

Worst-Case Hardness for Succinct and Decomposed Learning with Errors

Damiano Abram ^{*}
University of Edinburgh

Gal Arnon [†]
NTT Research

Valerio Cini [‡]
Bocconi University

Paul Lou [§]
Bocconi University

Giulio Malavolta [¶]
Bocconi University

Lawrence Roy ^{||}
IBM Research Zürich

Abstract

The succinct LWE (CRYPTO '24) and decomposed LWE (CRYPTO '25) assumptions are two recently introduced variants of the celebrated LWE problem, known to imply the existence of numerous new cryptographic primitives. Yet, the hardness of these problems has thus far remained poorly understood, since there is no known reduction to LWE.

In this work, we study these lattice problems from a foundational perspective. Our main technical contribution is a worst-case-to-average-case reduction from *superlattice SVP* (SuperSVP) to succinct SIS. We introduce SuperSVP as a variant of SVP for general lattices in which the solver is additionally given, as a hint, a short basis for a related superlattice. Combining this result with our quantum reduction from succinct SIS to succinct LWE, which is an analog of Regev's SIS-to-LWE reduction, and with our search-to-decision reduction establishes worst-case hardness for succinct LWE. Finally, we show that succinct and decomposed LWE are equivalent, extending this conclusion to decomposed LWE.

Overall, our work places succinct and decomposed LWE on firmer theoretical ground and, as a direct consequence, establishes the *worst-case* hardness of a variety of cryptographic primitives, known to be implied by these assumptions.

^{*}abram.damiano@protonmail.com

[†]galarnon42@gmail.com. The majority of this work was done while the author was a postdoc at Bocconi University

[‡]cini.valerio@gmail.com

[§]paul96lou@gmail.com

[¶]giulio.malavolta@hotmail.it

^{||}ldr709@gmail.com

Contents

1	Introduction	3
1.1	The “Hinted Lattice” Framework	4
1.2	Our Results	6
1.3	The Superlattice SVP	7
2	Technical Overview	8
2.1	Worst-Case to Average-Case Reduction	8
2.2	Proof of Equivalence	15
2.3	Search-to-Decision Reduction	17
2.4	Quantum Reduction from Succinct SIS to Succinct LWE	19
3	Preliminaries	20
3.1	Statistical Distance Facts	21
3.2	Lattice Background	21
3.3	Computational Assumptions	24
3.4	Additional Lattice Tools	26
3.5	Quotient Decomposition and Smith Normal Form	30
4	A Worst-Case to Average-Case Reduction for Succinct SIS	31
4.1	The Worst-Case Problem	31
4.2	The Reduction and Its Parameters	31
4.3	From One Vector to a Short Basis	35
4.4	Converting the Hinted Instance to Succinct SIS	36
4.5	Lifting a Succinct SIS Solution	41
4.6	The Opposite Implication	44
5	Proof of Equivalence	46
5.1	Rerandomization of the Public Matrix	53
6	A Search-to-Decision Reduction for Decomposed LWE	58
7	A Quantum Reduction from Succinct SIS to Succinct LWE	64
8	Generalizations of Decomposed LWE	68
8.1	Extended Succinct LWE	70
8.2	Succinct LWE with a Small Secret	71
8.3	Generalized Decomposed LWE	72
8.4	Decomposed LWE with a Structured Secret	74
9	References	75

1 Introduction

The Learning With Errors (LWE) problem is the pillar of lattice-based cryptography and one of the foundations of post-quantum cryptography, due to its conjectured hardness against quantum algorithms. The (decisional) LWE assumption asserts that a noisy linear encoding $\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$ is computationally indistinguishable from a uniformly random vector even when the matrix $\mathbf{A}$ is public. Due to its rich algebraic structure, the LWE problem enables fundamental cryptographic primitives, such as fully homomorphic encryption [Gen09, vDGHV10, BV11, BGV12, GSW13], attribute-based encryption [GVW13, BGG⁺14], predicate encryption [GVW15], or classical verification of quantum computation [Mah18], to mention a few.

