114	0.450	4	5	0.618	4.459	2.665	0.295
13	40	20	3	0.125	0.480	5	5	0.464	4.432	2.940	0.193
14	45	23	3	0.111	0.470	5	5	0.464	4.427	2.868	0.193
15	50	25	3	0.100	0.470	5	5	0.464	4.529	2.940	0.193
16	55	37	3	0.109	0.450	6	≥ 8	0.445	3.863	2.893	0.152
17	60	40	3	0.100	0.450	6	≥ 8	0.518	3.937	2.893	0.152

Note: the entries $\hat{d}^{\perp} - 1 \geq 8$ are certified lower bounds. Here $\beta_{k_p} := \max_U \|\hat{\mathbf{G}}_U\|_2$ and $\gamma_{k_p} := \min_U \sigma_{\min}(\mathbf{M}_U)$ over $1 \leq |U| \leq k_p$, where $\mathbf{M}_U$ is the peeled unimodular submatrix.

6.2.2 Reconstruction Thresholds and Certification Probability

Deterministic Reconstruction Threshold

A deterministic correctness threshold k_c^{det} satisfies $\mathbf{e}_1^{\top} \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)$ for every $T \subseteq [n]$ with $|T| \geq k_c^{\mathsf{det}}$. Since certification is preserved under enlarging a coalition, it is sufficient to verify this condition for subsets of size exactly k_c^{det}. In the empirical search, the deterministic threshold is obtained either from the planted certified supports through Lemma 8 or by exhaustive verification whenever the number of relevant subsets is sufficiently small. In the exhaustive case, we report the smallest verified threshold. When the planted-support argument is used, the reported value is a certified deterministic threshold but need not be the smallest possible one.

Certification Probability and Empirical Correctness

Definition 12 (Certification probability and average-case correctness) Fix public parameters pp containing $\mathbf{G} \in \mathbb{Z}^{n \times k}$. For $t \in [n]$, define

$$\Pr_{\mathrm{cert}}(t, \mathbf{G}) := \Pr_{T \xleftarrow{\$} \binom{[n]}{t}} \left[\mathbf{e}_1^{\top} \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T) \right]. \tag{62}$$

For $k_c \in [n]$ and $\epsilon_c \in [0, 1]$, the fixed instance $(\mathsf{SEC}, \mathsf{pp})$ is ϵ_c-correct for random coalitions above threshold k_c if $\Pr_{\mathrm{cert}}(t, \mathbf{G}) \geq 1 - \epsilon_c$ for every $t \geq k_c$.

Thus, $\Pr_{\mathrm{cert}}(t, \mathbf{G})$ is the fraction of size-t coalitions certified by the integer row-span condition, and it is nondecreasing in t because certification is monotone. We compute it exactly when $\binom{n}{t} \le 10^6$, and otherwise estimate it from N_t distinct uniformly sampled subsets, with $10^5 \le N_t \le 10^6$ and $N_t = 10^6$ near $t = n/2$. The empirical threshold k_c is the smallest tested value for which the exact or estimated probability is at least 0.99 for every tested $t \ge k_c$, that is, $\epsilon_c \le 0.01$; when sampling is used, k_c is an empirical estimate rather than an exhaustive certification.

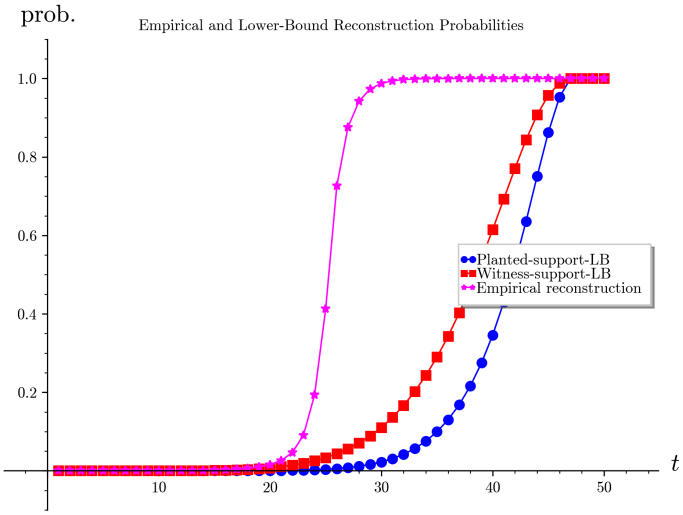
Figure 1 compares the exact or estimated certification probability of the selected instances with $n = 50$ and $n = 60$ with the planted-support and witness-support lower bounds of Corollaries 2 and 3, which describe only sufficient conditions for certification. For $n = 50$, the empirical threshold is $k_c = 31$, whereas the deterministic threshold reported in Table 5 is $k_c^{\mathsf{det}} = 45$. For $n = 60$, the corresponding thresholds are $k_c = 46$ and $k_c^{\mathsf{det}} = 56$. These results illustrate that random coalitions typically reconstruct substantially earlier than the conservative worst-case deterministic threshold.

To examine how the reconstruction failure probability decreases as the coalition size exceeds the code dimension, let $p_{c,e}(t) := 1 - \Pr_{\mathrm{cert}}(t, \mathbf{G})$, and let $\Delta_c := t - k$ denote the excess coalition size over $k = m+1$. Figure 2 plots $\lg p_{c,e}$ against $\Delta_c \in [0, 15]$ for the same two instances. For both instances, $\lg p_{c,e}$ is approximately linear in Δ_c: a least-squares fit over $0 \le \Delta_c \le 12$ gives slopes of -1.15 for $n = 50$ and -1.11 for $n = 60$, so each additional party reduces the failure probability by a factor of approximately $2^{1.1}$. The two curves nearly coincide and remain close to $2^{-(\Delta_c+1)}$, suggesting that, for these instances, the failure probability depends primarily on Δ_c, rather than separately on n or k. This empirical tail is consistent with the observed thresholds: the failure probability first drops below $\epsilon_c = 0.01$ at $\Delta_c = 6$, giving $k_c = k + 6$ in both cases. It also suggests that each halving of ϵ_c requires roughly one additional party, until k_c^{det} guarantees reconstruction for every coalition. The points with $\Delta_c \ge 13$ are based on only a few observed failures among at most 10^6 samples and are therefore excluded from the fit.

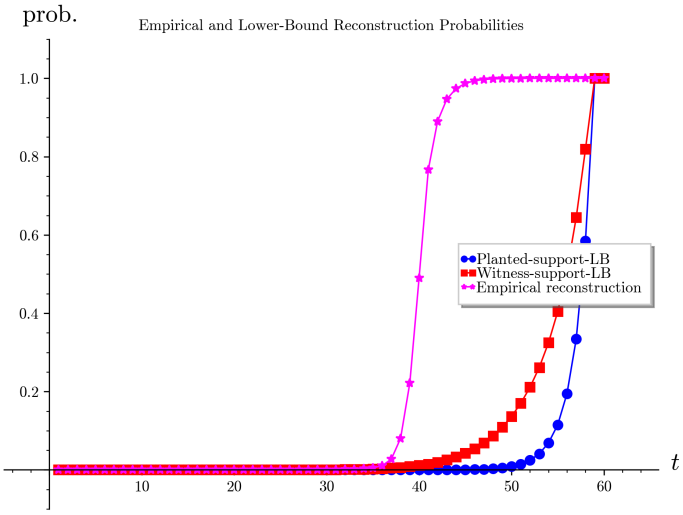
6.3 Reconstruction-Coefficient Evaluation

6.3.1 Reconstruction-Vector Statistics

The reconstruction-vector experiment considers coalition sizes $t \ge k_c$, corresponding to the empirical high-certification regime of Section 6.2.2. Computing a reconstruction vector requires the row-HNF of $\mathbf{G}_T$, LLL reduction of a basis of $\ker_{\mathbb{Z}}(\mathbf{G}_T^\top)$, and the decoding and refinement steps of Algorithm 1, which is considerably more expensive than the row-span test used for the certification probability in Section 6.2.2. We therefore use a smaller sample per coalition size: for each tested size t, let $\mathcal{T}_{\mathrm{test}}^t(\mathbf{G})$ be the set of tested size-t subsets, obtained by enumerating all size-t subsets when $\binom{n}{t} \le 30000$, and otherwise by sampling at most 30000 distinct size-t subsets uniformly without replacement. Every tested subset T is first checked for certification, and only certified subsets contribute reconstruction vectors; when sampling is used, the number contributing at a given size may therefore be smaller than 30000. For each certified subset, we apply Algorithm 1 with $\delta_{\mathsf{LLL}} = 0.99$, obtaining a vector $\boldsymbol{\lambda}_T$ that satisfies $\boldsymbol{\lambda}_T^\top \mathbf{G}_T = \mathbf{e}_1^\top$ by the exact certification step. We report $\hat{\beta}_{\boldsymbol{\lambda}} := \max_T \|\boldsymbol{\lambda}_T\|_\infty$



(a) $n = 50$



(b) $n = 60$

Fig. 1: Empirical certification probability and the planted- and witness-support lower bounds for the selected instances with $n = 50$ and $n = 60$.

and $\hat{\omega}_{\boldsymbol{\lambda}} := \max_T wt_H(\boldsymbol{\lambda}_T)$, where both maxima range over the tested certified subsets with $|T| \geq k_c$, not over every certified coalition of size at least k_c. The reconstruction statistics of all selected instances are reported in Table 5.

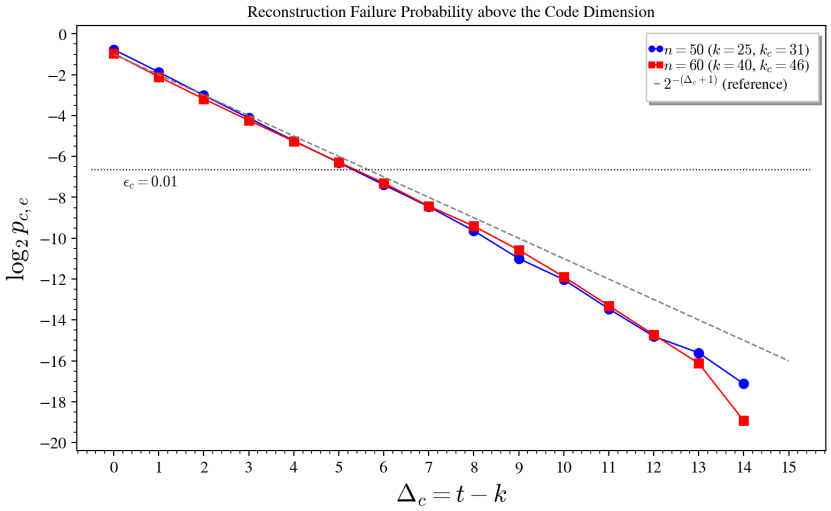


Fig. 2: Reconstruction failure probability $p_{c,e}(t) = 1 - \Pr_{\text{cert}}(t, \mathbf{G})$ against the excess coalition size $\Delta_c = t - k$ for the selected instances with $n = 50$ and $n = 60$. The dashed line is $2^{-(\Delta_c+1)}$, and the dotted line marks $\epsilon_c = 0.01$.

Table 5: Deterministic and empirical reconstruction statistics for the selected SEC instances

No.	n	$k = m + 1$	k_c^{det}	k_c	$\tau_c = k_c/n$	$\hat{\beta}_{\boldsymbol{\lambda}}$	$\hat{\omega}_{\boldsymbol{\lambda}}$
1	10	6	9	8	0.800	1	4
2	10	6	9	8	0.800	1	4
3	15	8	13	12	0.800	1	6
4	20	10	17	14	0.700	1	6
5	20	10	16	14	0.700	2	12
6	25	13	21	16	0.640	1	10
7	25	13	22	18	0.720	4	17
8	30	15	25	19	0.633	1	8
9	30	15	26	20	0.667	3	17
10	30	20	28	25	0.833	1	12
11	30	20	28	25	0.833	7	25
12	35	18	31	24	0.686	4	22
13	40	20	35	26	0.650	4	24
14	45	23	41	29	0.644	7	27
15	50	25	45	31	0.620	4	28
16	55	37	51	43	0.782	6	39
17	60	40	56	46	0.767	8	42

Note: $\hat{\beta}_{\boldsymbol{\lambda}}$ and $\hat{\omega}_{\boldsymbol{\lambda}}$ are observed maxima over the tested certified subsets of sizes at least k_c.

6.3.2 Observed Reconstruction-Coefficient Bounds

To examine how representative the largest observed coefficient is, we also consider the empirical distribution of $\|\boldsymbol{\lambda}_T\|_\infty$. For each coalition size t, let $\mathcal{T}^t_{\mathrm{cert}}(\mathbf{G}) := \{T \in \mathcal{T}^t_{\mathrm{test}}(\mathbf{G}) : \mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_T)\}$ denote the tested certified subsets of size t. For a range $I \subseteq [k_c, n]$, define $F^I_\infty(b) := \sum_{t\in I} |\{T \in \mathcal{T}^t_{\mathrm{cert}}(\mathbf{G}) : \|\boldsymbol{\lambda}_T\|_\infty = b\}|$ and $\widehat{F}^I_\infty(b) := \frac{F^I_\infty(b)}{\sum_{b'\geq 0} F^I_\infty(b')}$. Thus, $\widehat{F}^I_\infty$ is the pooled empirical distribution of $\|\boldsymbol{\lambda}_T\|_\infty$ over the tested certified subsets whose sizes lie in I.

For the selected $n = 60$ instance, we report the range $I = [46, 50]$, which begins at the empirical threshold $k_c = 46$, in Figure 3. For this instance, $\|\boldsymbol{\lambda}_T\|_\infty \leq 2$ for 99.18% of the tested certified subsets, and the distribution shifts toward smaller values as the coalition size increases, consistent with the additional freedom provided by larger reconstructing coalitions. In particular, $\|\boldsymbol{\lambda}_T\|_\infty \leq 2$ for all observed T such that $|T| \in [51, 60]$. The observed maximum $\hat{\beta}_{\boldsymbol{\lambda}} = 8$ is attained by a single subset, of the smallest size $|T| = 46$.

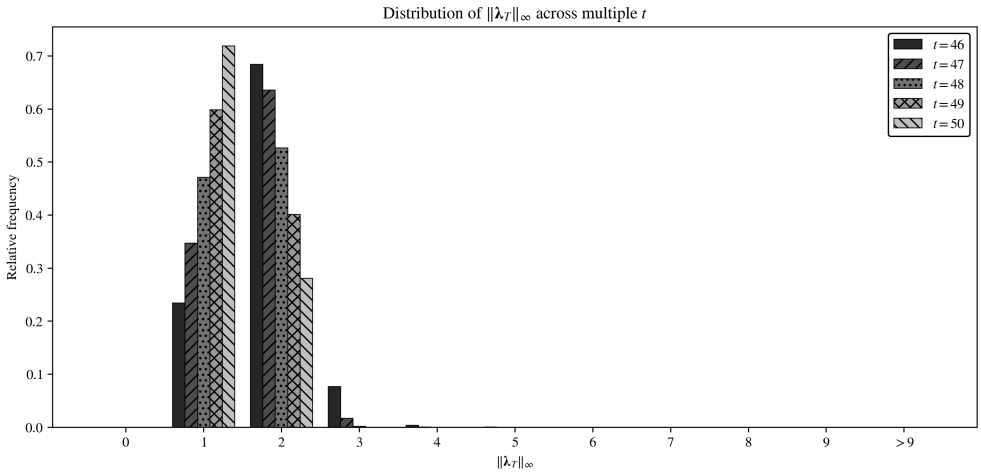


Fig. 3: Pooled relative-frequency distribution of the observed values of $\|\boldsymbol{\lambda}_T\|_\infty$ over the tested certified subsets with $|T| \in [46, 50]$ for the selected $n = 60$ instance.

The $n = 40$ instance is even more concentrated: $\|\boldsymbol{\lambda}_T\|_\infty \leq 2$ for 99.99% of the tested certified subsets, and the observed maximum $\hat{\beta}_{\boldsymbol{\lambda}} = 4$ is attained by only three subsets, all of the smallest size $|T| = 26$.

We therefore set $\beta_{\boldsymbol{\lambda}} := \hat{\beta}_{\boldsymbol{\lambda}} + 1$, that is, $\beta_{\boldsymbol{\lambda}} = 5$ for $n = 40$ and $\beta_{\boldsymbol{\lambda}} = 9$ for $n = 60$. For the Hamming weight, which is trivially bounded by $|T| \leq n$, we use $\omega_{\boldsymbol{\lambda}} = 25$ and $\omega_{\boldsymbol{\lambda}} = 45$, slightly above the observed maxima $\hat{\omega}_{\boldsymbol{\lambda}} = 24$ and 42, respectively. If an application enforces both bounds, as discussed in Section 4.2.2, they become additional admissibility conditions for authorized sets; none of the tested subsets is rejected under the selected values.

6.4 Concrete Privacy Instantiations

6.4.1 Gaussian Tail Parameters

The share bound of Section 4.2.1 and the requirements below use explicit tail bounds for the three families of Gaussian samples in SEC: the secret $\boldsymbol{s}$, the masking matrix $\mathbf{R}$, and the leakage masks $\boldsymbol{r}_{i,s}$. We use the following discrete Gaussian tail bound [61, Lemma 4.4], stated for the standard-deviation convention of (1), together with an explicit sufficient condition on the tail factor. We use $\sigma_{\boldsymbol{s}} = \sigma_{\mathbf{s}}$ and $\beta_{\boldsymbol{s}} = \beta_{\mathbf{s}}$, since $\boldsymbol{s} \in \mathcal{R}^{\ell}$ and its coefficient embedding $\mathbf{s} \in \mathbb{Z}^D$ represent the same secret. The choice of notation indicates only the representation.

Lemma 35 (Discrete Gaussian tail bound) *Let $\sigma > 0$, let $N \geq 1$, and let $\mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^N,\sigma}$. For $\theta > 1$, define $C_{\mathrm{tail}}(\theta) := \theta \exp\left(\frac{1-\theta^2}{2}\right)$. Then $C_{\mathrm{tail}}(\theta) < 1$ and*

$$\Pr\left[\|\mathbf{x}\|_2 > \theta\sigma\sqrt{N}\right] < C_{\mathrm{tail}}(\theta)^N. \tag{63}$$

Moreover, for any $\kappa_{\mathrm{tail}} > 0$, if $\theta \geq 1 + \sqrt{2\kappa_{\mathrm{tail}} \ln 2/N}$, then $C_{\mathrm{tail}}(\theta)^N \leq 2^{-\kappa_{\mathrm{tail}}}$.

The second claim follows from $-\ln C_{\mathrm{tail}}(1+x) \geq x^2/2$ for $x > 0$; the proof is given in Appendix I. Fix the security parameter κ. For the secret, we choose

$$\theta_{\boldsymbol{s}} \geq 1 + \sqrt{\frac{2\kappa \ln 2}{D}}. \tag{64}$$

For the masking matrix, let $\kappa_{\mathbf{R}} := \kappa + \lceil \log_2((k-1)D) \rceil$ and choose

$$\theta_{\mathbf{R}} \geq 1 + \sqrt{2\kappa_{\mathbf{R}} \ln 2}. \tag{65}$$

For the leakage masks, when $Q \geq 1$, let $\kappa_{\boldsymbol{r}} := \kappa + \lceil \log_2(nQ) \rceil$ and choose

$$\theta_{\boldsymbol{r}} \geq 1 + \sqrt{\frac{2\kappa_{\boldsymbol{r}} \ln 2}{D}}. \tag{66}$$

We define $\beta_{\boldsymbol{s}} := \theta_{\boldsymbol{s}}\sigma_{\boldsymbol{s}}\sqrt{D}$ and $\beta_{\mathbf{R}} := \theta_{\mathbf{R}}\sigma_{\mathbf{R}}$. If $Q \geq 1$, we additionally define $\beta_{\boldsymbol{r}} := \theta_{\boldsymbol{r}}\sigma_{\boldsymbol{r}}\sqrt{D}$. The secret parameter $\theta_{\boldsymbol{s}}$ controls a single sample in $\mathbb{Z}^D$, whereas $\theta_{\mathbf{R}}$ absorbs a union bound over the $(k-1)D$ entries of $\mathbf{R}$. If $Q \geq 1$, $\theta_{\boldsymbol{r}}$ absorbs a union bound over at most nQ leakage masks. Lemma 35 then gives

$$\Pr\left[\|\boldsymbol{s}\|_\infty > \beta_{\boldsymbol{s}}\right] \leq 2^{-\kappa}, \tag{67}$$

$$\Pr\left[\exists (i',j') \in [k-1] \times [D] : |\mathbf{R}[i',j']| > \beta_{\mathbf{R}}\right] \leq 2^{-\kappa}, \tag{68}$$

$$\Pr\left[\exists s \in [Q],\ \exists i \in [n] : \|\boldsymbol{r}_{i,s}\|_2 > \beta_{\boldsymbol{r}}\right] \leq 2^{-\kappa} \text{ (when } Q \geq 1), \tag{69}$$

where (67) uses $\|\boldsymbol{s}\|_\infty \leq \|\boldsymbol{s}\|_2$, and the derivations are given in Appendix I. We collect these bounds in the following event.

Definition 13 (Good-tail event) Let $\mathsf{Good}_{\mathsf{tail}}$ denote the event that (1) the secret satisfies $\|\boldsymbol{s}\|_\infty \le \beta_{\boldsymbol{s}}$; (2) every entry of the masking matrix satisfies $|\mathbf{R}[i',j']| \le \beta_{\mathbf{R}}$; and (3) when $Q \ge 1$, for every query $s \in [Q]$ and every $i \in [n]$, the leakage mask satisfies $\|\boldsymbol{r}_{i,s}\|_2 \le \beta_{\boldsymbol{r}}$.