Nevertheless, in the current state of affairs, not all cryptographic primitives are known to be reducible from LWE. Motivated both by the construction of advanced cryptographic primitives and by the goal of improving the concrete efficiency of existing schemes, the research community has recently investigated numerous *new variants* of the SIS/LWE problem and related lattice assumptions, e.g., [ACL⁺22, AKSY22, BLNS23, CLM23, WW23b, FMN24, ARS24, Wee24, AMR25, ABCR26]. This proliferation has also prompted community efforts to catalog the resulting assumptions and clarify the reductions, generalizations, and other relationships among them. Examples include the ongoing Lattice Assumption Zoo project [LatZoo26] and the 2026 ICMS workshop *Novel Lattice Assumptions* [ICMS26].

Among these assumptions, succinct LWE [Wee24] and decomposed LWE [AMR25], together with their SIS counterparts [WW23a, Wee24, AMR26a], have enabled a growing collection of constructions. Examples include compact attribute-based encryption [Wee24, Wee25, AMR25, LZF26], distributed and adaptively secure broadcast encryption [CW24, CHW25, WW25, GY26c], flexible and identity-based broadcast encryption [GY26a], registered attribute-based encryption [AMR25, CHW25, WW25], laconic function evaluation [Wee24, AMR25], and succinct randomized encodings and garbled circuits [AMR25, LLL26]. Further applications include distributed monotone-policy encryption [CW26a, AGY26a, CW26b, AGY26b, GY26b], silent threshold encryption and bounded-collusion registered functional encryption [CWY26], succinct computational secret sharing [CW26b], secure group messaging [BBD⁺26], functional commitments [WW23a], and correlation-intractable hash functions for batched relations [AMR26a]. Under additional assumptions, succinct LWE also supports extensions of registered ABE with certified deletion and everlasting security [MSM25]. Yet, the hardness of these problems has thus far remained opaque. Given their usefulness in the realm of cryptography, an outstanding open question is to relate their difficulty to regular lattice problems and to shed some light on where such hardness may come from.

The objective of this work is to understand and solidify the foundations of decomposed LWE and succinct LWE: We investigate how these assumptions relate to one another, how variants of these assumptions connect, and which worst-case lattice problems can support their hardness. In more detail, we show that (1) the (average-case) succinct SIS problem is at least as hard as a *worst-case* shortest vector problem over general lattices; (2) succinct and decomposed LWE are equivalent and their search and decision variants are also equivalent; (3) succinct SIS reduces to succinct LWE with a quantum reduction. Before stating our results more precisely, we recall the definitions of succinct and decomposed LWE and review their common structure, through the lens of *hinted* lattice problems.

1.1 The “Hinted Lattice” Framework

The following “hinted lattice” framework was introduced by Wu [Wu25] as a common language for several recent lattice assumptions. In this framework, an LWE public matrix $\mathbf{A}$ is accompanied by an auxiliary matrix constructed from $\mathbf{A}$, together with correlated short vectors forming a gadget trapdoor for the auxiliary matrix. The security question is whether SIS or LWE with respect to $\mathbf{A}$ remains hard even when this additional information is revealed.

In more detail, consider a wide matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and matrices $\mathbf{W}_1, \dots, \mathbf{W}_\ell \in \mathbb{Z}_q^{n \times k}$, and define

$$\mathbf{D}_\ell := \left[\begin{array}{c|c} \mathbf{A} & \mathbf{W}_1 \\ & \vdots \\ & \mathbf{W}_\ell \end{array} \right] = [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}] \in \mathbb{Z}_q^{\ell n \times (\ell m + k)}, \quad \mathbf{W} := \begin{bmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_\ell \end{bmatrix} \in \mathbb{Z}_q^{\ell n \times k}.$$

The associated trapdoor is a short matrix $\mathbf{T} \in \mathbb{Z}^{(\ell m + k) \times \ell n \lceil \log q \rceil}$ satisfying

$$\mathbf{D}_\ell \mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q, \tag{1}$$

where $\mathbf{G}$ is the usual gadget matrix [MP12]. The general security question is whether SIS or LWE with respect to $\mathbf{A}$ remains hard when both $\mathbf{D}_\ell$ and the short trapdoor $\mathbf{T}$ are made public.

For SIS, the task is to find a short nonzero integer vector $\mathbf{x}$ such that $\mathbf{A}\mathbf{x} = \mathbf{0} \bmod q$. For the LWE variants considered here, the corresponding security assumption states that

$$(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{W}, \mathbf{T}) \approx_c (\mathbf{A}, \mathbf{u}^\top, \mathbf{W}, \mathbf{T}),$$

where $\mathbf{s}$ is uniform over $\mathbb{Z}_q^n$, $\mathbf{u}$ is uniform over $\mathbb{Z}_q^m$, and $\mathbf{e} \in \mathbb{Z}^m$ is sampled from a discrete Gaussian.

Succinct LWE and SIS. The ℓ -succinct LWE assumption, introduced in [Wee24], considers the matrices $\mathbf{A}$ and $\mathbf{W}_i$ uniformly sampled, and a short preimage $\mathbf{T} \leftarrow \mathbf{D}_\ell^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$.¹ In other words, $\mathbf{T}$ is sampled from a discrete Gaussian of parameter σ , conditioned on the relation in Equation (1) holding. Although $\mathbf{T}$ provides short preimages under $\mathbf{D}_\ell$ of all the columns of $\mathbf{I}_\ell \otimes \mathbf{G}$, it does not immediately help distinguish LWE samples associated with the uniformly sampled matrix $\mathbf{A}$ from uniform. The succinct SIS assumption [Wee24] uses the same public tuple $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ and asserts that finding a short nonzero vector in the kernel of $\mathbf{A}$ modulo q remains hard.

The adjective “succinct” reflects the fact that this correlated trapdoor can be used to compress large collections of LWE samples and matrices. This compression has enabled constructions with parameters that depend only logarithmically, or not at all, on the size of the represented object, including succinct attribute-based encryption, functional commitments, registered attribute-based encryption, and distributed broadcast encryption [WW23a, Wee24, CW24, Wee25, CHW25, WW25]. When $k = \Omega(\ell n \log q)$ is sufficiently large, decisional succinct LWE is known to be equivalent to standard decisional LWE [Wee24, Remark 1]. In applications, by contrast, one needs the bit representation of each $\mathbf{W}_i$ to have size $o(\ell) \cdot \text{poly}(n, \log q)$; see, for example, [Wu25].

¹This notation is formally introduced in Equation (3).

Decomposed LWE and SIS. Abram, Malavolta, and Roy [AMR25] introduced decomposed LWE directly as a hardness assumption for LWE with respect to a structured public matrix. Let

$$\begin{aligned}\overline{\mathbf{W}} &:= [\mathbf{W}_1 \ \cdots \ \mathbf{W}_\ell] \in \mathbb{Z}_q^{n \times \ell k}, \\ \mathbf{R} &:= [\mathbf{R}_1 \ \cdots \ \mathbf{R}_\ell] \in \mathbb{Z}^{k \times \ell n \lceil \log q \rceil},\end{aligned}$$

where the matrices $\mathbf{W}_i \in \mathbb{Z}_q^{n \times k}$ are uniform and the matrices $\mathbf{R}_j \in \mathbb{Z}^{k \times n \lceil \log q \rceil}$ are short, typically sampled from a discrete Gaussian. Define

$$\mathbf{A} := \overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G} \in \mathbb{Z}_q^{n \times \ell^2 n \lceil \log q \rceil}. \quad (2)$$