Lemma 36 (Failure probability of the good-tail event) *For the choices of $\theta_{\boldsymbol{s}}$ and $\theta_{\mathbf{R}}$ in (64) and (65), and, when $Q \ge 1$, of $\theta_{\boldsymbol{r}}$ in (66), we have $\Pr[\neg\mathsf{Good}_{\mathsf{tail}}] \le 3 \cdot 2^{-\kappa}$. When $Q = 0$, the leakage-mask condition in Definition 13 is vacuous, and $\Pr[\neg\mathsf{Good}_{\mathsf{tail}}] \le 2 \cdot 2^{-\kappa}$.*

The lemma follows from (67)–(69) by a union bound; the leakage-mask condition is vacuous when $Q = 0$. The secret and masking-matrix conditions, together with Lemma 12, give $\max_{i\in[n]} \|\boldsymbol{s}_i\|_\infty \le \beta_{\mathsf{sh}} = \beta_{\mathbf{g}}\beta_{\boldsymbol{s}} + c_{\exp}\beta_{\mathbf{R}}$, so the shortness error of Definition 7 satisfies $\epsilon_{\mathsf{sh}} \le 2 \cdot 2^{-\kappa}$. The leakage-mask condition is not needed for the privacy or shortness of SEC; when $Q \ge 1$, it bounds the released leakage blocks by $\|\boldsymbol{z}_{i,s}\|_2 \le \omega_c\sqrt{D}\beta_{\mathsf{sh}} + \beta_{\boldsymbol{r}}$, which an application that verifies partial responses must accept.

6.4.2 Instantiation Methodology

We instantiate the selected instances with $n \in \{40, 60\}$ from Table 4, for $\kappa \in \{128, 192, 256\}$. We fix the ring degree $d = 512$ and the systematic hard-core matrix $\boldsymbol{A} = \left[\boldsymbol{I}_p \ \boldsymbol{A}'\right] \in \mathcal{R}_q^{p\times\ell}$, and write $\ell = p + \ell'$, so that the hardness assumptions are $\mathsf{MLWE}_{p,\ell',q,d,\sigma}$ for the relevant widths σ. We consider two query budgets, $Q \in \{0, 2^{64}\}$. The budget $Q = 0$ corresponds to SEC as a standalone sharing scheme, in which the adversary observes only the shares of at most k_p parties; the budget $Q = 2^{64}$ matches the signing-query budgets targeted by lattice-based threshold signatures. The leakage-description set is $\mathcal{C} := \{c \in \mathcal{R} : \|c\|_\infty = 1, \ \|c\|_1 = \omega_c\}$, so that $|\mathcal{C}| = 2^{\omega_c}\binom{d}{\omega_c}$ and $\|\mathbf{L}_s\|_2 \le \omega_c$. Following the convention of [19, Section 7], we set $\omega_c := \min\{\omega \in [d] : 2^\omega\binom{d}{\omega} \ge 2^\kappa\}$, which gives $\omega_c = 19$, 31, and 44 for $\kappa = 128$, 192, and 256, respectively. When $Q = 0$, the second term of Ξ_{MQ} in (43) vanishes, $\sigma_{\mathrm{eff}} = \sigma_{\mathbf{R}}$ in (49), and the width $\sigma_{\boldsymbol{r}}$ plays no role. For each of these instances, each pair (κ, Q), and each candidate ℓ, we choose widths $\sigma_{\boldsymbol{s}}$, $\sigma_{\mathbf{R}}$, and $\sigma_{\boldsymbol{r}}$ satisfying the requirements of Section 6.4.3 so as to minimize the effective share width σ_{sh}, defined in Appendix I, which satisfies $\sigma_{\mathsf{sh}} \ge \beta_{\mathbf{g}}\sigma_{\boldsymbol{s}}$ and determines the size of one share. When $Q \ge 1$, minimizing σ_{sh} alone yields no finite optimizer, because increasing $\sigma_{\boldsymbol{r}}$ continually weakens the query-dependent restriction and drives the objective toward its $Q = 0$ limit. We therefore minimize $\sigma_{\mathsf{sh}}\sigma_{\boldsymbol{r}}$, which balances share size against the size of the released leakage masks. The resulting parameter sets and their sizes are reported in Tables 6 and 7 and discussed in Section 6.4.4.

6.4.3 Concrete Parameter Requirements

For every admissible pattern, Theorem 4 bounds the distinguishing advantage by $\epsilon_{\mathrm{MQ}} = \mathsf{Adv}_{\mathsf{MLWE}} + 2\epsilon' + 2\epsilon$ (see (61)), where $\mathsf{Adv}_{\mathsf{MLWE}}$ denotes the maximum advantage, over the good setup outputs, of the corresponding non-uniform reduction against the relevant MLWE problem, and $2\epsilon' + 2\epsilon$ is the statistical loss of the Gaussian approximation and posterior resampling (Lemma 34). For a concrete instantiation, we target

Table 6: Selected SEC parameter sets for the application-driven and fixed 32-bit modulus regimes

(n,k)	κ	$(k_p,k_c,k_c^{\mathsf{det}})$	Q	ω_c	$\lceil\lg q\rceil$	ℓ	$p=\ell'$	$\lg\sigma_{\boldsymbol{s}}$	$\lg\sigma_{\mathbf{R}}$	$\lg\sigma_{\boldsymbol{r}}$
application-driven threshold-signature moduli										
$(40,20)$	128	$(5,26,35)$	0	19	61	12	6	2.957	4.736	–
$(40,20)$	128	$(5,26,35)$	2^{64}	19	61	12	6	3.730	5.222	43.646
$(40,20)$	192	$(5,26,35)$	0	31	63	16	8	3.236	4.995	–
$(40,20)$	192	$(5,26,35)$	2^{64}	31	63	16	8	4.039	5.478	44.615
$(40,20)$	256	$(5,26,35)$	0	44	64	20	10	3.436	5.186	–
$(40,20)$	256	$(5,26,35)$	2^{64}	44	64	20	10	4.257	5.666	45.313
$(60,40)$	128	$(6,46,56)$	0	19	64	12	6	2.463	5.210	–
$(60,40)$	128	$(6,46,56)$	2^{64}	19	64	12	6	2.670	5.711	43.934
$(60,40)$	192	$(6,46,56)$	0	31	66	16	8	2.733	5.468	–
$(60,40)$	192	$(6,46,56)$	2^{64}	31	66	16	8	2.944	5.969	44.898
$(60,40)$	256	$(6,46,56)$	0	44	66	20	10	2.928	5.657	–
$(60,40)$	256	$(6,46,56)$	2^{64}	44	66	20	10	3.142	6.158	45.593
fixed modulus $q=2^{32}-10239$										
$(40,20)$	128	$(5,26,35)$	0	19	32	6	3	2.952	4.731	–
$(40,20)$	128	$(5,26,35)$	2^{64}	19	32	6	3	3.724	5.217	43.641
$(40,20)$	192	$(5,26,35)$	0	31	32	8	4	3.232	4.992	–
$(40,20)$	192	$(5,26,35)$	2^{64}	31	32	8	4	4.035	5.475	44.612
$(40,20)$	256	$(5,26,35)$	0	44	32	10	5	3.434	5.183	–
$(40,20)$	256	$(5,26,35)$	2^{64}	44	32	10	5	4.254	5.664	45.310
$(60,40)$	128	$(6,46,56)$	0	19	32	6	3	2.458	5.205	–
$(60,40)$	128	$(6,46,56)$	2^{64}	19	32	6	3	2.665	5.706	43.929
$(60,40)$	192	$(6,46,56)$	0	31	32	8	4	2.729	5.465	–
$(60,40)$	192	$(6,46,56)$	2^{64}	31	32	8	4	2.941	5.966	44.895
$(60,40)$	256	$(6,46,56)$	0	44	32	10	5	2.926	5.655	–
$(60,40)$	256	$(6,46,56)$	2^{64}	44	32	10	5	3.139	6.156	45.590

Note: $\ell = p + \ell'$ and $p = \ell' = \ell/2$. All logarithms are base two and are computed from the full-precision parameter values. The application-driven rows use the exact threshold-signature prime moduli corresponding to the displayed bit lengths, while the fixed 32-bit rows use $q = 2^{32} - 10239$. For $Q = 0$, no leakage mask is sampled. For $Q = 2^{64}$, the leakage blocks are released over $\mathcal{R}^{\ell}$, not modulo q, so $\sigma_{\boldsymbol{r}}$ may exceed q.

a total advantage of at most $2^{-\kappa}$. For the privacy reduction, we set $\epsilon' = \epsilon := 2^{-(\kappa+3)}$ and $\kappa_{\mathsf{MLWE}} := \kappa + 1$, so that its statistical and computational contributions are each bounded by $2^{-(\kappa+1)}$. Here, following the usual convention, κ_{MLWE} bits of MLWE security are read as $\mathsf{Adv}_{\mathsf{MLWE}} \le 2^{-\kappa_{\mathsf{MLWE}}}$. A parameter set $(q, p, \ell', \sigma_{\boldsymbol{s}}, \sigma_{\mathbf{R}}, \sigma_{\boldsymbol{r}})$, with $\sigma_{\boldsymbol{r}}$ omitted when $Q = 0$, is admissible for an instance and a pair (κ, Q) if it satisfies the following requirements, which are the hypotheses of Theorem 4 together with the tail requirements of Section 6.4.1.

1. **Kernel smoothing.** The computable condition (56) holds with target ϵ', using the underlying full-precision values of β_{k_p} and γ_{k_p} reported in rounded form in Table 4.
2. **Posterior resampling.** The width σ_{MQ} in (58) satisfies $\sigma_{\mathrm{MQ}} \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^D)$, evaluated using the estimate of $\eta_\epsilon(\mathbb{Z}^N)$ from Section 2.2 and the underlying full-precision values of α_{k_p} and $\|\hat{\mathbf{G}}\|_2$ reported in rounded form in Table 4.
3. **Computational hardness.** Both $\mathsf{MLWE}_{p,\ell',q,d,\sigma_{\mathrm{MQ}}}$, for nonempty auxiliary views, and $\mathsf{MLWE}_{p,\ell',q,d,\sigma_{\boldsymbol{s}}}$, for the empty view, provide at least κ_{MLWE} bits of security against the known attacks evaluated by the lattice estimator [62] under the cost model of [63]. For this concrete evaluation, we assume that the polynomial-size advice pp, which is independent of the random MLWE matrix $\boldsymbol{A}'$, does not improve these attacks; the formal non-uniform hardness assumption remains that of Theorem 4. Since $\sigma_{\mathrm{MQ}} \leq \sigma_{\boldsymbol{s}}/\sqrt{2}$, the requirement for $\mathsf{MLWE}_{p,\ell',q,d,\sigma_{\boldsymbol{s}}}$ is implied in practice by the one for $\mathsf{MLWE}_{p,\ell',q,d,\sigma_{\mathrm{MQ}}}$, because the estimated security does not decrease with the width.
4. **Tail bounds and shortness.** The tail parameters satisfy (64), (65), and, when $Q \geq 1$, (66); the resulting share bound is $\beta_{\mathsf{sh}} = \beta_{\mathbf{g}}\beta_{\boldsymbol{s}} + c_{\exp}\beta_{\mathbf{R}}$ from Lemma 12.

Table 7: Ideal entropy-coded sizes, in KiB, of the secret and of one share for the selected parameter sets

				threshold-signature moduli		$\lceil \lg q \rceil = 32$	
(n, k)	$(k_p, k_c, k_c^{\mathsf{det}})$	κ	Q	Secret $\boldsymbol{s}$	Share $\boldsymbol{s}_i$	Secret $\boldsymbol{s}$	Share $\boldsymbol{s}_i$
$(40, 20)$	$(5, 26, 35)$	128	0	3.753	5.697	1.875	2.847
$(40, 20)$	$(5, 26, 35)$	128	2^{64}	4.333	6.068	2.164	3.032
$(40, 20)$	$(5, 26, 35)$	192	0	5.283	7.856	2.640	3.926
$(40, 20)$	$(5, 26, 35)$	192	2^{64}	6.087	8.350	3.041	4.173
$(40, 20)$	$(5, 26, 35)$	256	0	6.854	10.058	3.425	5.027
$(40, 20)$	$(5, 26, 35)$	256	2^{64}	7.880	10.674	3.938	5.335
$(60, 40)$	$(6, 46, 56)$	128	0	3.383	6.041	1.689	3.019
$(60, 40)$	$(6, 46, 56)$	128	2^{64}	3.538	6.416	1.767	3.206
$(60, 40)$	$(6, 46, 56)$	192	0	4.780	8.313	2.388	4.155
$(60, 40)$	$(6, 46, 56)$	192	2^{64}	4.991	8.812	2.494	4.404
$(60, 40)$	$(6, 46, 56)$	256	0	6.219	10.628	3.108	5.312
$(60, 40)$	$(6, 46, 56)$	256	2^{64}	6.486	11.252	3.242	5.624

Note: sizes exclude implementation-specific coding, framing, and byte-alignment overheads; each participant stores exactly one share.

6.4.4 Security Estimates and Selected Parameter Sets

Privacy imposes no lower bound on the hard-core modulus q, and the requirements of Section 6.4.3 are independent of it. We therefore evaluate two modulus regimes: the

application-driven moduli obtained from the MSIS estimates for the corresponding threshold-signature instantiations, and the fixed smaller prime $q = 2^{32} - 10239$, for which $\lceil \lg q \rceil = 32$ and $q \equiv 1 \pmod{2d}$. For each pair (n, κ) and each modulus regime, we select the smallest even ℓ, with $p = \ell' = \ell/2$, for which $\mathsf{MLWE}_{p,\ell',q,d,\sigma_{\mathrm{MQ}}}$ provides at least κ_{MLWE} bits of security. This choice also minimizes the sizes, which increase with ℓ. At the selected widths, the posterior-resampling requirement holds with equality, so $\sigma_{\mathrm{MQ}} = \sqrt{2}\eta_\epsilon(\mathbb{Z}^D)$ depends only on κ and ℓ, and the same security estimate applies to both query budgets. Every retained tuple is evaluated at its exact prime modulus using the full LWE attack set of the *lattice estimator* introduced in [62], including primal, BDD, dual, and hybrid attacks under the reduction-cost model of [63].

For the threshold-signature moduli, the selected parameters remain private for every hard-core modulus satisfying $\lceil \lg q \rceil \leq 68$ when $\kappa \in \{128, 192\}$ and $\lceil \lg q \rceil \leq 69$ when $\kappa = 256$. For the fixed 32-bit modulus, the selected module dimensions are $\ell = 6$, 8, and 10 for $\kappa = 128$, 192, and 256, respectively, for $n \in \{40, 60\}$. These parameter sets remain private for every hard-core modulus satisfying $\lceil \lg q \rceil \leq 35$, and their secrets and shares are approximately half the size of those obtained with the threshold-signature moduli. Table 6 reports the parameter sets for both modulus regimes. The corresponding ideal entropy-coded sizes are reported in Table 7 and computed as described in Appendix I.

7 Discussion and Conclusion

SEC shows that lossless bipartite expanders yield everywhere-short secret sharing over $\mathbb{Z}$-modules in which each participant receives exactly one share, with bounded share coefficients and bounded entries in the integer reconstruction vectors. Its correctness and shortness follow from integer linear algebra over the expander-derived distribution matrix, while its hard-core privacy against correlated multi-query local leakage reduces to standard MLWE. We conclude by discussing the ramp gap, limitations, and open problems.

7.1 The Ramp Gap

The principal trade-off in assigning exactly one bounded-coefficient share to each party is the gap between the privacy threshold k_p and the reconstruction threshold k_c, which remains the main limitation of the construction. The privacy threshold is governed by the expansion of the underlying bipartite graph, whereas the reconstruction threshold depends on the resulting distribution matrix and the construction of its secret-bearing column. For the selected instances in Tables 4 and 5, $\tau_p \in [0.100, 0.200]$ and $\tau_c \in [0.620, 0.833]$. The near-threshold black-box schemes of [13, 14] and the random-assignment ramp sharing of [32] likewise trade a gap between the thresholds for compact shares, although their reconstruction coefficients are either integers without a stated bound [13] or binary [14, 32], rather than bounded signed integers.

7.2 Open Problems and Future Directions

The main open problem is to construct an everywhere-short secret sharing scheme over $\mathbb{Z}$-modules with one share per participant, or more generally $O(1)$ shares per participant, that realizes a perfect (non-ramp) threshold access structure, or at least a substantially narrower ramp gap than SEC. A further question is how to instantiate SEC for larger committees, beyond the exhaustively certified regime considered here. Finally, applications of SEC to lattice-based distributed cryptography, such as threshold-signing protocols, are left for future work.

Statements and Declarations

Funding

This work was partially supported by Australian Research Council Discovery Grants DP250100229 and DP260100245. M. A. is supported by a scholarship from the Indonesia Endowment Fund for Education (LPDP), grant no. SKPB9126/LPDP/LPDP.3/2023.

Competing interests

The authors have no competing interests to declare that are relevant to the content of this article.

Data availability

The SageMath scripts used for the expander search, certification, and evaluation in Section 6 are publicly available at https://github.com/arzaki-cstheory/SEC-LaTSEC under the GPL-3.0 license.

Author contributions

M. A. contributed to the construction, developed the proofs, implemented the experiments and parameter search, and wrote the initial manuscript. A. S. and R. S. contributed to the conception of the work and the interpretation of the results, supervised the project, and revised the manuscript. M. F. E. contributed to the interpretation of the results, supervised the project, and revised the manuscript. All authors read and approved the final manuscript.

Use of AI tools

During the preparation of this manuscript, the authors used Claude (Anthropic) and Microsoft Copilot to assist with editing the text, improving the technical writing, checking the consistency of proofs and notation, drafting revisions of some passages, and preparing and debugging the SageMath scripts used for the experiments and parameter search. All such content and code were reviewed, verified, and revised by the authors, who take full responsibility for the content of the publication.

Appendix A Additional Notation and Linear-Algebra Conventions

For distributions X and Y, their statistical distance is denoted by $\Delta(X, Y)$. For real vectors $\mathbf{u}$ and $\mathbf{v}$ of the same dimension, let $\langle \mathbf{u}, \mathbf{v} \rangle := \mathbf{u}^\top \mathbf{v}$. We write $\mathbf{A}^\top$ for transpose, and $\mathbf{I}$ and $\mathbf{O}$ for identity and zero matrices of the required dimensions. The symbols $\otimes$ and $\oplus$ denote the Kronecker product and block-diagonal direct sum, respectively, and $\mathrm{vec}(\mathbf{A})$ denotes column-wise vectorization. A matrix of the form $\begin{bmatrix}\mathbf{I} & \mathbf{A}\end{bmatrix}$ is in systematic form.

For vectors, $\|\cdot\|_1$, $\|\cdot\|_2$, and $\|\cdot\|_\infty$ denote the 1-norm, Euclidean norm, and infinity norm, respectively, applied after coefficient embedding when necessary. For real matrices, $\|\cdot\|_2$, $\|\cdot\|_F$, $\sigma_{\min}(\cdot)$, and $\sigma^+_{\min}(\cdot)$ denote the spectral norm, Frobenius norm, smallest singular value, and smallest positive singular value, respectively, and $\|\mathbf{A}\|_{\max} := \max_{i,j} |\mathbf{A}[i,j]|$. We write $\mathbf{A} \succeq \mathbf{B}$ (resp. $\mathbf{A} \succ \mathbf{B}$) when $\mathbf{A} - \mathbf{B}$ is positive semidefinite (resp. positive definite), and $\mathbf{A}^+$ for the Moore-Penrose pseudoinverse; if $\mathbf{A}$ has full row rank, then $\mathbf{A}^+ = \mathbf{A}^\top(\mathbf{A}\mathbf{A}^\top)^{-1}$ and $\|\mathbf{A}^+\|_2 = 1/\sigma_{\min}(\mathbf{A})$. Recall that $\mathbf{A} \succeq \mathbf{B} \succ 0$ implies $\mathbf{A}^{-1} \preceq \mathbf{B}^{-1}$.

Appendix B Gaussian Normalizations and Matrix Hint-MLWE Proofs

B.1 Normalizations

Since $\rho_{1/s}(\mathbf{y}) = \exp(-\pi s^2 \|\mathbf{y}\|_2^2)$, substituting $s = \sqrt{2\pi}\sigma$ in (2) gives $\eta_\epsilon(\Lambda) = \frac{\eta^\pi_\epsilon(\Lambda)}{\sqrt{2\pi}}$, where $\eta^\pi_\epsilon(\Lambda) := \inf\left\{s > 0 : \rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) \le \epsilon\right\}$ is the normalization of [41, Definition 5]. Similarly, for $\mathbf{\Sigma} \succ 0$ and $\mathbf{S} := \sqrt{2\pi}\mathbf{\Sigma}^{1/2}$, $\exp\left(-\pi\|\mathbf{S}^{-1}\mathbf{x}\|_2^2\right) = \exp\left(-\frac{1}{2}\mathbf{x}^\top\mathbf{\Sigma}^{-1}\mathbf{x}\right)$. Thus, the parameter matrix $\mathbf{\Sigma}$ in (1) corresponds to the matrix parameter $\sqrt{2\pi}\mathbf{\Sigma}^{1/2}$ under the convention of [40].