Equivalently, the block of $\mathbf{A}$ indexed by $(i, j) \in [\ell]^2$ is $\mathbf{W}_i \mathbf{R}_j + \delta_{ij} \mathbf{G}$. The decomposed LWE assumption states that

$$(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \overline{\mathbf{W}}, \mathbf{R}) \approx_c (\mathbf{u}^\top, \overline{\mathbf{W}}, \mathbf{R}),$$

where $\mathbf{u} \sim \mathbb{Z}_q^{\ell^2 n \lceil \log q \rceil}$ is uniform. Thus, the original definition does not refer to an auxiliary matrix or a trapdoor. It simply requires LWE to remain hard for the structured matrix (2), even when its decomposition $(\overline{\mathbf{W}}, \mathbf{R})$ is revealed. The decomposed SIS assumption [AMR26a] asserts that finding a short nonzero integer vector $\mathbf{x}$ satisfying $\mathbf{A}\mathbf{x} = \mathbf{0} \bmod q$ for the same structured matrix, given its decomposition $(\overline{\mathbf{W}}, \mathbf{R})$, is hard.

Nevertheless, Wu [Wu25] observed that this short matrix $\mathbf{R}$ implicitly determines a short gadget trapdoor for a related matrix, and hence that decomposed LWE can also be expressed within the hinted LWE framework. Let us now elaborate on this point. Let

$$\mathbf{W} := \begin{bmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_\ell \end{bmatrix} \in \mathbb{Z}_q^{\ell n \times k}$$

be the block transpose of $\overline{\mathbf{W}}$, and define

$$\mathbf{D}_\ell := [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}] \in \mathbb{Z}_q^{\ell n \times (\ell^3 n \lceil \log q \rceil + k)}.$$

Consider the matrix

$$\mathbf{T} := \begin{bmatrix} \text{vec}(\mathbf{I}_\ell) \otimes \mathbf{I}_{\ell n \lceil \log q \rceil} \\ -\mathbf{R} \end{bmatrix} \in \mathbb{Z}^{(\ell^3 n \lceil \log q \rceil + k) \times \ell n \lceil \log q \rceil}.$$

Since $\mathbf{R}$ is short and every entry of the upper block belongs to $\{0, 1\}$, we have $\|\mathbf{T}\| = \max\{1, \|\mathbf{R}\|\}$, and hence $\mathbf{T}$ is short. Moreover, the blockwise definition of $\mathbf{A}$ gives

$$\mathbf{D}_\ell \mathbf{T} = (\mathbf{I}_\ell \otimes \mathbf{A}) (\text{vec}(\mathbf{I}_\ell) \otimes \mathbf{I}_{\ell n \lceil \log q \rceil}) - \mathbf{W} \mathbf{R} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q.$$

Therefore, $\mathbf{T}$ is a short gadget trapdoor for $\mathbf{D}_\ell$,² and decomposed LWE can equivalently be written in the hinted LWE form

$$(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{W}, \mathbf{T}) \approx_c (\mathbf{A}, \mathbf{u}^\top, \mathbf{W}, \mathbf{T}).$$

²Under the dimension and Gaussian width conditions of Lemma 3.6, $\mathbf{T}$ can be used by the standard lattice preimage-sampling procedure to sample a fresh trapdoor whose distribution is statistically close to $\mathbf{D}_\ell^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$ if necessary.

Indeed, $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ can be computed from $(\mathbf{W}, \mathbf{R})$, while $\mathbf{W}$ is obtained by reordering the blocks of $\mathbf{W}$ and $\mathbf{R}$ is recovered from the bottom k rows of $-\mathbf{T}$. Let us remark, however, that casting the assumption in the hinted lattice framework does not automatically make it identical to succinct LWE/SIS, since the matrices are sampled from a different distribution. Indeed, the exact relation between these two problems has thus far remained opaque.




Figure 1: An arrow from one problem to another denotes a reduction that solves instances of the problem at its tail using a solver for the problem at its head. Blue arrows indicate our contributions. Parameters may change along the arrows, as specified in the cited results; the decomposed search-to-decision reduction and its composition with the LWE-to-SIS reduction assume standard LWE hardness. The LWE-to-SIS direction also requires the noise and solution norm bounds in Lemma 5.6. The search LWE and SIS analogs use the corresponding LWE transformations. Dashed arrows marked “quantum” use quantum computation; all other reductions are classical.

1.2 Our Results

We are now ready to state our results more precisely. Our main technical contribution is a *worst-case to average-case* reduction for succinct SIS. We define a worst-case lattice problem, which we call *superlattice SVP* (SuperSVP), as follows: The input is a worst-case lattice Λ together with a short basis of a superlattice $\Lambda' \supseteq \Lambda^{\oplus \ell}$ of bounded index, where $\Lambda^{\oplus \ell} := \Lambda \oplus \dots \oplus \Lambda$ denotes the lattice obtained by taking the orthogonal direct sum of ℓ copies of Λ . Put equivalently, a basis of $\Lambda^{\oplus \ell}$ is given by $\mathbf{I}_\ell \otimes \mathbf{L}$ where $\mathbf{L}$ is a basis of Λ and where $\mathbf{I}_\ell$ is the $\ell \times \ell$ identity matrix. The task is to find a short nonzero vector in Λ , even with this auxiliary basis available. We give reductions in both directions between this worst-case problem and average-case succinct SIS, with the parameter changes specified in the formal statements.

Theorem 1.1 (Informal). *Worst-case SuperSVP $\equiv$ average-case succinct SIS.*

Here $\equiv$ denotes polynomial-time reductions in both directions, under suitable parameter conditions. The worst-case to average-case reduction preserves the lattice dimension. If $|\Lambda'/\Lambda^{\oplus \ell}| \leq$

2^k , it uses at most $2k$ auxiliary columns, so the number of columns of $\mathbf{W}$ is controlled by the logarithm of the size of the quotient. This gives a geometric source of hardness for succinct SIS that explicitly accounts for its public trapdoor. See Theorem 4.2 for the precise parameter conditions and the resulting norm bounds. Conversely, Theorem 4.10 reduces succinct SIS with an $n \times m$ public matrix to SuperSVP in dimension m , with quotient bound $k' \geq k \log q$.

To connect this worst-case hardness problem to the *decomposed* family of assumptions, we prove that succinct LWE and decomposed LWE are equivalent. Previously, the hardness of decomposed LWE was known to follow from that of an appropriate instantiation of succinct LWE [AMR25, Theorem 4.2]. We establish the converse reduction, namely that the hardness of succinct LWE follows from an appropriate instantiation of decomposed LWE.

Theorem 1.2 (Informal). *Decisional succinct LWE $\equiv$ decisional decomposed LWE.*

Our new reduction preserves the secret dimension n , the number of blocks ℓ , the auxiliary dimension k , the challenge length $m = \ell^2 n \lceil \log q \rceil$, and the modulus q . Since $\mathbf{W}_i \in \mathbb{Z}_q^{n \times k}$, the bit representation of each matrix $\mathbf{W}_i$ has size $nk \lceil \log q \rceil$ in both the original decomposed LWE instance and the resulting succinct LWE instance. Consequently, if the decomposed LWE instance lies in the application-relevant regime $nk \lceil \log q \rceil = o(\ell) \cdot \text{poly}(n, \log q)$, then the succinct LWE instance produced by the reduction lies in *the same regime*. The reduction is compatible with polynomial moduli; see Theorem 5.1 for the precise conditions. The same transformations establish the analogous equivalence for succinct and decomposed SIS, with polynomial losses in the solution norm.