B.2 Deferred Proofs

Proof of Theorem 3 Fix the tuple of hint matrices from the theorem statement and define $\mathbf{\Xi} := \mathbf{\Sigma}_\mathbf{s}^{-1} + \sum_{i=1}^h \mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1}\mathbf{M}_i$. The posterior-resampling proofs underlying the Matrix Hint-MLWE hardness results of [24, 48] use $\mathbf{\Xi}$ as the precision matrix of the secret conditioned on the hints. The argument requires an upper bound on $\|\mathbf{\Xi}\|_2$. By the triangle inequality and the definition of the effective source width, $\|\mathbf{\Xi}\|_2 \le \left\|\mathbf{\Sigma}_\mathbf{s}^{-1}\right\|_2 + \sum_{i=1}^h \left\|\mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1}\mathbf{M}_i\right\|_2 = \frac{1}{2\sigma^2}$. Thus, the precision bound required by the correction and smoothing steps holds directly. In contrast, the standard separate-product formulation further applies $\|\mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1}\mathbf{M}_i\|_2 \le \|\mathbf{M}_i^\top\mathbf{M}_i\|_2\|\mathbf{\Sigma}_{\mathbf{y}_i}^{-1}\|_2$. Retaining the joint term yields the condition stated in Theorem 3. The remaining posterior-resampling, convolution, uniform-target, and advantage-accounting arguments are unchanged and incur additive distinguishing loss at most 2ϵ. □

Proof of Corollary 1 Fix the tuple of hint matrices and write $\mathbf{\Xi} := \mathbf{\Sigma}_\mathbf{s}^{-1} + \sum_{i=1}^h \mathbf{M}_i^\top \mathbf{\Sigma}_{\mathbf{y}_i}^{-1}\mathbf{M}_i$ and $\mathbf{\Sigma}_0 := \mathbf{\Xi}^{-1}$. By the definition of σ, we have $\|\mathbf{\Xi}\|_2 \le 1/(2\sigma^2)$, and hence $\lambda_{\min}(\mathbf{\Sigma}_0) = 1/\|\mathbf{\Xi}\|_2 \ge 2\sigma^2$. Run the posterior-resampling reduction underlying Theorem 3 with source width σ' in place of σ. The corresponding correction parameter matrix is $\mathbf{\Sigma}_0 - \sigma'^2\mathbf{I}_N$, which

is positive definite because $\lambda_{\min}\left(\mathbf{\Sigma}_0 - \sigma'^2\mathbf{I}_N\right) \geq 2\sigma^2 - \sigma'^2 \geq \sigma'^2 > 0$. It remains to verify the smoothing condition in the convolution step. Set $\mathbf{\Sigma}_1 := \sigma'^2\mathbf{I}_N$, $\mathbf{\Sigma}_2 := \mathbf{\Sigma}_0 - \sigma'^2\mathbf{I}_N$, and define $\mathbf{\Sigma}_3^{-1} := \mathbf{\Sigma}_1^{-1} + \mathbf{\Sigma}_2^{-1}$. The preceding bound gives $\lambda_{\min}(\mathbf{\Sigma}_2) \geq \sigma'^2$, and therefore $\left\|\mathbf{\Sigma}_3^{-1}\right\|_2 \leq \frac{1}{\sigma'^2} + \frac{1}{\lambda_{\min}(\mathbf{\Sigma}_2)} \leq \frac{2}{\sigma'^2} \leq \frac{1}{\eta_\epsilon(\mathbb{Z}^N)^2}$, where the last inequality follows from $\sigma' \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^N)$. Since $\exp(-\frac{1}{2}\mathbf{x}^\top\mathbf{\Sigma}^{-1}\mathbf{x}) = \exp(-\pi\,\mathbf{x}^\top(2\pi\mathbf{\Sigma})^{-1}\mathbf{x})$ and $\eta_\epsilon^\pi = \sqrt{2\pi}\,\eta_\epsilon$, this is exactly the hypothesis of the convolution lemma [41, Lemmas 3 and 5] in the normalization (2). Thus, the same convolution argument applies with statistical loss at most 2ϵ. The posterior-resampling identity, the uniform-target branch, and the advantage accounting are independent of the choice between σ and σ', so the remainder of the reduction is unchanged. □

Appendix C Conventional Privacy

Conventional secret-sharing privacy compares the share views induced by two arbitrarily fixed secrets.

Definition 14 (Conventional privacy) Let $\mathsf{Sharing} = (\mathsf{Setup}, \mathsf{Deal}, \mathsf{Rec})$ be a secret sharing scheme with secret space S. Fix pp in the support of $\mathsf{Setup}(1^\kappa)$. For $U \in \Gamma_0(\mathsf{pp})$, secrets $s, s' \in S$, and an adversary $\mathcal{A}$, define

$$\mathsf{Adv}^{\mathsf{priv}}_{\mathcal{A},\mathsf{Sharing}}(\mathsf{pp}, U, s, s') := \Big|\Pr\Big[\mathcal{A}(\mathsf{pp}, (s_i)_{i\in U}) = 1 \;\Big|\; (s_i)_{i\in[n]} \leftarrow \mathsf{Deal}_{\mathsf{pp}}(s)\Big] - \Pr\Big[\mathcal{A}(\mathsf{pp}, (s'_i)_{i\in U}) = 1 \;\Big|\; (s'_i)_{i\in[n]} \leftarrow \mathsf{Deal}_{\mathsf{pp}}(s')\Big]\Big|,$$

where the probabilities are over the randomness of Deal and $\mathcal{A}$. The scheme is *computationally δ-private* for pp if $\mathsf{Adv}^{\mathsf{priv}}_{\mathcal{A},\mathsf{Sharing}}(\mathsf{pp}, U, s, s') \leq \delta$ for every PPT adversary $\mathcal{A}$, every $U \in \Gamma_0(\mathsf{pp})$, and every $s, s' \in S$. It is *statistically δ-private* if the same condition holds for every possibly unbounded adversary, and it is *perfectly private* if it is statistically 0-private.

Appendix D Deferred Proofs for Codes, Construction, and Correctness

D.1 Codes and Expanders

Proof of Lemma 1 Let $\mathbf{c} \in C_{\mathbb{Q}} \setminus \{\mathbf{0}\}$, and let $\mathbf{a} := D\mathbf{c} \in \mathbb{Z}^n$ for a positive common denominator D of its entries, so that $\mathbf{Ha} = \mathbf{0}$ over $\mathbb{Z}$ and $\mathrm{supp}(\mathbf{a}) = \mathrm{supp}(\mathbf{c})$. With $t := \min_{j\in\mathrm{supp}(\mathbf{a})} v_p(\mathbf{a}[j])$, where v_p denotes the p-adic valuation, the vector $\mathbf{b} := p^{-t}\mathbf{a} \in \mathbb{Z}^n$ satisfies $\mathbf{Hb} = \mathbf{0}$ and has an entry that is nonzero modulo p. Hence $\overline{\mathbf{b}} := \mathbf{b} \bmod p$ is a nonzero codeword of C_p with $\mathrm{supp}(\overline{\mathbf{b}}) \subseteq \mathrm{supp}(\mathbf{c})$, so $d(C_p) \leq wt_H(\overline{\mathbf{b}}) \leq wt_H(\mathbf{c})$. Taking the minimum over all nonzero $\mathbf{c} \in C_{\mathbb{Q}}$ gives $d(C_{\mathbb{Q}}) \geq d(C_p)$. □

Proof of Lemma 3 Let $S \subseteq L$ be nonempty and satisfy $|S| < 2(1-\epsilon_{\mathrm{exp}})\lfloor\gamma n\rfloor$. If $|S| \leq \gamma n$, then Lemma 2 gives $|U(S)| \geq c_{\mathrm{exp}}(1-2\epsilon_{\mathrm{exp}})|S| > 0$, where the strict inequality uses $|S| \geq 1$, that is, the nonemptiness of S, together with $c_{\mathrm{exp}} \geq 1$ and $\epsilon_{\mathrm{exp}} < 1/2$. Otherwise, $|S| > \gamma n \geq \lfloor\gamma n\rfloor \geq 1$, so we may choose $T \subsetneq S$ with $|T| = \lfloor\gamma n\rfloor$. Then $|U(T)| \geq c_{\mathrm{exp}}(1-2\epsilon_{\mathrm{exp}})\lfloor\gamma n\rfloor$ by Lemma 2, while $|S \setminus T| < 2(1-\epsilon_{\mathrm{exp}})\lfloor\gamma n\rfloor - \lfloor\gamma n\rfloor = (1-2\epsilon_{\mathrm{exp}})\lfloor\gamma n\rfloor$. By left-regularity,

$|N(S \setminus T)| \le c_{\exp}|S \setminus T| < |U(T)|$, so some vertex of $U(T) \setminus N(S \setminus T)$ has exactly one neighbor in T and none in $S \setminus T$, and hence exactly one neighbor in S. □

Deterministic Peeling Pseudocode

Algorithm 3 states the deterministic peeling procedure of Section 2.3.2. The rows of $\hat{\mathbf{G}}_U$ are indexed by $[r]$, where $r := |U|$, in increasing order of the elements of U.

Algorithm 3 Peeling-based construction of a unimodular submatrix.

Require: A matrix $\hat{\mathbf{G}}_U \in \mathbb{Z}^{r \times (k-1)}$, where $r := |U|$, satisfying the unit-valued unique-neighbor property: for every nonempty $S \subseteq [r]$, there exists $j \in [k-1]$ such that $\left|\left\{i \in S : \hat{\mathbf{G}}_U[i,j] \neq 0\right\}\right| = 1$, and the unique nonzero entry belongs to $\{-1, 1\}$.

Ensure: An ordered column tuple $J_U = (j_1, \ldots, j_r)$ such that $\mathbf{M}_U := \hat{\mathbf{G}}_U[\cdot, J_U]$ is unimodular.

$S := [r]$, $J_U := ()$, and $I_U := ()$.
for $t = 1$ **to** r **do**
  for $j \in [k-1] \setminus \{j_1, \ldots, j_{t-1}\}$ **do**
    $\mathrm{supp}_S(j) := \{i \in S : \hat{\mathbf{G}}_U[i,j] \neq 0\}$.
  end for
  $C := \left\{ j \in [k-1] \setminus \{j_1, \ldots, j_{t-1}\} : \begin{array}{l} |\mathrm{supp}_S(j)| = 1, \\ \hat{\mathbf{G}}_U[i(j), j] \in \{-1, 1\} \end{array} \right\}$, where $i(j)$ is the unique row in $\mathrm{supp}_S(j)$.
  if $C = \emptyset$ **then**
    return $\perp$. ▷ This cannot occur under the assumed property.
  end if
  $j_t := \min C$.
  Let i_t be the unique row in $\mathrm{supp}_S(j_t)$.
  $J_U := (j_1, \ldots, j_t)$ and $I_U := (i_1, \ldots, i_t)$.
  $S := S \setminus \{i_t\}$.
end for
$\mathbf{M}_U := \hat{\mathbf{G}}_U[\cdot, J_U]$.
return $(\mathbf{M}_U, J_U, I_U)$.

For $1 \le |U| \le \lceil 2(1 - \epsilon_{\exp})\lfloor \gamma n \rfloor \rceil - 1$, the row restriction $\hat{\mathbf{G}}_U$ of $\hat{\mathbf{G}} = (\mathbf{PH})^\top$ satisfies the unit-valued unique-neighbor property: a column with exactly one nonzero entry in S corresponds to a unique neighbor of the corresponding nonempty set of left vertices, which exists by Lemma 3, and every nonzero entry of $\hat{\mathbf{G}}$ lies in $\{-1, 1\}$ because $\mathbf{P}$ is a signed permutation matrix. Moreover, once a column j_t and its unique active row i_t are selected, the column j_t has no nonzero entry among the remaining active rows. Thus, no previously selected column can have singleton support in a later iteration, so excluding the previously selected columns does not affect the nonemptiness of C. Hence, the candidate set C in Algorithm 3 is nonempty at every iteration.

Under the peeling row order $I_U = (i_1, \ldots, i_r)$, the matrix $\mathbf{M}_U^{\triangle} := \hat{\mathbf{G}}_U[I_U, J_U]$ is triangular with diagonal entries in $\{-1, 1\}$, so $\left|\det\left(\mathbf{M}_U^{\triangle}\right)\right| = 1$. Since $\mathbf{M}_U$ differs from $\mathbf{M}_U^{\triangle}$ only by a row permutation, $|\det(\mathbf{M}_U)| = 1$, and thus $\mathbf{M}_U \in GL_r(\mathbb{Z})$. The

deterministic choice $j_t := \min C$ and the fixed ordering conventions make $(\mathbf{M}_U, J_U, I_U)$ uniquely determined by $\hat{\mathbf{G}}$ and U.

D.2 Construction

Proof of Lemma 6 Since $\mathbf{U_A}$ is unimodular, $\mathrm{Row}_{\mathbb{Z}}(\mathbf{R_A}) = \mathrm{Row}_{\mathbb{Z}}(\mathbf{A})$. If $\mathbf{R_A}[1,\cdot] = \mathbf{e}_1^\top$, then $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{R_A}) = \mathrm{Row}_{\mathbb{Z}}(\mathbf{A})$. Conversely, suppose that $\mathbf{e}_1^\top = \sum_{i=1}^{r} \alpha_i \mathbf{R_A}[i,\cdot]$ with $\alpha_1, \ldots, \alpha_r \in \mathbb{Z}$. Since $\mathbf{e}_1^\top$ has a nonzero first coordinate, some row of $\mathbf{R_A}$ is nonzero in column 1, so $j_1 = 1$, and since $j_1 < \cdots < j_r$, row 1 is the only such row. Comparing first coordinates gives $1 = \alpha_1 p_1$, so $\alpha_1 = p_1 = 1$. Suppose, for contradiction, that $\alpha_i \neq 0$ for some $i \geq 2$, and let i be the smallest such index. Rows $\iota > i$ vanish in column j_i, and $\mathbf{e}_1^\top$ vanishes in column $j_i > 1$, so comparing coordinates j_i gives $0 = \mathbf{R_A}[1, j_i] + \alpha_i p_i$. Since $0 \leq \mathbf{R_A}[1, j_i] < p_i$, this forces $\alpha_i = 0$, a contradiction. Hence $\alpha_2 = \cdots = \alpha_r = 0$, and therefore $\mathbf{R_A}[1,\cdot] = \mathbf{e}_1^\top$. Finally, in this case, $\boldsymbol{\lambda}_0^\top \mathbf{A} = \mathbf{U_A}[1,\cdot]\mathbf{A} = \mathbf{R_A}[1,\cdot] = \mathbf{e}_1^\top$. □

Proof of Lemma 9 Every $r+1$ columns of $\mathbf{H}'$ are linearly dependent over $\mathbb{Q}$, and clearing denominators in a nonzero rational kernel vector gives a nonzero vector in $\ker_{\mathbb{Z}}(\mathbf{H}'_{(T)})$ for every $T \subseteq [n]$ with $|T| = r+1$. Partitioning $[n]$ into $b := \lfloor n/(r+1) \rfloor$ pairwise disjoint sets $T_1, \ldots, T_b$ of size $r+1$ and a leftover set R with $|R| \leq r$ proves the claim. □

Complexity of Algorithm 2

The baseline construction using supports of size $t_{\max} = r+1$, where $r := \mathrm{rank}_{\mathbb{Q}}(\mathbf{H}')$, processes at most $\lfloor n/(r+1) \rfloor$ supports and uses polynomial-time integer linear-algebra and generalized-Bézout computations. The optional search for smaller supports examines at most *maxTrials* candidates for each size $t \in [t_0, t_{\max} - 1]$ in every iteration. For fixed t_0 and *maxTrials*, the number of rank tests in any execution is $O(b \cdot \mathit{maxTrials} \cdot (t_{\max} - t_0))$, where $b := |\mathcal{T}|$. Each rank or kernel computation is polynomial in m, t, and the bit length of the entries of $\mathbf{H}'$. Thus, the construction has polynomial bit complexity for fixed search parameters.

D.3 Reconstruction Coverage

Deferred Proofs

Proof of Corollaries 2 and 3 For each $i \in [b]$, $\Pr[T_i \subseteq S] = \binom{n-t_i}{s-t_i} / \binom{n}{s}$, and, since the supports are pairwise disjoint, for every $\emptyset \neq J \subseteq [b]$, $\Pr[\bigcap_{i \in J} (T_i \subseteq S)] = \binom{n-\sum_{i\in J} t_i}{s-\sum_{i\in J} t_i} / \binom{n}{s}$. Inclusion-exclusion gives the union probability in (5), and the second-order Bonferroni inequality gives (6). For the witness version, suppose that $W_i \subseteq S$. Restrict $\boldsymbol{\lambda}_{T_i}$ to W_i and extend it by zero to the remaining coordinates of S. The resulting vector $\boldsymbol{\lambda}_S \in \mathbb{Z}^{|S|}$ satisfies $\boldsymbol{\lambda}_S \in \ker_{\mathbb{Z}}(\mathbf{H}'_{(S)})$ and $\boldsymbol{\lambda}_S^\top(\mathbf{g})_S = 1$. Therefore, $W_i \subseteq S$ implies $\mathbf{e}_1^\top \in \mathrm{Row}_{\mathbb{Z}}(\mathbf{G}_S)$. Since the sets T_i are pairwise disjoint and $W_i \subseteq T_i$, the witness supports $W_1, \ldots, W_b$ are also pairwise disjoint. Hence the same inclusion-exclusion and Bonferroni calculations apply with t_i replaced by $w_i := |W_i|$. □

Heuristic Design Rules

If all planted supports have a common size $t_i = t$, write $p(t) := \binom{n-t}{s-t} / \binom{n}{s}$ and $p(2t) := \binom{n-2t}{s-2t} / \binom{n}{s}$. Then (6) gives $\Pr[\bigcup_{i=1}^{b} (T_i \subseteq S)] \geq bp(t) - \binom{b}{2} p(2t)$. For

parameter-selection intuition, when $t \ll n$, we have $p(t) = (s)_t/(n)_t \approx (s/n)^t$. If the containment events are treated heuristically as independent, then $\Pr[\bigcup_{i=1}^{b}(T_i \subseteq S)] \approx 1-(1-p(t))^b \approx 1-(1-(s/n)^t)^b$. Accordingly, a target reconstruction error ϵ_c suggests $\epsilon_c \gtrapprox (1-(s/n)^t)^b$. Setting $s = k_c$ gives $k_c \gtrapprox (1-\epsilon_c^{1/b})^{1/t} n$. When ϵ_c is small and b is moderately large, $k_c \gtrapprox (\ln(1/\epsilon_c)/b)^{1/t} n$. For witness supports of a common size $w_i = w$, the same heuristic gives the rules with t replaced by w. These independence-based expressions are heuristic approximations used only for parameter-selection intuition; the inclusion-exclusion and Bonferroni expressions are rigorous lower bounds.

D.4 Linearity and Public-Map Compatibility

Proof of Lemmas 17 and 18 Let $\mathbf{X}_T, \mathbf{Y}_T \in \mathbb{Z}^{|T|\times D}$ have rows $\mathbf{x}_i^\top$ and $\mathbf{y}_i^\top$, ordered according to the indices in T. The matrix with rows $(\alpha\mathbf{x}_i + \mathbf{y}_i)^\top$ is $\alpha\mathbf{X}_T + \mathbf{Y}_T$, and the matrix with rows $(\mathbf{M}\mathbf{x}_i)^\top = \mathbf{x}_i^\top\mathbf{M}^\top$ is $\mathbf{X}_T\mathbf{M}^\top$. Hence $(\boldsymbol{\lambda}_T^\top(\alpha\mathbf{X}_T + \mathbf{Y}_T))^\top = \alpha(\boldsymbol{\lambda}_T^\top\mathbf{X}_T)^\top + (\boldsymbol{\lambda}_T^\top\mathbf{Y}_T)^\top$ and $(\boldsymbol{\lambda}_T^\top\mathbf{X}_T\mathbf{M}^\top)^\top = \mathbf{M}(\boldsymbol{\lambda}_T^\top\mathbf{X}_T)^\top$, which are the two claims. □

Appendix E Babai Decoding and Greedy Refinement

This appendix specifies the procedure BabaiGreedy (Algorithm 4) invoked by Algorithm 1 and justifies its correctness, complexity, and determinism.

The first stage of Algorithm 4 subtracts an element of K_T from $\boldsymbol{\lambda}_{T,0}$, and hence returns $\boldsymbol{\lambda}_{\mathrm{dec}} \in \boldsymbol{\lambda}_{T,0} + K_T$. Every update in the second stage also adds an element of K_T, so the final output $\boldsymbol{\lambda}_T$ remains in the same affine coset. Consequently, if $\boldsymbol{\lambda}_{T,0}^\top\mathbf{G}_T = \mathbf{e}_1^\top$, then $\boldsymbol{\lambda}_T^\top\mathbf{G}_T = \mathbf{e}_1^\top$. The greedy procedure accepts only updates that strictly decrease $\|\boldsymbol{\lambda}\|_\infty$. Since this norm is a nonnegative integer and the number of passes is fixed, Stage 2 terminates and returns a vector satisfying $\|\boldsymbol{\lambda}_T\|_\infty \le \|\boldsymbol{\lambda}_{\mathrm{dec}}\|_\infty$. The exact certification step in Algorithm 1 verifies the reconstruction equation before the vector is used by $\mathsf{Rec}_{\mathsf{pp}}$.

The Gram-Schmidt orthogonalization and Babai decoding in Stage 1 require $O(r^2|T|)$ arithmetic operations. In Stage 2, each of the *passes* passes examines r basis vectors and at most $2 \cdot \mathit{window} + 1$ candidate multiples per basis vector, while evaluating each candidate requires $O(|T|)$ arithmetic operations. Since *passes* and *window* are fixed small parameters, the overall arithmetic complexity is $O(r^2|T|)$. Its bit complexity is polynomial in r, $|T|$, and the bit lengths of the rational and integer values encountered during the computation.

It remains to justify the deterministic-selection claim in Remark 4. Different transformation matrices $\mathbf{U}_{\mathbf{G}_T}$ may yield different particular solutions $\boldsymbol{\lambda}_{T,0}$, but any two such solutions differ by an element of K_T and therefore belong to the same affine coset. Let $V_T := \mathrm{span}_{\mathbb{R}}(K_T)$, and let $\mathbf{b}_1^*, \ldots, \mathbf{b}_r^*$ be the Gram-Schmidt vectors of the ordered reduced basis $\mathbf{B}'_K$. Define $\mathcal{P}_{1/2}(\mathbf{B}'^*_K) := \{\sum_{i=1}^r \alpha_i\mathbf{b}_i^* : \alpha_i \in [-1/2, 1/2)\} \subseteq V_T$. For $\mathbf{t} \in \mathbb{R}^{|T|}$, denote by $\mathsf{NP}_{\mathbf{B}'_K}(\mathbf{t})$ the residual returned by Stage 1 on input $\mathbf{t}$. With the rounding convention $\widehat{x} := \lfloor x + 1/2 \rfloor$, Babai's nearest-plane procedure, applied within V_T, selects the unique residual satisfying $\mathsf{NP}_{\mathbf{B}'_K}(\mathbf{t}) \in (\mathbf{t}+K_T) \cap \big(\mathcal{P}_{1/2}(\mathbf{B}'^*_K) + V_T^\perp\big)$ [56, 57].

Algorithm 4 Babai decoding and greedy ℓ_∞-refinement for a reconstruction vector.

Require: A vector $\boldsymbol{\lambda}_{T,0} \in \mathbb{Z}^{|T|}$; an $\mathsf{LLL}_{\delta_{\mathsf{LLL}}}$-reduced basis $\mathbf{B}'_K \in \mathbb{Z}^{r\times|T|}$ of $K_T := \ker_{\mathbb{Z}}(\mathbf{G}_T^\top)$; and fixed small parameters *passes*, *window* $\in \mathbb{Z}^+$.

Ensure: A vector $\boldsymbol{\lambda}_T \in \boldsymbol{\lambda}_{T,0} + K_T$ satisfying $\|\boldsymbol{\lambda}_T\|_\infty \le \|\boldsymbol{\lambda}_{\mathrm{dec}}\|_\infty$.

Stage 1: Babai nearest-plane decoding.

Let $\mathbf{b}_i := (\mathbf{B}'_K[i,\cdot])^\top$ for $i \in [r]$, and compute the Gram-Schmidt vectors $\mathbf{b}_1^*, \dots, \mathbf{b}_r^*$ in exact rational arithmetic.
Set $\boldsymbol{\lambda} := \boldsymbol{\lambda}_{T,0}$.
for $i = r$ **down to** 1 **do**
  Set $c_i := \langle \boldsymbol{\lambda}, \mathbf{b}_i^* \rangle / \|\mathbf{b}_i^*\|_2^2$.
  Set $\widehat{c}_i := \lfloor c_i + \frac{1}{2} \rfloor$. ▷ half-integers are rounded toward $+\infty$
  Set $\boldsymbol{\lambda} := \boldsymbol{\lambda} - \widehat{c}_i \mathbf{b}_i$.
end for
Set $\boldsymbol{\lambda}_{\mathrm{dec}} := \boldsymbol{\lambda}$.

Stage 2: Greedy ℓ_∞-local refinement.

for $p = 1$ **to** *passes* **do**
  Set *improved* := **False**.
  for $i = 1$ **to** r **do**
    Let j be the smallest index satisfying $|\boldsymbol{\lambda}[j]| = \|\boldsymbol{\lambda}\|_\infty$.
    if $\mathbf{b}_i[j] = 0$ **then**
      continue.
    end if
    Set $\alpha_0 := \lfloor -\boldsymbol{\lambda}[j] / \mathbf{b}_i[j] + \frac{1}{2} \rfloor$.
    Set $\boldsymbol{\lambda}_{\mathrm{best}} := \boldsymbol{\lambda}$ and $w_{\mathrm{best}} := \|\boldsymbol{\lambda}\|_\infty$.
    for $\alpha = \alpha_0 - \mathit{window}$ **to** $\alpha_0 + \mathit{window}$ **do**
      Set $\boldsymbol{\lambda}_\alpha := \boldsymbol{\lambda} + \alpha \mathbf{b}_i$ and $w_\alpha := \|\boldsymbol{\lambda}_\alpha\|_\infty$.
      if $w_\alpha < w_{\mathrm{best}}$ **then**
        Set $\boldsymbol{\lambda}_{\mathrm{best}} := \boldsymbol{\lambda}_\alpha$ and $w_{\mathrm{best}} := w_\alpha$.
      end if
    end for
    if $w_{\mathrm{best}} < \|\boldsymbol{\lambda}\|_\infty$ **then**
      Set $\boldsymbol{\lambda} := \boldsymbol{\lambda}_{\mathrm{best}}$ and *improved* := **True**.
    end if
  end for
  if not *improved* **then**
    break.
  end if
end for
return $\boldsymbol{\lambda}$.

Consequently, if $\mathbf{a} - \mathbf{b} \in K_T$, then $\mathbf{a} + K_T = \mathbf{b} + K_T$, and hence $\mathsf{NP}_{\mathbf{B}'_K}(\mathbf{a}) = \mathsf{NP}_{\mathbf{B}'_K}(\mathbf{b})$. Therefore, $\boldsymbol{\lambda}_{\mathrm{dec}}$ is independent of the transformation matrix used to obtain $\boldsymbol{\lambda}_{T,0}$. Since the row-HNF basis $\mathbf{B}_K$, the exact-arithmetic LLL conventions, and all iteration and tie-breaking rules in Algorithm 4 are fixed, the final output is a deterministic function of $\mathbf{G}_T$ and the stated public parameters.

Appendix F Heuristic Estimates for Reconstruction Coefficients

This appendix heuristically relates the certificate in (8) to the sparsity of $\mathbf{G}$. These estimates are not used in any proof.

F.1 Kernel-Lattice Scale from Sparsity

The certified bound depends on the conditioning of $\mathbf{G}_T$ and the Gram-Schmidt lengths of the reduced basis of $K_T := \ker_{\mathbb{Z}}(\mathbf{G}_T^\top)$. We first relate the determinant scale of K_T to the sparsity of $\mathbf{G}$.

Assume that $\|\mathbf{g}\|_\infty \le \beta_{\mathbf{g}}$. Since every row of $\hat{\mathbf{G}}$ has exactly $c_{\exp}$ nonzero entries in $\{-1, 1\}$, $\|\mathbf{G}[i,\cdot]\|_2^2 \le \beta_{\mathbf{g}}^2 + c_{\exp}$. Define $R_{\mathbf{G}} := \sqrt{\beta_{\mathbf{g}}^2 + c_{\exp}}$. Let $r_T := \operatorname{rank}_{\mathbb{R}}(\mathbf{A}_T)$, where $\mathbf{A}_T := \mathbf{G}_T^\top$. Choose $I_T \subseteq [k]$ with $|I_T| = r_T$ such that $\mathbf{A}'_T := \mathbf{A}_T[I_T, \cdot] \in \mathbb{Z}^{r_T \times |T|}$ has full row rank. Because the rows of $\mathbf{A}'_T$ span $\operatorname{Row}_{\mathbb{Q}}(\mathbf{A}_T)$, $\ker_{\mathbb{Z}}(\mathbf{A}'_T) = \ker_{\mathbb{Z}}(\mathbf{A}_T) = K_T$. Let Δ_T be the greatest common divisor of the $r_T \times r_T$ minors of $\mathbf{A}'_T$. By the standard determinant formula for integer kernel lattices,

$$\det(K_T) = \frac{\sqrt{\det(\mathbf{A}'_T \mathbf{A}'^\top_T)}}{\Delta_T}. \tag{F1}$$

By Cauchy-Binet, the squared numerator in (F1) is the sum of the squares of the $r_T \times r_T$ minors of $\mathbf{A}'_T$. For every $J \subseteq T$ with $|J| = r_T$, the minor $\det(\mathbf{A}'_T[\cdot, J])$ is the determinant of the transpose of $\mathbf{G}_T[J, I_T]$, whose rows have Euclidean norm at most $R_{\mathbf{G}}$. Hence, Hadamard's inequality gives $|\det(\mathbf{A}'_T[\cdot, J])| \le R_{\mathbf{G}}^{r_T}$. There are $\binom{|T|}{r_T}$ such minors, and therefore

$$\det(K_T) \le \frac{\binom{|T|}{r_T}^{1/2} R_{\mathbf{G}}^{r_T}}{\Delta_T}. \tag{F2}$$

Dropping $\Delta_T \ge 1$ gives a conservative bound, whereas $\Delta_T > 1$ makes the kernel lattice denser than the corresponding Euclidean minor scale alone suggests.

If $\mathbf{G}_T$ has full column rank, then $r_T = k$, so $I_T = [k]$, $\mathbf{A}'_T = \mathbf{A}_T$, and $\det(\mathbf{A}'_T \mathbf{A}'^\top_T) = \det(\mathbf{G}_T^\top \mathbf{G}_T)$. Hadamard's inequality for the positive semidefinite Gram matrix $\mathbf{G}_T^\top \mathbf{G}_T$ then bounds this determinant by the product of the squared column norms of $\mathbf{G}_T$. The first column satisfies $\|(\mathbf{g})_T\|_2^2 \le \beta_{\mathbf{g}}^2 |T|$. For $j \in [m]$, where $m = k - 1$, let $\rho_j := \|\hat{\mathbf{G}}_T[\cdot, j]\|_2^2$ be the squared norm of the j-th masking column, which equals its number of nonzero entries. Since every row of $\hat{\mathbf{G}}_T$ has exactly $c_{\exp}$ nonzero entries in $\{-1, 1\}$, we have $\sum_{j=1}^m \rho_j = c_{\exp}|T|$. The inequality of arithmetic and geometric means then gives $\prod_{j=1}^m \rho_j \le \left(\frac{1}{m}\sum_{j=1}^m \rho_j\right)^m = (c_{\exp}|T|/m)^m$, and hence $\det(\mathbf{G}_T^\top \mathbf{G}_T) \le \beta_{\mathbf{g}}^2 |T| \, (c_{\exp}|T|/m)^m$. Substituting this bound into (F1) yields

$$\det(K_T) \le \frac{\beta_{\mathbf{g}} \sqrt{|T|} \, (c_{\exp}|T|/m)^{m/2}}{\Delta_T}, \tag{F3}$$

which avoids the binomial factor in the full-rank case $\det(K_T) \le \binom{|T|}{k}^{1/2} R_{\mathbf{G}}^k/\Delta_T$ of (F2).

For $d_T := |T| - r_T > 0$, define the determinant scale $D_T := \det(K_T)^{1/d_T}$. Thus, sparsity controls a coarse average scale of K_T through $R_{\mathbf{G}}$, or through $c_{\exp}|T|/m$ in the full-rank case, but does not determine the shape of its reduced basis or the ℓ_∞-norm of the decoded reconstruction vector.

F.2 Heuristic Gram-Schmidt Profile

To estimate the Babai term from the determinant scale, we use the geometric-series assumption for the Gram-Schmidt profile of the reduced kernel basis. Assume that $d_T > 0$, and let $\delta_{\mathrm{RHF}} > 1$ be the modeled root-Hermite factor of the LLL-reduced basis. Then

$$\|\mathbf{b}_i^*\|_2 \approx \delta_{\mathrm{RHF}}^{d_T-2i+1} D_T, \qquad i \in [d_T]. \tag{F4}$$

Substituting (F4) into (9) and (10) gives

$$B_{\boldsymbol{\lambda}}(T) \lesssim \left(\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 + \frac{1}{4}\det(K_T)^{2/d_T} \sum_{i=1}^{d_T} \delta_{\mathrm{RHF}}^{2(d_T-2i+1)} \right)^{1/2}. \tag{F5}$$

If the reduced basis is approximately balanced, then $\sum_{i=1}^{d_T} \delta_{\mathrm{RHF}}^{2(d_T-2i+1)} \approx d_T$, and hence $R_{\mathsf{Babai}}(\mathbf{B}'_K) \approx \frac{\sqrt{d_T}}{2} D_T$ and $B_{\boldsymbol{\lambda}}(T) \lesssim \left(\|\boldsymbol{\lambda}_{T,\mathbb{R}}\|_2^2 + \frac{d_T}{4} D_T^2\right)^{1/2}$. This approximation is most relevant when $d_T = |T| - \mathrm{rank}(\mathbf{G}_T)$ is small. The exponents in (F4) decrease symmetrically from $d_T - 1$ at $i = 1$ to $-(d_T - 1)$ at $i = d_T$, so pairing symmetric terms shows that the sum is at least d_T, with equality only if $\delta_{\mathrm{RHF}} = 1$, and its deviation from d_T is governed by the largest term $\delta_{\mathrm{RHF}}^{2(d_T-1)}$. Explicitly, $\sum_{i=1}^{d_T} \delta_{\mathrm{RHF}}^{2(d_T-2i+1)} = \left(\delta_{\mathrm{RHF}}^{2d_T} - \delta_{\mathrm{RHF}}^{-2d_T}\right)/\left(\delta_{\mathrm{RHF}}^{2} - \delta_{\mathrm{RHF}}^{-2}\right)$, which remains close to d_T only while $\delta_{\mathrm{RHF}}^{2(d_T-1)} \approx 1$.

For LLL with Lovász parameter δ_{LLL}, the first basis vector provably satisfies $\|\mathbf{b}_1\|_2 \le (\delta_{\mathsf{LLL}} - 1/4)^{-(d_T-1)/4} \det(K_T)^{1/d_T}$ [58], so the root-Hermite factor is at most $(\delta_{\mathsf{LLL}} - 1/4)^{-1/4} \approx 1.078$ for $\delta_{\mathsf{LLL}} = 0.99$ (as in Section 6.3), and at most $(4/3)^{1/4} \approx 1.075$ in the limit $\delta_{\mathsf{LLL}} \to 1$. In practice, LLL performs considerably better on average, with $\delta_{\mathrm{RHF}} \approx 1.02$, as first observed experimentally by Nguyen and Stehlé [59] and later confirmed by Gama and Nguyen [60]. These values concern the first basis vector; the geometric-series profile itself remains a heuristic. This estimate provides parameter-selection intuition but does not replace the instance-specific certificate in (8) or the observed values returned by Algorithm 1. For the concrete instantiations, we select $\beta_{\boldsymbol{\lambda}}$ from the observed values $B_{\boldsymbol{\lambda}}(T)$ using the safety margin specified in Section 6.3.2.

Appendix G Structural Noise-Cancellation Analysis

This appendix supports the structural setup condition in Section 5.1. We first characterize exact noise cancellation over $\mathbb{Z}$ and $\mathbb{Q}$, and then show that failure of integral

column-module membership yields a cancellation attack after reduction modulo an adversarially chosen modulus. For a fixed private coalition $U \in \Gamma_0$, recall that the restricted share matrix is

$$\mathbf{S}_U = (\mathbf{g})_U \mathbf{s}^\top + \hat{\mathbf{G}}_U \mathbf{R}. \tag{G1}$$

G.1 Exact Noise Cancellation over $\mathbb{Z}$ and $\mathbb{Q}$

Fix public parameters pp and a private coalition $U \in \Gamma_0$. The following lemma characterizes an exact cancellation witness.

Lemma 37 *Let $\Sigma \in \{\mathbb{Z}, \mathbb{Q}\}$. Suppose that there exists $\boldsymbol{\alpha}_U \in \Sigma^{|U|}$ such that $\hat{\mathbf{G}}_U^\top \boldsymbol{\alpha}_U = \mathbf{0}$ and $c := \boldsymbol{\alpha}_U^\top (\mathbf{g})_U \neq 0$. Then $\boldsymbol{\alpha}_U^\top \mathbf{S}_U = c\mathbf{s}^\top$. Thus, the shares indexed by U reveal a noiseless nonzero scaled secret $c\mathbf{s}$. Over $\mathbb{Q}$, the secret is recoverable as $\mathbf{s} = c^{-1}\mathbf{S}_U^\top \boldsymbol{\alpha}_U$.*

Proof Multiplying (G1) by $\boldsymbol{\alpha}_U^\top$ gives $\boldsymbol{\alpha}_U^\top \mathbf{S}_U = \boldsymbol{\alpha}_U^\top (\mathbf{g})_U \mathbf{s}^\top + \boldsymbol{\alpha}_U^\top \hat{\mathbf{G}}_U \mathbf{R} = c\mathbf{s}^\top$. If the computation is performed over $\mathbb{Q}$, then $c \neq 0$ is invertible, giving the stated recovery expression. □

Since $\mathbf{G}_U = \begin{bmatrix}(\mathbf{g})_U & \hat{\mathbf{G}}_U\end{bmatrix}$, we always have $\ker_\Sigma(\mathbf{G}_U^\top) \subseteq \ker_\Sigma(\hat{\mathbf{G}}_U^\top)$. Strict containment is equivalent to the existence of a cancellation witness as in Lemma 37.

Corollary 5 *Let $\Sigma \in \{\mathbb{Z}, \mathbb{Q}\}$. If $\ker_\Sigma(\mathbf{G}_U^\top) \subsetneq \ker_\Sigma(\hat{\mathbf{G}}_U^\top)$, then an exact noise-cancellation witness exists over Σ. In particular, the shares indexed by U reveal a noiseless nonzero scaled secret.*

We next express rational cancellation through column-space separation.

Lemma 38 (Column-space separation) *Let Σ be a field, let $W \subseteq \Sigma^r$ be a subspace, and let $\mathbf{v} \in \Sigma^r \setminus W$. Then there exists $\mathbf{x} \in \Sigma^r$ such that $\langle \mathbf{x}, \mathbf{w} \rangle = 0$ for every $\mathbf{w} \in W$, but $\langle \mathbf{x}, \mathbf{v} \rangle \neq 0$.*

Proof The double-annihilator property gives $(W^\perp)^\perp = W$. If every $\mathbf{x} \in W^\perp$ were also orthogonal to $\mathbf{v}$, then $\mathbf{v} \in (W^\perp)^\perp = W$, contradicting $\mathbf{v} \notin W$. □

Lemma 39 (Rational cancellation criterion) *Rational noise cancellation is impossible for U if and only if $(\mathbf{g})_U \in \mathrm{Col}_\mathbb{Q}(\hat{\mathbf{G}}_U)$.*

Proof If $(\mathbf{g})_U \notin \mathrm{Col}_\mathbb{Q}(\hat{\mathbf{G}}_U)$, Lemma 38 gives $\boldsymbol{\alpha}_U \in \ker_\mathbb{Q}(\hat{\mathbf{G}}_U^\top)$ with $\boldsymbol{\alpha}_U^\top (\mathbf{g})_U \neq 0$, and Lemma 37 gives a cancellation attack. Conversely, if $(\mathbf{g})_U = \hat{\mathbf{G}}_U \mathbf{a}$ for some $\mathbf{a} \in \mathbb{Q}^{k-1}$, then every $\mathbf{x} \in \ker_\mathbb{Q}(\hat{\mathbf{G}}_U^\top)$ satisfies $\mathbf{x}^\top (\mathbf{g})_U = \mathbf{x}^\top \hat{\mathbf{G}}_U \mathbf{a} = 0$. Hence no rational cancellation witness exists. □

G.2 Modular Cancellation under Adversarial Post-Processing

For an integer matrix or vector, an overline denotes component-wise reduction modulo the modulus under consideration. We use the Smith normal form to show that failure of integral column-module membership is detected modulo some integer.

Lemma 40 *Let* $\mathbf{A} \in \mathbb{Z}^{r\times m}$*, let* $\mathbf{v} \in \mathbb{Z}^r$*, and let* $\mathbf{PAQ} = \begin{bmatrix} d_1 & & & \\ & \ddots & & \\ & & d_\rho & \\ & & & \mathbf{0} \end{bmatrix}$ *be a Smith normal form of* $\mathbf{A}$*, where* $\mathbf{P} \in GL_r(\mathbb{Z})$*,* $\mathbf{Q} \in GL_m(\mathbb{Z})$*,* $0 < d_1 \mid \cdots \mid d_\rho$*, and* $0 \le \rho \le \min\{r, m\}$*. Write* $\mathbf{v}' := \mathbf{Pv}$*. Then* $\mathbf{v} \in \mathrm{Col}_{\mathbb{Z}}(\mathbf{A})$ *if and only if* $d_i \mid \mathbf{v}'[i]$ *for every* $i \in [\rho]$ *and* $\mathbf{v}'[i] = 0$ *for every* $i > \rho$*.*

Proof Since $\mathbf{P}$ and $\mathbf{Q}$ are unimodular, they induce isomorphisms of the relevant integer modules. The claim therefore follows by solving the diagonal Smith-normal-form system. □

Lemma 41 *Let* $\mathbf{A} \in \mathbb{Z}^{r\times m}$ *and* $\mathbf{v} \in \mathbb{Z}^r$*. If* $\mathbf{v} \notin \mathrm{Col}_{\mathbb{Z}}(\mathbf{A})$*, then there exists* $q \ge 2$ *such that* $\overline{\mathbf{v}} \notin \mathrm{Col}_{\mathbb{Z}_q}(\overline{\mathbf{A}})$*.*

Proof Use the notation of Lemma 40. Integral non-membership implies either that $\mathbf{v}'[i] \neq 0$ for some $i > \rho$, or that $d_i \nmid \mathbf{v}'[i]$ for some $i \le \rho$. In the first case, choose a prime q not dividing $\mathbf{v}'[i]$. The i-th row of the Smith normal form is zero modulo q, whereas $\mathbf{v}'[i] \not\equiv 0 \pmod q$. In the second case, choose a prime p for which $v_p(\mathbf{v}'[i]) < v_p(d_i)$, where $v_p(\cdot)$ denotes the p-adic valuation function, and put $q := p^{v_p(\mathbf{v}'[i])+1}$. Then $d_i \equiv 0 \pmod q$, while $\mathbf{v}'[i] \not\equiv 0 \pmod q$. Thus, in either case, $\overline{\mathbf{v}'}$ does not lie in the modular column module of the reduced Smith normal form. Since $\mathbf{P}$ and $\mathbf{Q}$ remain invertible modulo q, the same conclusion holds for $\mathbf{v}$ and $\mathbf{A}$. □

Lemma 42 (Noise-cancellation attack modulo q) *If* $(\mathbf{g})_U \notin \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$*, then there exist* $q \ge 2$ *and a nonzero vector* $\overline{\mathbf{x}} \in \mathbb{Z}_q^{|U|}$ *such that* $\overline{\mathbf{x}}^\top \overline{\hat{\mathbf{G}}_U} = \mathbf{0}^\top$ *and* $\overline{c} := \overline{\mathbf{x}}^\top \overline{(\mathbf{g})_U} \neq 0$*. Consequently,* $\overline{\mathbf{x}}^\top \overline{\mathbf{S}_U} = \overline{c}\,\overline{\mathbf{s}}^\top \pmod q$*. Thus, any two secrets having different images under multiplication by* $\overline{c}$ *induce perfectly distinguishable modular share views. If* $\overline{c}$ *is a unit in* $\mathbb{Z}_q$*, then* $\overline{\mathbf{s}}$ *is fully recoverable.*

Proof By Lemma 41, there exists $q \ge 2$ such that $\overline{(\mathbf{g})_U} \notin \mathrm{Col}_{\mathbb{Z}_q}(\overline{\hat{\mathbf{G}}_U})$. The double-annihilator property over the finite Frobenius ring $\mathbb{Z}_q$ (see, e.g., [64] and [65, Corollary 3.2]) gives a vector $\overline{\mathbf{x}} \in \mathbb{Z}_q^{|U|}$ that annihilates every column of $\overline{\hat{\mathbf{G}}_U}$ but not $\overline{(\mathbf{g})_U}$. Reducing (G1) modulo q and multiplying by $\overline{\mathbf{x}}^\top$ gives

$$\overline{\mathbf{x}}^\top \overline{\mathbf{S}_U} = \overline{\mathbf{x}}^\top \overline{(\mathbf{g})_U}\,\overline{\mathbf{s}}^\top + \overline{\mathbf{x}}^\top \overline{\hat{\mathbf{G}}_U}\,\overline{\mathbf{R}} = \overline{c}\,\overline{\mathbf{s}}^\top \pmod q. \tag{G2}$$

The distinguishability and recovery statements follow immediately. □

G.3 Proof of the Structural Proposition

Proof of Proposition 1 Suppose that $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$. Then there exists $\mathbf{a} \in \mathbb{Z}^{k-1}$ such that $(\mathbf{g})_U = \hat{\mathbf{G}}_U \mathbf{a}$. First, the inclusion $\mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U) \subseteq \mathrm{Col}_{\mathbb{Q}}(\hat{\mathbf{G}}_U)$ implies $(\mathbf{g})_U \in \mathrm{Col}_{\mathbb{Q}}(\hat{\mathbf{G}}_U)$. Rational noise cancellation is therefore impossible by Lemma 39.

Second, $\mathbf{G}_U = \left[(\mathbf{g})_U \ \hat{\mathbf{G}}_U\right] = \hat{\mathbf{G}}_U \left[\mathbf{a} \ \mathbf{I}_{k-1}\right]$. Hence $\mathrm{Col}_{\mathbb{Z}}(\mathbf{G}_U) = \mathrm{Col}_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$, the reverse inclusion being immediate. Moreover, $\ker_{\mathbb{Z}}(\mathbf{G}_U^\top) \subseteq \ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U^\top)$ always holds, and conversely, if $\mathbf{x} \in \ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U^\top)$, then $(\mathbf{g})_U^\top \mathbf{x} = \mathbf{a}^\top \hat{\mathbf{G}}_U^\top \mathbf{x} = 0$, and hence $\mathbf{x} \in \ker_{\mathbb{Z}}(\mathbf{G}_U^\top)$. Therefore, $\ker_{\mathbb{Z}}(\mathbf{G}_U^\top) = \ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U^\top)$, which excludes integer noise cancellation.

Finally, reducing $(\mathbf{g})_U = \hat{\mathbf{G}}_U \mathbf{a}$ modulo any $q' \geq 2$ gives $\overline{(\mathbf{g})_U} = \overline{\hat{\mathbf{G}}_U}\,\overline{\mathbf{a}}$, so $\overline{(\mathbf{g})_U} \in \mathrm{Col}_{\mathbb{Z}_{q'}}(\overline{\hat{\mathbf{G}}_U})$. Consequently, every vector annihilating $\overline{\hat{\mathbf{G}}_U}$ also annihilates $\overline{(\mathbf{g})_U}$, and modular noise cancellation is impossible. □

Appendix H Deferred Proofs for the Privacy Analysis

This appendix contains the proofs omitted from Section 5. Throughout, $(\mathcal{U}, \mathbf{c})$ is a fixed admissible public realization, the notation is that of Table 3, and empty block rows and columns are omitted.

H.1 Joint Precision Bound

Proof of Lemma 20 By (28), $\tilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}$ has blocks $\mathbf{y}_{\mathrm{raw}}$, $\mathbf{r}_{\mathrm{in}}$, and $\mathbf{y}_{\mathrm{out}}$. The masks in $\mathbf{r}_{\mathrm{in}}$ are independent of $\mathbf{r}_{\mathbf{R}}$ and of the distinct masks appearing in $\mathbf{y}_{\mathrm{out}}$, so all parameter blocks between the $\mathcal{S}_{\mathrm{in}}$ coordinates and the remaining coordinates vanish, and the $\mathcal{S}_{\mathrm{in}}$ block equals $\sigma_{\mathbf{r}}^2\mathbf{I}$. Let $\mathbf{P}_{\mathrm{cor}}$ be the permutation that groups the raw-share and $\mathcal{S}_{\mathrm{out}}$ coordinates first. Then $\mathbf{P}_{\mathrm{cor}}\boldsymbol{\Sigma}_{\tilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}\mathbf{P}_{\mathrm{cor}}^\top = \boldsymbol{\Sigma}_{\mathrm{cor}} \oplus \sigma_{\mathbf{r}}^2\mathbf{I}$. Both diagonal blocks are invertible, so $\mathbf{P}_{\mathrm{cor}}\boldsymbol{\Sigma}_{\tilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}^{-1}\mathbf{P}_{\mathrm{cor}}^\top = \boldsymbol{\Sigma}_{\mathrm{cor}}^{-1} \oplus \sigma_{\mathbf{r}}^{-2}\mathbf{I}$, and undoing the permutation gives (31). The norm identity holds because permutations preserve the spectral norm and the spectral norm of a block-diagonal positive-definite matrix is the maximum over its diagonal blocks. □

Proof of Lemma 21 By (29), the left-hand side of (32) equals $\widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}}^\top \boldsymbol{\Sigma}_{\tilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}^{-1} \widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}}$. Substituting (28) and (31), every term involving an $\mathcal{S}_{\mathrm{in}}$ row or column is multiplied by the zero block of $\widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}}$ and vanishes. The remaining blocks give $\widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}}^\top \boldsymbol{\Sigma}_{\tilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}^{-1} \widetilde{\mathbf{M}}_{\mathcal{U},\mathbf{c}} = \begin{bmatrix}\mathbf{M}_{\mathrm{raw}} \\ \mathbf{M}_{\mathrm{out}}\end{bmatrix}^\top \boldsymbol{\Sigma}_{\mathrm{cor}}^{-1} \begin{bmatrix}\mathbf{M}_{\mathrm{raw}} \\ \mathbf{M}_{\mathrm{out}}\end{bmatrix}$, which proves (32). □

Proof of Lemma 22 By (33) and (34), $\boldsymbol{\Sigma}_{\mathrm{raw}}\mathbf{B} = \boldsymbol{\Sigma}_{\mathrm{raw,out}}$ and $\mathbf{B}^\top\boldsymbol{\Sigma}_{\mathrm{raw}}\mathbf{B} + \mathbf{C} = \boldsymbol{\Sigma}_{\mathrm{out}}$, and hence

$$\boldsymbol{\Sigma}_{\mathrm{cor}} = \begin{bmatrix}\mathbf{I} & \mathbf{O} \\ \mathbf{B}^\top & \mathbf{I}\end{bmatrix} \begin{bmatrix}\boldsymbol{\Sigma}_{\mathrm{raw}} & \mathbf{O} \\ \mathbf{O} & \mathbf{C}\end{bmatrix} \begin{bmatrix}\mathbf{I} & \mathbf{B} \\ \mathbf{O} & \mathbf{I}\end{bmatrix}.$$

Taking inverses gives $\boldsymbol{\Sigma}_{\mathrm{cor}}^{-1} = \mathbf{K}^\top(\boldsymbol{\Sigma}_{\mathrm{raw}}^{-1} \oplus \mathbf{C}^{-1})\mathbf{K}$ with $\mathbf{K} := \begin{bmatrix}\mathbf{I} & \mathbf{O} \\ -\mathbf{B}^\top & \mathbf{I}\end{bmatrix}$. Since $\mathbf{K}\mathbf{M}_{\mathrm{cor}} = \left[\mathbf{M}_{\mathrm{raw}}^\top \ (\mathbf{M}_{\mathrm{out}} - \mathbf{B}^\top\mathbf{M}_{\mathrm{raw}})^\top\right]^\top$, expanding the block-diagonal quadratic form proves (35). □

Proof of Lemma 23 By (37), $\mathbf{\Sigma}_{\text{raw}} = \sigma_{\mathbf{R}}^2(\hat{\mathbf{G}}_{U_{\text{raw}}}\hat{\mathbf{G}}_{U_{\text{raw}}}^\top) \otimes \mathbf{I}_D$. Full row rank of $\hat{\mathbf{G}}_{U_{\text{raw}}}$ gives $\mathbf{\Sigma}_{\text{raw}} \succ 0$ with $\lambda_{\min}(\mathbf{\Sigma}_{\text{raw}}) = \sigma_{\mathbf{R}}^2\sigma_{\min}^2(\hat{\mathbf{G}}_{U_{\text{raw}}})$, which proves the inverse-norm identity.

Suppose that $\mathcal{S}_{\text{out}} \neq \emptyset$. By (33), (37), and (38), $\mathbf{B} = (\hat{\mathbf{G}}_{\text{raw}}^+)^\top \mathbf{T}_{\text{out}}^\top$. Let $\hat{\mathbf{G}}_{\mathcal{S}_{\text{out}}}$ stack the rows $\hat{\mathbf{g}}_i^\top$ over $(s,i) \in \mathcal{S}_{\text{out}}$, with one copy per occurrence of i, and let $\mathbf{L}_{\mathcal{S}_{\text{out}}}$ be block diagonal with block $\mathbf{L}_s$ at coordinate (s,i). Then

$$\mathbf{T}_{\text{out}} = \mathbf{L}_{\mathcal{S}_{\text{out}}}(\hat{\mathbf{G}}_{\mathcal{S}_{\text{out}}} \otimes \mathbf{I}_D), \tag{H1}$$

with $\|\mathbf{L}_{\mathcal{S}_{\text{out}}}\|_2 \leq \omega_c$ and $\|\hat{\mathbf{G}}_{\mathcal{S}_{\text{out}}}\|_2 \leq \sqrt{Q}\|\hat{\mathbf{G}}\|_2$; the last bound holds because $\hat{\mathbf{G}}_{\mathcal{S}_{\text{out}}} = \mathbf{E}_{\text{out}}\hat{\mathbf{G}}$ for a row-selection matrix $\mathbf{E}_{\text{out}}$ in which each row of $\hat{\mathbf{G}}$ is selected at most Q times, so $\|\mathbf{E}_{\text{out}}\|_2 \leq \sqrt{Q}$. Hence $\|\mathbf{B}\|_2 \leq \|\mathbf{L}_{\mathcal{S}_{\text{out}}}\|_2\|\hat{\mathbf{G}}_{\mathcal{S}_{\text{out}}}\|_2\|\hat{\mathbf{G}}_{U_{\text{raw}}}^+\|_2 \leq \omega_c\sqrt{Q}\|\hat{\mathbf{G}}\|_2/\sigma_{\min}(\hat{\mathbf{G}}_{U_{\text{raw}}})$. Finally, $\|\hat{\mathbf{G}}\|_2 = \|\mathbf{PH}\|_2 = \|\mathbf{H}\|_2 \leq \|\mathbf{H}\|_F = \sqrt{c_{\text{exp}}n}$. □

Proof of Lemma 24 Substituting (37), (38), and (39) into the definition of $\mathbf{C}$ in (34) gives

$$\mathbf{C} = \sigma_{\mathbf{R}}^2\mathbf{T}_{\text{out}}\left(\mathbf{I} - \hat{\mathbf{G}}_{\text{raw}}^+\hat{\mathbf{G}}_{\text{raw}}\right)\mathbf{T}_{\text{out}}^\top + \sigma_{\mathbf{r}}^2\mathbf{I}.$$

Since $\hat{\mathbf{G}}_{\text{raw}}$ has full row rank, $\hat{\mathbf{G}}_{\text{raw}}^+\hat{\mathbf{G}}_{\text{raw}}$ is the orthogonal projector onto its row space, so $\mathbf{I} - \hat{\mathbf{G}}_{\text{raw}}^+\hat{\mathbf{G}}_{\text{raw}} \succeq 0$, and congruence by $\mathbf{T}_{\text{out}}$ preserves this. Thus, $\mathbf{C} \succeq \sigma_{\mathbf{r}}^2\mathbf{I} \succ 0$ and $\|\mathbf{C}^{-1}\|_2 = 1/\lambda_{\min}(\mathbf{C}) \leq \sigma_{\mathbf{r}}^{-2}$. □

Proof of Lemma 25 Since $\mathbf{M}_{\text{raw}} = (\mathbf{g})_{U_{\text{raw}}} \otimes \mathbf{I}_D$, we have $\|\mathbf{M}_{\text{raw}}\|_2 = \|(\mathbf{g})_{U_{\text{raw}}}\|_2 \leq \beta_{\mathbf{g}}\sqrt{k_p}$. The matrix $\mathbf{M}_{\text{out}}$ stacks $\mathbf{g}[i]\mathbf{L}_s$ over $(s,i) \in \mathcal{S}_{\text{out}}$, so $\|\mathbf{M}_{\text{out}}\|_2^2 \leq \sum_{(s,i)\in\mathcal{S}_{\text{out}}} \mathbf{g}[i]^2\|\mathbf{L}_s\|_2^2 \leq \beta_{\mathbf{g}}^2\omega_c^2|\mathcal{S}_{\text{out}}| \leq \beta_{\mathbf{g}}^2\omega_c^2 nQ$. Finally, $\|\mathbf{M}_{\text{out}} - \mathbf{B}^\top\mathbf{M}_{\text{raw}}\|_2 \leq \|\mathbf{M}_{\text{out}}\|_2 + \|\mathbf{B}\|_2\|\mathbf{M}_{\text{raw}}\|_2$, and Lemma 23 gives (40). □

Proof of Lemma 26 We first verify the hypotheses of Lemmas 20 to 22. By Lemma 23, $\mathbf{\Sigma}_{\text{raw}} \succ 0$. If $\mathcal{S}_{\text{out}} = \emptyset$, then $\mathbf{\Sigma}_{\text{cor}} = \mathbf{\Sigma}_{\text{raw}} \succ 0$. Otherwise, Lemma 24 gives $\mathbf{C} \succ 0$, and the block factorization of $\mathbf{\Sigma}_{\text{cor}}$ in the proof of Lemma 22, a congruence of $\mathbf{\Sigma}_{\text{raw}} \oplus \mathbf{C}$ by an invertible matrix, gives $\mathbf{\Sigma}_{\text{cor}} \succ 0$. Together with $\sigma_{\mathbf{r}} > 0$, these are the required hypotheses.

If $\mathcal{S}_{\text{out}} = \emptyset$, Lemma 21 leaves only the raw-share coordinates, so the precision term equals $\mathbf{M}_{\text{raw}}^\top\mathbf{\Sigma}_{\text{raw}}^{-1}\mathbf{M}_{\text{raw}}$, whose norm is at most $\|\mathbf{M}_{\text{raw}}\|_2^2\|\mathbf{\Sigma}_{\text{raw}}^{-1}\|_2 \leq \beta_{\mathbf{g}}^2 k_p/(\sigma_{\mathbf{R}}^2\sigma_{\min}^2(\hat{\mathbf{G}}_{U_{\text{raw}}}))$ by Lemmas 23 and 25; the second term of (41) is nonnegative. If $\mathcal{S}_{\text{out}} \neq \emptyset$, Lemmas 21 and 22 and the triangle inequality give

$$\left\|\mathbf{M}_{\mathcal{U},\mathbf{c}}^\top\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}^{-1}\mathbf{M}_{\mathcal{U},\mathbf{c}}\right\|_2 \leq \|\mathbf{M}_{\text{raw}}\|_2^2\|\mathbf{\Sigma}_{\text{raw}}^{-1}\|_2 + \left\|\mathbf{M}_{\text{out}} - \mathbf{B}^\top\mathbf{M}_{\text{raw}}\right\|_2^2\|\mathbf{C}^{-1}\|_2.$$

The first term is bounded as before, and the second by Lemma 24 and (40). □

Verification of Remark 5

If $U_{\text{raw}} = \emptyset$ and $\mathcal{S} \neq \emptyset$, then $\mathcal{S}_{\text{in}} = \emptyset$ and $\mathcal{S}_{\text{out}} = \mathcal{S}$, so there are no raw-share or $\mathcal{S}_{\text{in}}$ blocks, $\mathbf{M}_{\mathcal{U},\mathbf{c}} = \mathbf{M}_{\text{out}}$, and $\mathbf{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}} = \mathbf{\Sigma}_{\text{out}} = \sigma_{\mathbf{R}}^2\mathbf{T}_{\text{out}}\mathbf{T}_{\text{out}}^\top + \sigma_{\mathbf{r}}^2\mathbf{I} \succeq \sigma_{\mathbf{r}}^2\mathbf{I}$. Lemma 25 assumes $|U_{\text{raw}}| \geq 1$, but the argument for its bound $\|\mathbf{M}_{\text{out}}\|_2^2 \leq \beta_{\mathbf{g}}^2\omega_c^2|\mathcal{S}_{\text{out}}|$ uses only $\|\mathbf{g}\|_\infty \leq \beta_{\mathbf{g}}$ and $\|\mathbf{L}_s\|_2 \leq \omega_c$, and therefore also applies here. Hence the precision term is at most $\|\mathbf{M}_{\text{out}}\|_2^2\|\mathbf{\Sigma}_{\text{out}}^{-1}\|_2 \leq \beta_{\mathbf{g}}^2\omega_c^2 nQ/\sigma_{\mathbf{r}}^2$. If $U_{\text{raw}} \neq \emptyset$ and $\mathcal{S}_{\text{out}} = \emptyset$, including $\mathcal{S} = \emptyset$, every response block is eliminated by $\mathbf{W}$, and the first case of the proof of Lemma 26 applies. If $U_{\text{raw}} = \emptyset$ and $\mathcal{S} = \emptyset$, then $\mathbf{M}_{\mathcal{U},\mathbf{c}}$ and $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ are empty, and the precision term is 0 by convention. All remaining patterns are covered by Lemma 26.

H.2 Uniform Parameters

Proof of Lemma 27 Each $\hat{\mathbf{G}}_V$ with $|V| = k_p$ has full row rank, so α_{k_p} is a minimum of finitely many positive numbers. For a nonempty U with $|U| \le k_p$, choose $V \supseteq U$ with $|V| = k_p$. Since $\hat{\mathbf{G}}_U \hat{\mathbf{G}}_U^\top$ is a principal submatrix of $\hat{\mathbf{G}}_V \hat{\mathbf{G}}_V^\top$, eigenvalue interlacing gives $\sigma_{\min}^2(\hat{\mathbf{G}}_U) \ge \sigma_{\min}^2(\hat{\mathbf{G}}_V) \ge \alpha_{k_p}^2$. □

Proof of Lemma 28 If $U_{\text{raw}} \neq \emptyset$, then $1/\sigma_{\min}(\hat{\mathbf{G}}_{U_{\text{raw}}}) \le 1/\alpha_{k_p}$ by Lemma 27, and substituting into Lemma 26 gives Ξ_{MQ}. If $U_{\text{raw}} = \emptyset$ and $\mathcal{S} \neq \emptyset$, Remark 5 gives the bound $\beta_{\mathbf{g}}^2 \omega_c^2 nQ/\sigma_{\mathbf{r}}^2$, which is at most the second term of Ξ_{MQ} because in (43) we have $\left(\sqrt{n} + \sqrt{k_p}\frac{\|\hat{\mathbf{G}}\|_2}{\alpha_{k_p}}\right)^2 \ge n$. The empty view is immediate. The bound does not depend on $\mathbf{c}$, since each c_s enters only through $\|\mathbf{L}_s\|_2 \le \omega_c$. □

Proof of Lemma 29 The inclusion $\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}} \mathbb{Z}^M \subseteq \mathbb{Z}^{M'}$ is immediate. Suppose that $U_{\text{raw}} \neq \emptyset$. By Lemma 7, $\hat{\mathbf{G}}_{U_{\text{raw}}} \mathbb{Z}^{k-1} = \mathbb{Z}^{|U_{\text{raw}}|}$, and tensoring with $\mathbf{I}_D$ gives $\hat{\mathbf{G}}_{\text{raw}} \mathbb{Z}^{(k-1)D} = \mathbb{Z}^{|U_{\text{raw}}|D}$. Given an arbitrary $\begin{bmatrix}\mathbf{a}^\top & \mathbf{b}_{\text{in}}^\top & \mathbf{b}_{\text{out}}^\top\end{bmatrix}^\top \in \mathbb{Z}^{M'}$, choose $\mathbf{v}$ integral with $\hat{\mathbf{G}}_{\text{raw}}\mathbf{v} = \mathbf{a}$; then the source vector $\begin{bmatrix}\mathbf{v}^\top & \mathbf{b}_{\text{in}}^\top & (\mathbf{b}_{\text{out}} - \mathbf{T}_{\text{out}}\mathbf{v})^\top\end{bmatrix}^\top$ maps to it. If $U_{\text{raw}} = \emptyset$, then $\mathcal{S}_{\text{in}} = \emptyset$, and the output block contains $\mathbf{I}_{|\mathcal{S}|D}$. □

Proof of Lemma 30 The proof has three steps: we bound the Gram-Schmidt parameter of $\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_{\text{raw}})$, transfer this bound to the whitened kernel through a linear map of norm at most $1/\sigma_{\text{eff}}$, and apply a basis-dependent smoothing estimate.
Step 1: the unwhitened kernel. Coefficient expansion does not increase the Gram-Schmidt parameter: for every $U \subseteq [n]$,

$$\widetilde{bl}\left(\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U \otimes \mathbf{I}_D)\right) \le \widetilde{bl}\left(\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U)\right). \tag{H2}$$

Indeed, if $(\mathbf{b}_1, \dots, \mathbf{b}_r)$ is an ordered basis of $\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$ and $(\mathbf{f}_c)_{c\in[D]}$ is the standard basis of $\mathbb{Z}^D$, then, up to a coordinate permutation, $\{\mathbf{b}_j \otimes \mathbf{f}_c\}$ is a basis of $\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U \otimes \mathbf{I}_D)$; ordering it by c and then by j gives D mutually orthogonal copies of the original Gram-Schmidt process. If $U_{\text{raw}} \neq \emptyset$, then (H2) with $U = U_{\text{raw}}$ and (48) give $\widetilde{bl}(\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_{\text{raw}})) \le \gamma_{k_p}^{\text{ker}}$. If $U_{\text{raw}} = \emptyset$, then $\hat{\mathbf{G}}_{\text{raw}}$ has no rows, so $\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_{\text{raw}}) = \mathbb{Z}^{(k-1)D}$, whose standard basis gives $\widetilde{bl}(\mathbb{Z}^{(k-1)D}) = 1 \le \gamma_{k_p}^{\text{ker}}$ by the outer maximum with 1 in (48). In both cases,

$$\widetilde{bl}\left(\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_{\text{raw}})\right) \le \gamma_{k_p}^{\text{ker}}. \tag{H3}$$

Step 2: the whitened kernel. Define the injective map $\mathbf{F}_{\mathcal{U},\mathbf{c}}\mathbf{v} := \begin{bmatrix}\sigma_{\mathbf{R}}^{-1}\mathbf{v}^\top & \mathbf{0}^\top & -\sigma_{\mathbf{r}}^{-1}\mathbf{v}^\top\mathbf{T}_{\text{out}}^\top\end{bmatrix}^\top$. Since $\mathbf{\Sigma}_0^{-1/2}$ scales the source blocks by $\sigma_{\mathbf{R}}^{-1}$ and $\sigma_{\mathbf{r}}^{-1}$, (46) gives $\mathbf{\Sigma}_0^{-1/2}\ker_{\mathbb{Z}}(\widetilde{\mathbf{\Phi}}_{\mathcal{U},\mathbf{c}}) = \mathbf{F}_{\mathcal{U},\mathbf{c}}\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_{\text{raw}})$. For every $\mathbf{v}$, $\|\mathbf{F}_{\mathcal{U},\mathbf{c}}\mathbf{v}\|_2^2 \le (\sigma_{\mathbf{R}}^{-2} + \sigma_{\mathbf{r}}^{-2}\|\mathbf{T}_{\text{out}}\|_2^2)\|\mathbf{v}\|_2^2$. The factorization (H1), together with the bounds $\|\mathbf{L}_{\mathcal{S}_{\text{out}}}\|_2 \le \omega_c$ and $\|\hat{\mathbf{G}}_{\mathcal{S}_{\text{out}}}\|_2 \le \sqrt{Q}\|\hat{\mathbf{G}}\|_2$ established after it in the proof of Lemma 23 (an argument that does not use U_{raw}), gives $\|\mathbf{T}_{\text{out}}\|_2^2 \le \omega_c^2 Q\|\hat{\mathbf{G}}\|_2^2$; if $\mathcal{S}_{\text{out}} = \emptyset$, then $\mathbf{T}_{\text{out}}$ is empty and the bound is trivial. By the definition of σ_{eff} in (49), $\|\mathbf{F}_{\mathcal{U},\mathbf{c}}\|_2 \le 1/\sigma_{\text{eff}}$. For every lattice Λ and every linear map $\mathbf{F}$ injective on $\text{span}_{\mathbb{R}}(\Lambda)$, $\widetilde{bl}(\mathbf{F}\Lambda) \le \|\mathbf{F}\|_2\widetilde{bl}(\Lambda)$. Indeed, for an ordered basis $(\mathbf{b}_j)$ of Λ and $V_{j-1} := \text{span}_{\mathbb{R}}\{\mathbf{b}_1, \dots, \mathbf{b}_{j-1}\}$,

the j-th Gram-Schmidt norm of $(\mathbf{F}\mathbf{b}_j)$ is $\operatorname{dist}(\mathbf{F}\mathbf{b}_j, \mathbf{F}V_{j-1}) \le \|\mathbf{F}\|_2 \operatorname{dist}(\mathbf{b}_j, V_{j-1})$, where $\operatorname{dist}(\mathbf{x}, V) := \min_{\mathbf{v}\in V} \|\mathbf{x}-\mathbf{v}\|_2$. Applying this with $\mathbf{F} = \mathbf{F}_{\mathcal{U},\mathbf{c}}$ and (H3) gives

$$\widetilde{bl}\left(\boldsymbol{\Sigma}_0^{-1/2} \ker_{\mathbb{Z}}\left(\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}}\right)\right) \le \frac{\gamma_{k_p}^{\ker}}{\sigma_{\text{eff}}}. \tag{H4}$$

Step 3: smoothing. By Lemma 29, the kernel has rank $M - M' = (k - 1 - |U_{\text{raw}}|)D \le (k-1)D$. The basis-dependent smoothing estimate of [40, Lemma 2.8], rescaled by $1/\sqrt{2\pi}$ to the normalization (2), together with (H4), gives

$$\eta_{\epsilon_0}\left(\boldsymbol{\Sigma}_0^{-1/2} \ker_{\mathbb{Z}}\left(\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}}\right)\right) \le \frac{\gamma_{k_p}^{\ker}}{\sigma_{\text{eff}}} \sqrt{\frac{\ln\left(2(k-1)D\left(1+1/\epsilon_0\right)\right)}{2\pi^2}}.$$

Since $1 + 1/\epsilon_0 = 2 + 1/\epsilon'$, the condition (50) implies (51). □

Proof of Lemma 31 If $t = 0$, then $\hat{\mathbf{G}}_U = \mathbf{M}_U$ is invertible and the kernel is trivial. Suppose $t \ge 1$, and write $\mathbf{b}'_a = \mathbf{P}_U^{-1}\mathbf{b}_a$. Since $\mathbf{M}_U$ is unimodular, $\mathbf{M}_U^{-1}$ is integral, so $\mathbf{b}'_a \in \mathbb{Z}^r$, and $\begin{bmatrix}\mathbf{M}_U & \mathbf{N}_U\end{bmatrix}\mathbf{b}'_a = \mathbf{0}$. Conversely, if $\begin{bmatrix}\mathbf{M}_U & \mathbf{N}_U\end{bmatrix}\begin{bmatrix}\mathbf{y}\\ \mathbf{x}\end{bmatrix} = \mathbf{0}$ with $\mathbf{y} \in \mathbb{Z}^{|U|}$ and $\mathbf{x} \in \mathbb{Z}^t$, then $\mathbf{y} = -\mathbf{M}_U^{-1}\mathbf{N}_U\mathbf{x}$, and hence $\begin{bmatrix}\mathbf{y}\\ \mathbf{x}\end{bmatrix} = \begin{bmatrix}-\mathbf{M}_U^{-1}\mathbf{N}_U\mathbf{x}\\ \mathbf{x}\end{bmatrix} = \sum_{a=1}^t \mathbf{x}[a]\begin{bmatrix}-\mathbf{M}_U^{-1}\mathbf{N}_U\mathbf{e}_a\\ \mathbf{e}_a\end{bmatrix} = \sum_{a=1}^t \mathbf{x}[a]\mathbf{b}'_a$. The vectors $\mathbf{b}'_a$ are linearly independent because their last t coordinates are $\mathbf{e}_1, \dots, \mathbf{e}_t$, so they form a basis of the permuted kernel; since $\mathbf{P}_U$ is unimodular, the $\mathbf{b}_a$ form a basis of $\ker_{\mathbb{Z}}(\hat{\mathbf{G}}_U)$. Since $\mathbf{P}_U$ is an isometry,

$$\|\mathbf{b}_a\|_2^2 = 1 + \left\|\mathbf{M}_U^{-1}\mathbf{N}_U\mathbf{e}_a\right\|_2^2 \le 1 + \frac{\|\mathbf{N}_U\mathbf{e}_a\|_2^2}{\sigma_{\min}^2(\mathbf{M}_U)} \le 1 + \frac{\min\{k_p, \beta_{k_p}^2\}}{\gamma_{k_p}^2},$$

because $\mathbf{N}_U\mathbf{e}_a$ is a column of $\hat{\mathbf{G}}_U \in \{-1,0,1\}^{|U|\times r}$, whose squared norm is at most $\min\{|U|, \|\hat{\mathbf{G}}_U\|_2^2\}$. Each Gram-Schmidt vector is no longer than the corresponding basis vector, and the right-hand side is at least 1, which proves (55). □

Proof of Lemma 32 By (45), $\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}} = \widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}}\mathbf{r}_0$ with $\mathbf{r}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^M,\sqrt{\boldsymbol{\Sigma}_0}}$.
Trivial kernel. If $\ker_{\mathbb{Z}}(\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}}) = \{\mathbf{0}\}$, then Lemma 29 gives $M = M'$ and $\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}} \in GL_M(\mathbb{Z})$, so the pushforward is exactly $\mathcal{D}_{\mathbb{Z}^{M'},\sqrt{\boldsymbol{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}}}$ by (47).

Non-trivial kernel. Apply [40, Theorem 3.1] with source lattice $\mathbb{Z}^M$, linear map $\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}}$, and parameter $\mathbf{S} := \sqrt{2\pi}\,\boldsymbol{\Sigma}_0^{1/2}$. Under the convention of that theorem, the source mass at $\mathbf{x}$ is proportional to $\exp(-\pi\|\mathbf{S}^{-1}\mathbf{x}\|_2^2) = \exp(-\frac{1}{2}\mathbf{x}^\top\boldsymbol{\Sigma}_0^{-1}\mathbf{x})$, so the source distribution is $\mathcal{D}_{\mathbb{Z}^M,\sqrt{\boldsymbol{\Sigma}_0}}$ in our convention. Since $\mathbf{S}^{-1} = (2\pi)^{-1/2}\boldsymbol{\Sigma}_0^{-1/2}$, its smoothing hypothesis $\eta_{\epsilon_0}^{\pi}(\mathbf{S}^{-1}\ker_{\mathbb{Z}}(\widetilde{\boldsymbol{\Phi}}_{\mathcal{U},\mathbf{c}})) \le 1$ is exactly (51), which holds by Lemma 30; the image lattice is $\mathbb{Z}^{M'}$ by Lemma 29. The theorem gives $\Delta(\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}, \mathcal{D}_{\mathbb{Z}^{M'},\sqrt{\boldsymbol{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}}}) \le \epsilon_0/(1-\epsilon_0) = \epsilon'$.

Transfer. The map $\mathbf{x} \mapsto \mathbf{W}^{-1}\mathbf{x}$ is a bijection of $\mathbb{Z}^{M'}$ and preserves statistical distance. It maps $\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}$ to $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ and, by (47), $\mathcal{D}_{\mathbb{Z}^{M'},\sqrt{\boldsymbol{\Sigma}_{\widetilde{\mathbf{y}}_{\mathcal{U},\mathbf{c}}}}}$ to $\mathcal{D}_{\mathbb{Z}^{M'},\sqrt{\boldsymbol{\Sigma}_{\mathbf{y}_{\mathcal{U},\mathbf{c}}}}}$. All quantities in (50) are independent of $(\mathcal{U},\mathbf{c})$, so the bound is uniform. □

H.3 Reductions

Proof of Lemma 33 For $b \in \{0,1\}$, let H_b be branch b of the fixed-pattern $\mathsf{HC\text{-}Leak}$ experiment, whose auxiliary leakage is $\mathbf{z}_{\mathcal{U},\mathbf{c}} = \mathbf{M}_{\mathcal{U},\mathbf{c}}\mathbf{s} + \mathbf{y}_{\mathcal{U},\mathbf{c}}$, and let $\breve{H}_b$ be obtained by replacing $\mathbf{y}_{\mathcal{U},\mathbf{c}}$ with $\breve{\mathbf{y}}_{\mathcal{U},\mathbf{c}}$. Both perturbations are independent of the secret and all other variables are generated identically, so applying the deterministic map $(\mathbf{s},\mathbf{y}) \mapsto \mathbf{M}_{\mathcal{U},\mathbf{c}}\mathbf{s} + \mathbf{y}$ and Lemma 32 gives $\Delta(H_b, \breve{H}_b) \le \epsilon'$. The Gaussianized branches are the one-hint Matrix Hint-MLWE distributions $(\boldsymbol{A}', \boldsymbol{A}\boldsymbol{s}, \mathbf{M}_{\mathcal{U},\mathbf{c}}, \breve{\mathbf{z}}_{\mathcal{U},\mathbf{c}})$ and $(\boldsymbol{A}', \boldsymbol{u}, \mathbf{M}_{\mathcal{U},\mathbf{c}}, \breve{\mathbf{z}}_{\mathcal{U},\mathbf{c}})$ with $\boldsymbol{u} \xleftarrow{\$} \mathcal{R}_q^p$, where the integer hint dimension M' corresponds to module dimension $(|U_{\mathrm{raw}}| + |\mathcal{S}|)\ell$. On a challenge $(\boldsymbol{A}', \boldsymbol{v}, \mathbf{M}_{\mathcal{U},\mathbf{c}}, \breve{\mathbf{z}}_{\mathcal{U},\mathbf{c}})$, the adversary $\mathcal{B}$ parses $\breve{\mathbf{z}}_{\mathcal{U},\mathbf{c}}$ according to (14), reshaping the first $|U_{\mathrm{raw}}|D$ coordinates into the raw-share block and the rest into the response blocks indexed by $\mathcal{S}$; since the ordering and dimensions are public, this recovers the structured Gaussianized view exactly. It gives this view, the public realization, and $\boldsymbol{v}$ to $\mathcal{A}$ and returns its output. The simulation is distributed as $\breve{H}_1$ when $\boldsymbol{v} = \boldsymbol{A}\boldsymbol{s} \bmod q$ and as $\breve{H}_0$ when $\boldsymbol{v}$ is uniform, so $|\Pr[\mathcal{A}(\breve{H}_1) = 1] - \Pr[\mathcal{A}(\breve{H}_0) = 1]| = \mathsf{Adv}_{\mathcal{B}}^{\mathsf{MHMLWE}}(\kappa)$. The triangle inequality over the two branches completes the proof. □

Proof of Lemma 34 By Lemma 33, there is $\mathcal{B}_1$ with $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U},\mathbf{c})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}^{\mathsf{MHMLWE}}_{\mathcal{B}_1}(\kappa) + 2\epsilon'$. The lower bound $\sqrt{2}\eta_\epsilon(\mathbb{Z}^D) \le \sigma_{\mathrm{MQ}}$ is the posterior-resampling condition required by Corollary 1. The upper bound $\sigma_{\mathrm{MQ}} \le \sigma_{\mathcal{U},\mathbf{c}}$ follows by comparing (57) with (58), since Lemma 28 bounds the precision term in (57) by Ξ_{MQ}. Hence $\sqrt{2}\eta_\epsilon(\mathbb{Z}^D) \le \sigma_{\mathrm{MQ}} \le \sigma_{\mathcal{U},\mathbf{c}}$. Applying Corollary 1 with $n = p$, $m = \ell - p$, $d(n+m) = D$, $\sigma' = \sigma_{\mathrm{MQ}}$, and $\sigma = \sigma_{\mathcal{U},\mathbf{c}}$ gives $\mathcal{B}$ with $\mathsf{Adv}^{\mathsf{MHMLWE}}_{\mathcal{B}_1}(\kappa) \le \mathsf{Adv}_{\mathcal{B},\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathrm{MQ}}}}(\kappa) + 2\epsilon$. Composing the last inequality with the result from Lemma 33 concludes the lemma. □

Proof of Theorem 4 Fix κ, a setup output $\mathsf{pp} \in \mathsf{Good}_\kappa$, a PPT adversary $\mathcal{A}$, and a pattern $\mathcal{U} \in \mathcal{U}_{k_p,Q}$. Since $\mathsf{Good}_{\mathrm{MQ}}(\mathsf{pp})$ holds, the conclusions of Section 5.4 hold for all admissible patterns, so the hypotheses of Lemma 34 hold for every $\mathbf{c} \in \mathcal{C}^Q$ whenever the auxiliary view is nonempty.

Empty auxiliary view. If $U_{\mathrm{raw}} = \emptyset$ and $\mathcal{S} = \emptyset$, then, apart from public information generated identically in both branches, the branches are $(\boldsymbol{A}', \boldsymbol{A}\boldsymbol{s})$ and $(\boldsymbol{A}', \boldsymbol{u})$, the real and random distributions of $\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathbf{s}}}$. In this case, $\mathcal{B}_{\mathsf{pp}}$ samples $\mathbf{c} \xleftarrow{\$} \mathcal{C}^Q$ and the remaining public information itself, so $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}_{\mathcal{B}_{\mathsf{pp}},\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathbf{s}}}}(\kappa)$.

Nonempty auxiliary view. On its MLWE challenge, $\mathcal{B}_{\mathsf{pp}}$ samples $\mathbf{c} \xleftarrow{\$} \mathcal{C}^Q$ exactly as in the experiment and runs the reduction of Lemma 34 for $(\mathcal{U},\mathbf{c})$. This is well-defined for every sampled $\mathbf{c}$ because $\|\mathbf{L}_s\|_2 \le \omega_c$ and the smoothing and posterior-resampling conditions hold by assumption, and the width σ_{MQ} and losses are uniform over $\mathbf{c}$. Hence $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}_{\mathcal{B}_{\mathsf{pp}},\mathsf{MLWE}_{p,\ell-p,q,d,\sigma_{\mathrm{MQ}}}}(\kappa) + 2\epsilon' + 2\epsilon$.

Conclusion. In both cases, $\mathsf{Adv}^{(\mathsf{pp},\mathcal{U})}_{\mathcal{A},\mathsf{HC\text{-}Leak}}(\kappa) \le \mathsf{Adv}_{\mathcal{B}_{\mathsf{pp}},\mathsf{MLWE}}(\kappa) + 2\epsilon' + 2\epsilon \le \epsilon_{\mathrm{MQ}}(\kappa,\mathcal{A},\mathcal{U})$, where the last inequality holds because ϵ_{MQ} is the maximum over Good_κ. This holds for every $\mathsf{pp} \in \mathsf{Good}_\kappa$, every PPT adversary $\mathcal{A}$, and every $\mathcal{U} \in \mathcal{U}_{k_p,Q}$, and Good_κ depends neither on $\mathcal{A}$ nor on $\mathcal{U}$. Therefore, the event in Theorem 4 contains Good_κ, and its probability is at least $\Pr_{\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)}[\mathsf{pp} \in \mathsf{Good}_\kappa] = 1 - \delta_{\mathrm{MQ}}$ by (60). Finally, sampling $\boldsymbol{A}'$ uniformly and setting $\boldsymbol{A} := \begin{bmatrix}\boldsymbol{I}_p & \boldsymbol{A}'\end{bmatrix}$ induces the stated distribution on $\mathcal{F}_{\mathrm{MQ}}$, and the real and random branches above are exactly those of Definition 9 for $f_{\boldsymbol{A}}$. The preceding probability and advantage

bounds therefore establish the $(\delta_{\mathrm{MQ}}, \epsilon_{\mathrm{MQ}})$-privacy required by Definition 10. Under the third assumption of the theorem, $\epsilon_{\mathrm{MQ}}(\kappa, \mathcal{A}, \mathcal{U})$ is negligible in κ for every PPT adversary $\mathcal{A}$ and every $\mathcal{U}$, which gives δ_{MQ}-privacy. □

Appendix I Gaussian Tail Bounds and Size Estimates

This appendix proves the tail bounds of Section 6.4.1 and derives the size estimates reported in Table 7.

I.1 Discrete Gaussian Tail Bound

Proof of Lemma 35 The bound (63) is [61, Lemma 4.4], whose Gaussian function $\exp(-\|\mathbf{x}\|_2^2/(2\sigma^2))$ agrees with (1) for $\mathbf{\Sigma} = \sigma^2\mathbf{I}_N$. Write $\theta = 1 + x$ with $x > 0$. Then $-\ln C_{\mathrm{tail}}(\theta) = x + \frac{x^2}{2} - \ln(1+x) \geq \frac{x^2}{2} > 0$, since $\ln(1+x) \leq x$; in particular, $C_{\mathrm{tail}}(\theta) < 1$. If, moreover, $x \geq \sqrt{(2\kappa_{\mathrm{tail}} \ln 2)/N}$, then $-\ln C_{\mathrm{tail}}(\theta) \geq \frac{x^2}{2} \geq \frac{\kappa_{\mathrm{tail}} \ln 2}{N}$, and hence $C_{\mathrm{tail}}(\theta)^N \leq 2^{-\kappa_{\mathrm{tail}}}$. □

I.2 Failure Probability of the Good-Tail Event

We apply Lemma 35 to the three families of Gaussian samples in Definition 13.

Secret

Under the coefficient embedding, $\boldsymbol{s} \leftarrow \mathcal{D}_S = \mathcal{D}_{\sigma_{\boldsymbol{s}}}^D$ is a sample of $\mathcal{D}_{\mathbb{Z}^D, \sigma_{\boldsymbol{s}}}$. Applying Lemma 35 with $N = D$ gives $\Pr[\|\boldsymbol{s}\|_2 > \theta_{\boldsymbol{s}} \sigma_{\boldsymbol{s}} \sqrt{D}] \leq C_{\mathrm{tail}}(\theta_{\boldsymbol{s}})^D \leq 2^{-\kappa}$ by (64). Since $\|\boldsymbol{s}\|_\infty \leq \|\boldsymbol{s}\|_2$ and $\beta_{\boldsymbol{s}} = \theta_{\boldsymbol{s}} \sigma_{\boldsymbol{s}} \sqrt{D}$, this proves (67).

Masking Matrix

Each entry of $\mathbf{R} \in \mathbb{Z}^{(k-1)\times D}$ is sampled independently from $\mathcal{D}_{\sigma_{\mathbf{R}}}$. For a fixed (i', j'), applying Lemma 35 with $N = 1$ gives $\Pr[|\mathbf{R}[i', j']| > \theta_{\mathbf{R}} \sigma_{\mathbf{R}}] \leq C_{\mathrm{tail}}(\theta_{\mathbf{R}}) \leq 2^{-\kappa_{\mathbf{R}}}$ by (65). A union bound over the $(k-1)D$ entries gives a failure probability of at most $(k-1)D \cdot 2^{-\kappa_{\mathbf{R}}} \leq 2^{-\kappa}$, since $\kappa_{\mathbf{R}} = \kappa + \lceil \log_2((k-1)D) \rceil$. This proves (68).

Leakage Masks

Assume $Q \geq 1$. For a fixed pair (i, s), the mask $\boldsymbol{r}_{i,s}$ is a sample of $\mathcal{D}_{\mathbb{Z}^D, \sigma_{\boldsymbol{r}}}$, and Lemma 35 with $N = D$ gives $\Pr[\|\boldsymbol{r}_{i,s}\|_2 > \theta_{\boldsymbol{r}} \sigma_{\boldsymbol{r}} \sqrt{D}] \leq C_{\mathrm{tail}}(\theta_{\boldsymbol{r}})^D \leq 2^{-\kappa_{\boldsymbol{r}}}$ by (66). A union bound over at most nQ pairs gives a failure probability of at most $nQ \cdot 2^{-\kappa_{\boldsymbol{r}}} \leq 2^{-\kappa}$, since $\kappa_{\boldsymbol{r}} = \kappa + \lceil \log_2(nQ) \rceil$. This proves (69). If $Q \geq 1$, a union bound over the secret, masking-matrix, and leakage-mask conditions gives $\Pr[\neg \mathsf{Good}_{\mathsf{tail}}] \leq 3 \cdot 2^{-\kappa}$. If $Q = 0$, the leakage-mask condition is vacuous, and a union bound over the remaining two conditions gives $\Pr[\neg \mathsf{Good}_{\mathsf{tail}}] \leq 2 \cdot 2^{-\kappa}$. Together, they prove Lemma 36.

I.3 Size Estimates

We report ideal entropy-coded sizes. Arithmetic coding approaches the entropy of a stationary memoryless source up to a constant coding overhead [66, Theorem 2]. As the

ideal source model for a coefficient of standard deviation σ, we use a variance-matched discrete Gaussian, whose entropy is approximately $\lg \sigma + \frac{1}{2}\lg(2\pi e) \approx \lg \sigma + 2.047$ bits [67, Section 5.1]; since the Gaussian maximizes entropy for a given variance, this is an upper estimate whenever the coefficients are independent.

Secret

The secret $\boldsymbol{s} \in \mathcal{R}^{\ell}$ has $D = d\ell$ independent coefficients sampled from $\mathcal{D}_{\sigma_{\boldsymbol{s}}}$, so its size is approximately $D(\lg \sigma_{\boldsymbol{s}} + 2.047)$ bits.

Share

By (7), coefficient j of $\boldsymbol{s}_i$ is $\mathbf{g}[i]\mathbf{s}[j] + \sum_{t=1}^{k-1} \hat{\mathbf{G}}[i,t]\mathbf{R}[t,j]$. For a fixed i, the coefficients are independent across j, because they involve disjoint coordinates of $\mathbf{s}$ and disjoint columns of $\mathbf{R}$. Every row of $\hat{\mathbf{G}}$ has exactly $c_{\exp}$ nonzero entries in $\{-1, 1\}$, so, approximating the variance of $\mathcal{D}_{\sigma}$ by σ^2, each coefficient has variance at most $\sigma_{\mathsf{sh}}^2 := \beta_{\mathbf{g}}^2\sigma_{\boldsymbol{s}}^2 + c_{\exp}\sigma_{\mathbf{R}}^2$. The size of one share is therefore approximately $D(\lg \sigma_{\mathsf{sh}} + 2.047)$ bits, and every participant stores exactly one such share. The values in Table 7 are obtained by dividing these bit estimates by $8 \cdot 2^{10}$. These estimates exclude implementation-specific coding, framing, finite-block, and byte-alignment overheads.

Appendix J Comparison with Other Short-Share Secret Sharing Schemes

This appendix compares the ideal entropy-coded secret size and long-term per-participant secret-share storage of SEC with those of the replicated and ramp sDKG constructions of [32] and the Vandermonde secret-sharing construction used by Hermine [19]. For SEC, we use the application-driven parameter sets with $Q = 2^{64}$, since both comparison schemes are instantiated for threshold-signature applications. We compare only secret and long-term secret-share sizes, excluding public auxiliary keys, verification material, signing communication, and other components specific to the surrounding threshold-signature schemes.

J.1 Secret Size

We use the ideal entropy model of Appendix I. For a secret $\boldsymbol{s} \in \mathcal{R}^{\ell}$ whose $D = d\ell$ coefficients are independently sampled from $\mathcal{D}_{\sigma_{\boldsymbol{s}}}$, its ideal entropy-coded size is approximately $D\,(\lg \sigma_{\boldsymbol{s}} + 2.047)$ bits. All schemes use the same standard-deviation convention, so their reported Gaussian widths can be compared directly.

J.1.1 Secret Size in sDKG

The threshold-signature parameters of [32, Table 2] comprise two parameter sets at $\kappa = 128$ and one at $\kappa = 256$. The first $\kappa = 128$ set has $(d, p, \ell) = (512, 4, 4)$ and secret width $\sigma_{\boldsymbol{s}} = 2^{10}$, and the $\kappa = 256$ set has $(d, p, \ell) = (512, 7, 7)$ and $\sigma_{\boldsymbol{s}} = 2^{10}$; both support up to 16 participants. The second $\kappa = 128$ set has $(d, p, \ell) = (512, 5, 5)$ and $\sigma_{\boldsymbol{s}} = 2^{8}$, and it is the only reported set supporting more than 16 participants, namely

up to 1024 participants for ramp sDKG. The paper does not report a parameter set for $\kappa = 192$.

Each sDKG sub-share is a pair $(\boldsymbol{s}_{\mathsf{sub}}, \boldsymbol{e}_{\mathsf{sub}}) \in \mathcal{R}^{\ell} \times \mathcal{R}^{p}$, so we measure the shared secret as the corresponding pair $(\boldsymbol{s}, \boldsymbol{e}) \in \mathcal{R}^{\ell} \times \mathcal{R}^{p}$. This is also the like-for-like comparison with SEC, whose secret contains both parts of the systematic MLWE instance. Under the common ideal entropy model mentioned earlier, the sDKG secret size is $|(\boldsymbol{s}, \boldsymbol{e})|_{\mathsf{sDKG}} \approx d(p + \ell)\,(\lg \sigma_{\boldsymbol{s}} + 2.047)$ bits. This gives approximately 6.024 and 6.279 KiB for the first and second $\kappa = 128$ sets, respectively, and 10.541 KiB for $\kappa = 256$. Table J1 compares these values with the application-driven SEC parameters for $Q = 2^{64}$.

Table J1: Ideal entropy-coded secret sizes, in KiB, in sDKG and SEC

	sDKG				SEC, $n = 40$		SEC, $n = 60$	
κ	(d, p, ℓ)	$\lg \sigma_{\boldsymbol{s}}$	$n_{\max}$	Secret	Secret	Ratio	Secret	Ratio
128	$(512, 4, 4)$	10	16	6.024	4.333	1.390	3.538	1.703
128	$(512, 5, 5)$	8	1024	6.279	4.333	1.449	3.538	1.775
256	$(512, 7, 7)$	10	16	10.541	7.880	1.338	6.486	1.625

Note: the ratios are $|(\boldsymbol{s}, \boldsymbol{e})|_{\mathsf{sDKG}} / |\boldsymbol{s}|_{\mathsf{SEC}}$. $n_{\max}$ is the largest supported number of participants; the value 1024 applies to ramp sDKG only. The SEC values use the application-driven parameter sets with $Q = 2^{64}$. The sDKG paper reports no parameter set for $\kappa = 192$. All sizes use the ideal entropy model of Appendix I.

J.1.2 Secret Size in Vandermonde Secret Sharing

We consider the Vandermonde secret-sharing construction used by Hermine [19]. For the threshold-signature parameter sets in [19, Table 3], the ring degree is $d = 512$, the secret-module ranks are $\ell = 4$, 6, and 8 for $\kappa = 128$, 192, and 256, respectively, and $u_{\mathbf{s}} := \lceil \lg \sigma_{\mathbf{s}} \rceil = 13$. Since $\lg \sigma_{\mathbf{s}} \le u_{\mathbf{s}}$, replacing the exact logarithmic width by $u_{\mathbf{s}}$ gives a conservative ideal entropy estimate.

Under the common ideal entropy model mentioned earlier, the secret size is bounded approximately by $|\boldsymbol{s}|_{\mathsf{Vand}} \lesssim d\ell\,(u_{\mathbf{s}} + 2.047) = 15.047 d\ell$ bits. This gives approximately 3.762, 5.643, and 7.524 KiB for $\kappa = 128$, 192, and 256, respectively. Table J2 compares these estimates with the application-driven SEC parameters for $Q = 2^{64}$.

Measured consistently, the sDKG secret is 1.34 to 1.78 times as large as the SEC secret, whereas the Vandermonde upper estimate is comparable to it: slightly smaller for $n = 40$ and slightly larger for $n = 60$. The more substantial storage difference arises from the number of sub-shares held by each participant, as analyzed in Appendix J.2.

Table J2: Ideal entropy-coded secret-size upper estimates, in KiB, in Vandermonde secret sharing and SEC

	Vandermonde secret sharing			SEC, $n = 40$		SEC, $n = 60$	
κ	(d, ℓ)	$u_{\mathbf{s}}$	Secret	Secret	Ratio	Secret	Ratio
128	$(512, 4)$	13	3.762	4.333	0.868	3.538	1.063
192	$(512, 6)$	13	5.643	6.087	0.927	4.991	1.131
256	$(512, 8)$	13	7.524	7.880	0.955	6.486	1.160

Note: the Vandermonde values conservatively replace $\lg \sigma_{\mathbf{s}}$ by $u_{\mathbf{s}} = \lceil \lg \sigma_{\mathbf{s}} \rceil = 13$. The ratios are $|\boldsymbol{s}|_{\mathsf{Vand}}/|\boldsymbol{s}|_{\mathsf{SEC}}$. The SEC values use the application-driven parameter sets with $Q = 2^{64}$. All sizes use the ideal entropy model of Appendix I.

J.2 Long-Term Secret-Share Size

Since SEC has distinct privacy and empirical reconstruction thresholds k_p and k_c, no perfect t-out-of-n access structure simultaneously matches both thresholds. We therefore compare SEC with two profiles: $t_{\mathsf{pm}} := k_p + 1$ and $t_{\mathsf{qm}} := k_c$. The privacy-matched profile t_{pm} tolerates the same number k_p of corrupted participants, whereas the quorum-matched profile t_{qm} matches the empirical reconstruction threshold of SEC. More generally, every $t \in [t_{\mathsf{pm}}, t_{\mathsf{qm}}]$ provides at least the stated privacy of SEC without requiring more participants than its empirical reconstruction threshold. The two endpoint profiles are used when counting the sub-shares stored in the replicated and ramp sDKG constructions and in the Vandermonde construction used by Hermine. We consider $(n, t_{\mathsf{pm}}, t_{\mathsf{qm}}) \in \{(40, 6, 26), (60, 7, 46)\}$ in our comparative analysis.

J.2.1 Secret-Share Size in sDKG

Each short sub-share in the replicated and ramp sDKG constructions of [32] is a pair $(\boldsymbol{s}_{\mathsf{sub}}, \boldsymbol{e}_{\mathsf{sub}}) \in \mathcal{R}^{\ell} \times \mathcal{R}^{p}$. For a coefficient width σ, we estimate its ideal entropy-coded size by $\mathsf{ShareSize}_{p,\ell}(\sigma) := d(p+\ell)\,(\lg \sigma + 2.047)$ bits. The two constructions assign different numbers of sub-shares to each participant and use different reported parameter sets. For replicated sDKG, we use the first $\kappa = 128$ set of Appendix J.1.1, which gives the smallest sub-shares among the reported sets; for ramp sDKG, we use the second $\kappa = 128$ set, the only one supporting more than 16 participants. We compare their per-participant storage with the application-driven SEC share sizes for $Q = 2^{64}$ in Table 7. The sDKG paper reports parameter sets for $\kappa = 128$ and 256, but not for $\kappa = 192$.

Replicated sDKG

In a t-out-of-n replicated sDKG, there are $\binom{n}{t-1}$ short sub-shares in total, and each participant stores $k_{\mathsf{Repl}}(n, t) := \binom{n-1}{t-1}$ of them. Hence, the estimated per-participant storage is $|\mathsf{sk}_i^{\mathsf{Repl}}| \approx \binom{n-1}{t-1}\mathsf{ShareSize}_{p,\ell}(\sigma_{\boldsymbol{s}})$. At $\kappa = 128$, the reported replicated-sDKG parameters are $(d, p, \ell) = (512, 4, 4)$ and $\sigma_{\boldsymbol{s}} = 2^{10}$, giving $\mathsf{ShareSize}_{4,4}(2^{10}) = 512 \cdot 8 \cdot (10 + 2.047) \approx 6.024$ KiB. At $\kappa = 256$, the corresponding parameters are $(d, p, \ell) =$

$(512, 7, 7)$ with the same Gaussian width, so the individual sub-share size increases by the factor 7/4 to approximately 10.541 KiB.

Table J3 reports the resulting per-participant long-term secret-share storage for the aforementioned privacy- and quorum-matched profiles. The reported replicated-sDKG parameters support up to 16 participants, so the $n = 40$ and $n = 60$ values are extrapolations of the replicated construction.

Table J3: Estimated per-participant long-term secret-share storage for replicated sDKG, compared with SEC

κ	n	t	Profile	# sub-shares	sDKG storage	SEC share	Ratio
128	40	6	Privacy	575,757	3.307 GiB	6.068 KiB	571,534.655
128	40	26	Quorum	15,084,504,396	86,652.291 GiB	6.068 KiB	14,973,881,382.549
128	60	7	Privacy	45,057,474	258.831 GiB	6.416 KiB	42,301,074.601
128	60	46	Quorum	13,298,522,298,180	76,392,792.762 GiB	6.416 KiB	12,484,982,709,334.045
256	40	6	Privacy	575,757	5.788 GiB	10.674 KiB	568,589.705
256	40	26	Quorum	15,084,504,396	151,641.508 GiB	10.674 KiB	14,896,725,351.441
256	60	7	Privacy	45,057,474	452.954 GiB	11.252 KiB	42,210,848.349
256	60	46	Quorum	13,298,522,298,180	133,687,387.333 GiB	11.252 KiB	12,458,352,813,757.791

Note: the ratios are $|\mathsf{sk}_i^{\mathsf{Repl}}|/|\boldsymbol{s}_i^{\mathsf{SEC}}|$. The SEC values use the application-driven parameter sets with $Q = 2^{64}$. The sDKG values measure the ideal entropy-coded storage of all secret sub-shares held by one participant and exclude indexing, framing, implementation overheads, and the seed-and-tree compression proposed in [32]. The $n = 40$ and $n = 60$ profiles are extrapolated beyond the reported participant range of replicated sDKG.

The number of replicated sub-shares grows combinatorially with the matched threshold. Consequently, even the privacy-matched profiles require per-participant storage more than $5 \cdot 10^5$ times that of one SEC share, while the quorum-matched profiles are substantially larger.

Ramp sDKG

In the basic ramp sDKG, let ν be the number of additive share buckets and μ the number of parallel sharings. Each party stores one short sub-share from every parallel sharing, and hence stores μ sub-shares in total. Following the correctness-amplified construction of [32], we set $\mu := \kappa$, giving the estimated per-participant storage $|\mathsf{sk}_i^{\mathsf{Ramp}}| \approx \kappa \cdot \mathsf{ShareSize}_{p,\ell}(\sigma_{\mathsf{share}})$. At $\kappa = 128$, we use the reported ramp-sDKG parameter set supporting the larger participant range, where $(d, p, \ell) = (512, 5, 5)$ and $\sigma_{\mathsf{share}} = 2^{10}$. Thus, one ordinary sub-share has ideal entropy-coded size $\mathsf{ShareSize}_{5,5}(2^{10}) = 512 \cdot 10\left(10 + \frac{1}{2}\lg(2\pi e)\right) \approx 7.529$ KiB, and the total per-participant storage is approximately $128 \cdot 7.529 = 963.768$ KiB. At $\kappa = 256$, the corresponding parameters are $(d, p, \ell) = (512, 7, 7)$ with the same Gaussian width. The individual sub-share size increases by a factor of 7/5, while the number of parallel sharings doubles from 128 to 256. The total storage therefore increases by a factor of $256/128 \cdot 7/5 = 14/5$ to approximately 2698.549 KiB.

To obtain an sDKG-favorable estimate, we assume that every sub-share has width σ_{share}. This omits the larger dependent bucket in each additive sharing and therefore gives an optimistic estimate for the ramp construction. Table J4 compares these values with the application-driven SEC share sizes for $Q = 2^{64}$.

Table J4: Optimistic ideal entropy-coded per-participant long-term secret-share storage, in KiB, for ramp sDKG, compared with SEC

κ	n	μ	One sub-share	sDKG storage	SEC share	Ratio
128	40	128	7.529	963.768	6.068	158.828
128	60	128	7.529	963.768	6.416	150.213
256	40	256	10.541	2698.549	10.674	252.815
256	60	256	10.541	2698.549	11.252	239.828

Note: the ratios are $|\mathsf{sk}_i^{\mathsf{Ramp}}|/|\mathsf{s}_i^{\mathsf{SEC}}|$. The SEC values use the application-driven parameter sets with $Q = 2^{64}$. The ramp-sDKG estimates charge every sub-share at the ordinary width σ_{share} and omit the larger dependent bucket and implementation-specific overheads.

The ramp construction avoids the combinatorial sub-share growth of replicated sDKG, but the correctness-amplified instantiation stores $\mu = \kappa$ short sub-shares per participant. Even under the optimistic ordinary-sub-share estimate, its per-participant storage is approximately 150–159 times that of one SEC share at $\kappa = 128$ and 240–253 times at $\kappa = 256$.

J.2.2 Secret-Share Size in Vandermonde Secret Sharing

Number of Vandermonde Sub-Shares

The Vandermonde construction used by Hermine recursively partitions the participants, so an unbalanced partition tree may assign different numbers of sub-shares to different parties [19, 34]. Let $\mathbf{u}(n,t) := (u_1(n,t), \ldots, u_n(n,t)) \in \mathbb{Z}_{\geq 0}^n$, where $u_i(n,t)$ denotes the number of sub-shares assigned to party i in a t-out-of-n sharing. Define $L(n,t) := \min_{i\in[n]} u_i(n,t)$, $U(n,t) := \max_{i\in[n]} u_i(n,t)$, and $S(n,t) := \sum_{i=1}^n u_i(n,t)$ as the minimum number of sub-shares per party, the maximum number of sub-shares per party, and the total number of sub-shares across all parties, respectively.

Incorporating Hermine's $t = n$ optimization, set $\mathbf{u}(n,1) = \mathbf{u}(n,n) = \mathbf{1}_n$, and set $\mathbf{u}(n,t) = \mathbf{0}_n$ whenever $t \leq 0$ or $t > n$. For $t \in [2, n-1]$, let $n_{\mathsf{L}} := \lfloor n/2 \rfloor$, $n_{\mathsf{R}} := \lceil n/2 \rceil$, and define $\mathcal{T}_{\mathsf{L}}(n,t) := [\max\{1, t-n_{\mathsf{R}}\}, \min\{t, n_{\mathsf{L}}\}]$ and $\mathcal{T}_{\mathsf{R}}(n,t) := [\max\{1, t-n_{\mathsf{L}}\}, \min\{t, n_{\mathsf{R}}\}]$. The participant-wise count vector satisfies

$$\mathbf{u}(n,t) = \left(\sum_{t_{\mathsf{L}} \in \mathcal{T}_{\mathsf{L}}(n,t)} \mathbf{u}(n_{\mathsf{L}}, t_{\mathsf{L}}) \right) \| \left(\sum_{t_{\mathsf{R}} \in \mathcal{T}_{\mathsf{R}}(n,t)} \mathbf{u}(n_{\mathsf{R}}, t_{\mathsf{R}}) \right), \tag{J1}$$

where $\|$ denotes vector concatenation in the left-right order of the recursive partition.[1] Recurrence (J1) determines $L(n,t)$, $U(n,t)$, and $S(n,t)$ exactly. The minimum, average, and maximum per-party sub-share counts are $L(n,t)$, $S(n,t)/n$, and $U(n,t)$, respectively. We evaluate these quantities for the privacy- and quorum-matched profiles $(40,6)$, $(40,26)$, $(60,7)$, and $(60,46)$ in the subsequent storage comparison.

Size of an Individual Vandermonde Sub-Share

Each Vandermonde index idx corresponds to a pair $(\boldsymbol{s}_{\mathsf{idx}}, \boldsymbol{e}_{\mathsf{idx}}) \in \mathcal{R}^{\ell} \times \mathcal{R}^{p}$, but the long-term signing share stores only $\boldsymbol{s}_{\mathsf{idx}} \in \mathcal{R}^{\ell}$. The component $\boldsymbol{e}_{\mathsf{idx}}$ is used to construct the indexed partial verification key and is excluded consistently with the scope of this appendix. A party stores one such vector for each index assigned to it, and we measure the initial storage before any key refresh.

Following the convolution argument in the proof of [19, Lemma 6.1], each coefficient of an initial sub-share is modeled as a signed sum of at most $1 + \lceil \lg n \rceil$ independent Gaussian contributions with standard deviation $\sigma_{\mathbf{s}}$. We therefore use the uniform variance-matched width $\sigma_{\mathsf{sub}} := \sigma_{\mathbf{s}} \sqrt{1 + \lceil \lg n \rceil}$. Since $\boldsymbol{s}_{\mathsf{idx}}$ has $d\ell$ coefficients, its ideal entropy-coded size is approximately $|\boldsymbol{s}_{\mathsf{idx}}| \approx d\ell\big(\lg \sigma_{\mathbf{s}} + \frac{1}{2}\lg(1 + \lceil \lg n \rceil) + 2.047\big) \leq d\ell\big(u_{\mathbf{s}} + \frac{1}{2}\lg(1 + \lceil \lg n \rceil) + 2.047\big)$ bits, where $u_{\mathbf{s}} := \lceil \lg \sigma_{\mathbf{s}} \rceil = 13$ for the parameter sets considered. This uniform estimate may overestimate sub-shares arising from shorter or degenerate recursive paths, while replacing $\lg \sigma_{\mathbf{s}}$ by $u_{\mathbf{s}}$ adds less than one bit per coefficient. Measuring the initial sharing before any refresh is favorable to Hermine.

Secret-Share Storage Comparison

Since [19] does not provide concrete parameter sets for $n \in \{40, 60\}$, we extrapolate the Vandermonde sub-share counts and sizes using the parameter family with $\lceil \lg q \rceil = 48$ in [19, Table 3], which supports the larger reported signing-query budgets. For $\kappa \in \{128, 192, 256\}$, this family uses $(p, \ell) \in \{(4,4), (6,6), (7,8)\}$, respectively, with $d = 512$ and $u_{\mathbf{s}} = 13$. Since $\lceil \lg 40 \rceil = \lceil \lg 60 \rceil = 6$, the uniform estimate of the previous discussion gives $|\boldsymbol{s}_{\mathsf{idx}}| \leq 512\ell\left(13 + \frac{1}{2}\lg 7 + 2.047\right)$ bits. These values extrapolate Hermine beyond its analyzed range $n \leq 16$ and $t \leq 8$, and hence quantify the storage induced by its Vandermonde construction without constituting security-validated Hermine parameter sets for $n \in \{40, 60\}$. For a t-out-of-n profile, the minimum, average, and maximum per-party long-term secret-share storage are $L(n,t)|\boldsymbol{s}_{\mathsf{idx}}|$, $S(n,t)|\boldsymbol{s}_{\mathsf{idx}}|/n$, and $U(n,t)|\boldsymbol{s}_{\mathsf{idx}}|$, respectively. Table J5 reports these values at $\kappa = 128$ for the privacy- and quorum-matched profiles mentioned earlier, together with the ratio between the minimum Vandermonde storage and the application-driven SEC share size for $Q = 2^{64}$. Neither profile is strictly like-for-like: privacy matching allows reconstruction by fewer parties than the empirical SEC threshold, whereas quorum matching gives the Vandermonde scheme a stronger privacy threshold than SEC.

Even for the party with the smallest estimated storage, the extrapolated Vandermonde construction requires approximately 58–692 times the storage of one SEC share

[1] For odd n, the upper endpoint of the right sub-threshold range is $\min\{t, \lceil n/2 \rceil\}$, because the right subtree contains $\lceil n/2 \rceil$ participants. This corrects the expression $\min\{t, \lfloor n/2 \rfloor\}$ in [34, p. 15]; the two expressions coincide when n is even.

Table J5: Estimated initial per-party long-term secret-share storage, in KiB, in Vandermonde secret sharing at $\kappa = 128$, compared with SEC

n	t	Profile	# sub-shares	Vandermonde storage			SEC share	Ratio
			Min./Avg./Max.	Min.	Avg.	Max.		Min. Vand./SEC
40	6	Privacy	86/104.4/127	353.692	429.365	522.312	6.068	58.288
40	26	Quorum	474/581.6/713	1949.417	2391.942	2932.350	6.068	321.262
60	7	Privacy	170/261.6/280	699.158	1075.881	1151.554	6.416	108.971
60	46	Quorum	1080/1748.8/1920	4441.709	7192.278	7896.371	6.416	692.286

Note: the sub-share-count column reports $L(n,t)$, $S(n,t)/n$, and $U(n,t)$, respectively. The ratio is the minimum Vandermonde per-party storage divided by the application-driven SEC share size with $Q = 2^{64}$. The Vandermonde values are extrapolated beyond Hermine's analyzed range $n \leq 16$ and $t \leq 8$ and measure the initial sharing before any refresh.

at $\kappa = 128$, with the larger ratios arising under quorum matching. This difference is because Vandermonde sharing stores many indexed sub-shares per party, whereas SEC stores one share per party.

References

[1] Shamir, A.: How to share a secret. Communications of the ACM **22**(11), 612–613 (1979) https://doi.org/10.1145/359168.359176

[2] Blakley, G.R.: Safeguarding cryptographic keys. In: Proceedings of the National Computer Conference (AFIPS 1979). AFIPS Conference Proceedings, vol. 48, pp. 313–317 (1979). https://doi.org/10.1109/MARK.1979.8817296

[3] Blakley, G.R., Meadows, C.: Security of ramp schemes. In: Blakley, G.R., Chaum, D. (eds.) Advances in Cryptology – CRYPTO '84. Lecture Notes in Computer Science, vol. 196, pp. 242–268. Springer, Berlin, Heidelberg (1985). https://doi.org/10.1007/3-540-39568-7_20

[4] Yamamoto, H.: Secret sharing system using (k, L, n) threshold scheme. Electronics and Communications in Japan (Part I: Communications) **69**(9), 46–54 (1986) https://doi.org/10.1002/ecja.4410690906

[5] McEliece, R.J., Sarwate, D.V.: On sharing secrets and Reed–Solomon codes. Communications of the ACM **24**(9), 583–584 (1981) https://doi.org/10.1145/358746.358762

[6] Massey, J.L.: Minimal codewords and secret sharing. In: Proceedings of the 6th Joint Swedish–Russian International Workshop on Information Theory, Mölle, Sweden, pp. 276–279 (1993)

[7] Chen, H., Cramer, R., Goldwasser, S., Haan, R., Vaikuntanathan, V.: Secure

computation from random error correcting codes. In: Naor, M. (ed.) Advances in Cryptology – EUROCRYPT 2007. Lecture Notes in Computer Science, vol. 4515, pp. 291–310. Springer, Berlin, Heidelberg (2007). https://doi.org/10.1007/978-3-540-72540-4_17

[8] Cramer, R., Damgård, I.B., Döttling, N., Fehr, S., Spini, G.: Linear secret sharing schemes from error correcting codes and universal hash functions. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology – EUROCRYPT 2015, Part II. Lecture Notes in Computer Science, vol. 9057, pp. 313–336. Springer, Berlin, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_11

[9] Fehr, S.: Span programs over rings and how to share a secret from a module. Master's thesis, ETH Zurich (1998)

[10] Desmedt, Y., Frankel, Y.: Threshold cryptosystems. In: Brassard, G. (ed.) Advances in Cryptology – CRYPTO '89. Lecture Notes in Computer Science, vol. 435, pp. 307–315. Springer, New York, NY (1990). https://doi.org/10.1007/0-387-34805-0_28

[11] Cramer, R., Fehr, S.: Optimal black-box secret sharing over arbitrary Abelian groups. In: Yung, M. (ed.) Advances in Cryptology – CRYPTO 2002. Lecture Notes in Computer Science, vol. 2442, pp. 272–287. Springer, Berlin, Heidelberg (2002). https://doi.org/10.1007/3-540-45708-9_18

[12] Cramer, R., Fehr, S., Stam, M.: Black-box secret sharing from primitive sets in algebraic number fields. In: Shoup, V. (ed.) Advances in Cryptology – CRYPTO 2005. Lecture Notes in Computer Science, vol. 3621, pp. 344–360. Springer, Berlin, Heidelberg (2005). https://doi.org/10.1007/11535218_21

[13] Cramer, R., Xing, C.: Blackbox secret sharing revisited: A coding-theoretic approach with application to expansionless near-threshold schemes. In: Canteaut, A., Ishai, Y. (eds.) Advances in Cryptology – EUROCRYPT 2020, Part I. Lecture Notes in Computer Science, vol. 12105, pp. 499–528. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_18

[14] Applebaum, B., Nir, O., Pinkas, B.: How to recover a secret with $O(n)$ additions. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023, Part I. Lecture Notes in Computer Science, vol. 14081, pp. 236–262. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38557-5_8

[15] Brandão, L.T.A.N., Peralta, R.: NIST first call for multi-party threshold schemes. NIST Internal Report 8214C, National Institute of Standards and Technology (2026). https://doi.org/10.6028/NIST.IR.8214C

[16] Komlo, C., Goldberg, I.: FROST: Flexible round-optimized Schnorr threshold signatures. In: Dunkelman, O., Jacobson, M.J. Jr., O'Flynn, C. (eds.) Selected Areas in Cryptography – SAC 2020. Lecture Notes in Computer Science, vol. 12804, pp.

34–65. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-81652-0_2

[17] Pino, R., Katsumata, S., Prest, T., Rossi, M.: Raccoon: A masking-friendly signature proven in the probing model. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024, Part I. Lecture Notes in Computer Science, vol. 14920, pp. 409–444. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-68376-3_13

[18] Pino, R., Katsumata, S., Maller, M., Mouhartem, F., Prest, T., Saarinen, M.-J.O.: Threshold Raccoon: Practical threshold signatures from standard lattice assumptions. In: Joye, M., Leander, G. (eds.) Advances in Cryptology – EUROCRYPT 2024, Part II. Lecture Notes in Computer Science, vol. 14652, pp. 219–248. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-58723-8_8

[19] Borin, G., Celi, S., Pino, R., Espitau, T., Katsumata, S., Niot, G., Prest, T., Takemure, K.: Hermine: An efficient lattice-based FROST-like threshold signature. In: Advances in Cryptology – ASIACRYPT 2026 (2026). To appear. Full version: Cryptology ePrint Archive, Paper 2026/419. https://eprint.iacr.org/2026/419

[20] Katsumata, S., Reichle, M., Takemure, K.: Adaptively secure 5 round threshold signatures from MLWE/MSIS and DL with rewinding. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024, Part VII. Lecture Notes in Computer Science, vol. 14926, pp. 459–491. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-68394-7_15

[21] Espitau, T., Katsumata, S., Takemure, K.: Two-round threshold signature from algebraic one-more learning with errors. Journal of Cryptology **38**(4), 31 (2025) https://doi.org/10.1007/s00145-025-09549-2

[22] Zhu, C., Tessaro, S.: The algebraic one-more MISIS problem and applications to threshold signatures. In: Kalai, Y.T., Kamara, S.F. (eds.) Advances in Cryptology – CRYPTO 2025, Part I. Lecture Notes in Computer Science, vol. 16000, pp. 548–581. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-01855-7_18

[23] Boschini, C., Kaviani, D., Lai, R.W.F., Malavolta, G., Takahashi, A., Tibouchi, M.: Ringtail: Practical two-round threshold signatures from learning with errors. In: 2025 IEEE Symposium on Security and Privacy (SP), pp. 149–164. IEEE, Piscataway, NJ (2025). https://doi.org/10.1109/SP61157.2025.00070

[24] Espitau, T., Niot, G., Prest, T.: Flood and submerse: Distributed key generation and robust threshold signature from lattices. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024, Part VII. Lecture Notes in Computer Science, vol. 14926, pp. 425–458. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-68394-7_14

[25] Pino, R., Katsumata, S., Niot, G., Reichle, M., Takemure, K.: Unmasking TRaccoon: A lattice-based threshold signature with an efficient identifiable abort

protocol. In: Kalai, Y.T., Kamara, S.F. (eds.) Advances in Cryptology – CRYPTO 2025, Part VI. Lecture Notes in Computer Science, vol. 16005, pp. 423–456. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-01887-8_14

[26] Gur, K.D., Hough, P., Katz, J., Sandsbråten, C., Silde, T.: Olingo: Threshold lattice signatures with DKG and identifiable abort. Cryptology ePrint Archive, Paper 2025/1789. To appear at ACM CCS 2026 (2025). https://eprint.iacr.org/2025/1789

[27] Ball, M., Çakan, A., Malkin, T.: Linear threshold secret-sharing with binary reconstruction. In: Tessaro, S. (ed.) 2nd Conference on Information-Theoretic Cryptography (ITC 2021). Leibniz International Proceedings in Informatics (LIPIcs), vol. 199, pp. 12–11222. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2021). https://doi.org/10.4230/LIPIcs.ITC.2021.12

[28] Boneh, D., Gennaro, R., Goldfeder, S., Jain, A., Kim, S., Rasmussen, P.M.R., Sahai, A.: Threshold cryptosystems from threshold fully homomorphic encryption. In: Shacham, H., Boldyreva, A. (eds.) Advances in Cryptology – CRYPTO 2018, Part I. Lecture Notes in Computer Science, vol. 10991, pp. 565–596. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96884-1_19

[29] Benaloh, J.C., Leichter, J.: Generalized secret sharing and monotone functions. In: Goldwasser, S. (ed.) Advances in Cryptology – CRYPTO '88. Lecture Notes in Computer Science, vol. 403, pp. 27–35. Springer, New York, NY (1990). https://doi.org/10.1007/0-387-34799-2_3

[30] Boppana, R.B.: Amplification of probabilistic Boolean formulas. In: 26th Annual Symposium on Foundations of Computer Science (FOCS 1985), pp. 20–29. IEEE Computer Society, Los Alamitos, CA (1985). https://doi.org/10.1109/SFCS.1985.5

[31] Chairattana-Apirom, R., Tessaro, S., Zhu, C.: Partially non-interactive two-round lattice-based threshold signatures. In: Chung, K.-M., Sasaki, Y. (eds.) Advances in Cryptology – ASIACRYPT 2024, Part IV. Lecture Notes in Computer Science, vol. 15487, pp. 268–302. Springer, Singapore (2024). https://doi.org/10.1007/978-981-96-0894-2_9

[32] Pino, R., Espitau, T., Niot, G., Prest, T.: How to shortly share a short vector: DKG with short shares and application to lattice-based threshold signatures with identifiable aborts. Manuscript (2025). https://gniot.fr/assets/papers/sdkg.pdf

[33] Desmedt, Y., Di Crescenzo, G., Burmester, M.: Multiplicative non-abelian sharing schemes and their application to threshold cryptography. In: Pieprzyk, J., Safavi-Naini, R. (eds.) Advances in Cryptology – ASIACRYPT '94. Lecture Notes in Computer Science, vol. 917, pp. 19–32. Springer, Berlin, Heidelberg (1995). https://doi.org/10.1007/BFb0000421

[34] Battagliola, M., Borin, G., Di Crescenzo, G., Meneghetti, A., Persichetti, E.: Enhancing threshold group action signature schemes: Adaptive security and scalability improvements. In: Niederhagen, R., Saarinen, M.-J.O. (eds.) Post-Quantum Cryptography – PQCrypto 2025, Part I. Lecture Notes in Computer Science, vol. 15577, pp. 129–161. Springer, Cham (2025). https://doi.org/10.1007/978-3-031-86599-2_5

[35] Bormann, C., Ersue, M., Keranen, A.: Terminology for constrained-node networks. RFC 7228, RFC Editor (2014). https://doi.org/10.17487/RFC7228

[36] Kannwischer, M.J., Rijneveld, J., Schwabe, P., Stoffelen, K.: pqm4: Testing and benchmarking NIST PQC on ARM Cortex-M4. Cryptology ePrint Archive, Paper 2019/844 (2019). https://eprint.iacr.org/2019/844

[37] Bettale, L., De Oliveira, M., Dottax, E.: Post-quantum protocols for banking applications. In: Buhan, I., Schneider, T. (eds.) Smart Card Research and Advanced Applications – CARDIS 2022. Lecture Notes in Computer Science, vol. 13820, pp. 271–289. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-25319-5_14

[38] Puch, N., Pursche, M., Peters, S.N., Heinl, M.P.: SoK: The engineer's guide to post-quantum cryptography for embedded devices. In: 2026 IEEE 11th European Symposium on Security and Privacy (EuroS&P), pp. 23–42. IEEE, Piscataway, NJ (2026). https://doi.org/10.1109/EuroSP68448.2026.00014

[39] Nielsen, J.B., Simkin, M.: Lower bounds for leakage-resilient secret sharing. In: Canteaut, A., Ishai, Y. (eds.) Advances in Cryptology – EUROCRYPT 2020, Part I. Lecture Notes in Computer Science, vol. 12105, pp. 556–577. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_20

[40] Genise, N., Micciancio, D., Peikert, C., Walter, M.: Improved discrete Gaussian and subgaussian analysis for lattice cryptography. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) Public-Key Cryptography – PKC 2020, Part I. Lecture Notes in Computer Science, vol. 12110, pp. 623–651. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45374-9_21

[41] Kim, D., Lee, D., Seo, J., Song, Y.: Toward practical lattice-based proof of knowledge from Hint-MLWE. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023, Part V. Lecture Notes in Computer Science, vol. 14085, pp. 549–580. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38554-4_18

[42] Sipser, M., Spielman, D.A.: Expander codes. IEEE Transactions on Information Theory **42**(6), 1710–1722 (1996) https://doi.org/10.1109/18.556667

[43] Guruswami, V., Rudra, A., Sudan, M.: Essential Coding Theory. Book draft, available at https://cse.buffalo.edu/faculty/atri/courses/coding-theory/

book/ (2023)

[44] Hoory, S., Linial, N., Wigderson, A.: Expander graphs and their applications. Bulletin of the American Mathematical Society **43**(4), 439–561 (2006) https://doi.org/10.1090/S0273-0979-06-01126-8

[45] Xie, T., Zhang, Y., Song, D.: Orion: Zero knowledge proof with linear prover time. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part IV. Lecture Notes in Computer Science, vol. 13510, pp. 299–328. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15985-5_11

[46] Di, C., Proietti, D., Telatar, I.E., Richardson, T.J., Urbanke, R.L.: Finite-length analysis of low-density parity-check codes on the binary erasure channel. IEEE Transactions on Information Theory **48**(6), 1570–1579 (2002) https://doi.org/10.1109/TIT.2002.1003839

[47] Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. Designs, Codes and Cryptography **75**(3), 565–599 (2015) https://doi.org/10.1007/s10623-014-9938-4

[48] Wang, H., Steinfeld, R., Saarinen, M.-J.O., Esgin, M.F., Yiu, S.-M.: mmCipher: Batching post-quantum public key encryption made bandwidth-optimal. In: 35th USENIX Security Symposium (USENIX Security 26). USENIX Association, Berkeley, CA (2026). Full version: Cryptology ePrint Archive, Paper 2025/1000. https://eprint.iacr.org/2025/1000

[49] Goyal, V., Kumar, A.: Non-malleable secret sharing. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018), pp. 685–698. ACM, New York, NY (2018). https://doi.org/10.1145/3188745.3188872

[50] Benhamouda, F., Degwekar, A., Ishai, Y., Rabin, T.: On the local leakage resilience of linear secret sharing schemes. Journal of Cryptology **34**(2), 10 (2021) https://doi.org/10.1007/s00145-021-09375-2

[51] Karchmer, M., Wigderson, A.: On span programs. In: Proceedings of the Eighth Annual Structure in Complexity Theory Conference, pp. 102–111. IEEE Computer Society, Los Alamitos, CA (1993). https://doi.org/10.1109/SCT.1993.336536

[52] Massey, J.L.: Some applications of source coding in cryptography. European Transactions on Telecommunications **5**(4), 421–430 (1994) https://doi.org/10.1002/ett.4460050405

[53] Beimel, A.: Secure schemes for secret sharing and key distribution. PhD thesis, Technion – Israel Institute of Technology (1996)

[54] Emde Boas, P.: Another NP-complete partition problem and the complexity

of computing short vectors in a lattice. Technical Report 81-04, Mathematisch Instituut, Universiteit van Amsterdam (1981)

[55] Micciancio, D., Goldwasser, S.: Complexity of Lattice Problems: A Cryptographic Perspective. The Kluwer International Series in Engineering and Computer Science, vol. 671. Kluwer Academic Publishers, Boston, MA (2002). https://doi.org/10.1007/978-1-4615-0897-7

[56] Babai, L.: On Lovász' lattice reduction and the nearest lattice point problem. Combinatorica **6**(1), 1–13 (1986) https://doi.org/10.1007/BF02579403

[57] Hhan, M., Kim, J., Lee, C., Son, Y.: A hybrid of lattice-reduction and Meet-LWE via near-collision on Babai's plane. Designs, Codes and Cryptography **94**, 83 (2026) https://doi.org/10.1007/s10623-026-01813-1

[58] Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. Mathematische Annalen **261**(4), 515–534 (1982) https://doi.org/10.1007/BF01457454

[59] Nguyen, P.Q., Stehlé, D.: LLL on the average. In: Hess, F., Pauli, S., Pohst, M. (eds.) Algorithmic Number Theory – ANTS-VII. Lecture Notes in Computer Science, vol. 4076, pp. 238–256. Springer, Berlin, Heidelberg (2006). https://doi.org/10.1007/11792086_18

[60] Gama, N., Nguyen, P.Q.: Predicting lattice reduction. In: Smart, N.P. (ed.) Advances in Cryptology – EUROCRYPT 2008. Lecture Notes in Computer Science, vol. 4965, pp. 31–51. Springer, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_3

[61] Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) Advances in Cryptology – EUROCRYPT 2012. Lecture Notes in Computer Science, vol. 7237, pp. 738–755. Springer, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_43

[62] Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of Learning with Errors. Journal of Mathematical Cryptology **9**(3), 169–203 (2015) https://doi.org/10.1515/jmc-2015-0016

[63] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange – A new hope. In: 25th USENIX Security Symposium (USENIX Security 16), pp. 327–343. USENIX Association, Berkeley, CA (2016). https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim

[64] Wood, J.A.: Duality for modules over finite rings and applications to coding theory. American Journal of Mathematics **121**(3), 555–575 (1999) https://doi.org/10.1353/ajm.1999.0024

[65] Dougherty, S.T.: Algebraic Coding Theory Over Finite Commutative Rings. SpringerBriefs in Mathematics. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59806-2

[66] Saarinen, M.-J.O.: Arithmetic coding and blinding countermeasures for lattice signatures: Engineering a side-channel resistant post-quantum signature scheme with compact signatures. Journal of Cryptographic Engineering **8**(1), 71–84 (2018) https://doi.org/10.1007/s13389-017-0149-6

[67] Espitau, T., Tibouchi, M., Wallet, A., Yu, Y.: Shorter hash-and-sign lattice-based signatures. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part II. Lecture Notes in Computer Science, vol. 13508, pp. 245–275. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15979-4_9