To relate the search and decisional variants, we give a search-to-decision reduction for decomposed LWE. The reduction applies to polynomially bounded prime moduli and preserves the secret dimension, auxiliary dimension, modulus, and Gaussian parameters. We refer to Theorem 6.1 for a precise statement.

Theorem 1.3 (Informal). *Assuming standard LWE is hard, search decomposed LWE $\equiv$ decisional decomposed LWE.*

Finally, to complete the picture of the equivalence between these two families of assumptions, we apply Regev’s idea [Reg09] to prove a *quantum* reduction from succinct SIS to search succinct LWE. See Theorem 7.1 for the formal statement.

Theorem 1.4 (Informal). *An efficient solver for search succinct LWE yields an efficient quantum solver for succinct SIS.*

A schematic overview of our results is given in Figure 1. Taken together, our results establish a formal link between decomposed and succinct LWE and the worst-case superlattice SVP.

1.3 The Superlattice SVP

The superlattice SVP is the natural SVP analog in the context of the hinted lattice framework (Section 1.1), but given that we introduce this problem here for the first time, a few remarks are in order. The problem is worst-case over a basis $\mathbf{L}$ of $\Lambda = \mathbb{L}\mathbb{Z}^n$, together with a short basis $\mathbf{P}$ of a superlattice

$$\Lambda' = \mathbf{P}\mathbb{Z}^{\ell n} \supseteq \Lambda^{\oplus \ell}, \quad |\Lambda' / \Lambda^{\oplus \ell}| \leq 2^k.$$

The goal is to find a short nonzero vector of Λ . Although the columns of $\mathbf{P}$ are short (say β -bounded), they need not belong to $\Lambda^{\oplus \ell}$: obtaining a vector of the original lattice Λ requires a short combination whose class in $\Lambda'/\Lambda^{\oplus \ell}$ vanishes.

The bound on the index is informative about how much useful information the hint provides. When the allowed index is sufficiently large, a valid hint can be constructed directly from the basis $\mathbf{L}$. Indeed, choose an integer

$$t \geq \frac{\|\mathbf{L}\|}{\beta} \quad \text{and set} \quad \mathbf{P} = \frac{1}{t}(\mathbf{I}_\ell \otimes \mathbf{L}).$$

Here $\|\mathbf{L}\|$ is the maximum Euclidean column norm. This basis generates $\Lambda' = \frac{1}{t}\Lambda^{\oplus \ell}$ and has column lengths at most β . Moreover,

$$|\Lambda'/\Lambda^{\oplus \ell}| = t^{\ell n}.$$

Consequently, whenever $k \geq \ell n \log t$, any lattice Λ instance can be equipped efficiently with a valid hint, without changing either the lattice or the required output length. Thus, for a sufficiently permissive index bound, the problem is polynomial-time equivalent to ordinary SVP. This is consistent with known reductions for succinct LWE, in which the correlated trapdoor can be generated using a trapdoor for the auxiliary matrix itself, with appropriate Gaussian parameters [Wee24, Remark 1].

At the other extreme, the hint gives an elementary algorithm whenever the requested output bound is sufficiently large: since

$$|\Lambda'/\Lambda^{\oplus \ell}| \Lambda' \subseteq \Lambda^{\oplus \ell},$$

multiplying any column of $\mathbf{P}$ by this index and taking a nonzero block produces a nonzero vector of Λ of length at most $2^k \beta$. In particular, the problem is easy when the norm of the solution is allowed to be $2^k \beta$. Hardness must therefore be considered jointly as a function of the index bound, the hint length, and the required output length.

2 Technical Overview

2.1 Worst-Case to Average-Case Reduction

We begin with a simpler setting that isolates the mechanism of the reduction. Take $\ell = 1$ and suppose that, in addition to the input basis $\mathbf{L}$ of $\Lambda = \mathbf{L}\mathbb{Z}^n$, we are also given a short basis $\mathbf{P}$ of Λ itself (instead of the superlattice Λ' as in the actual definition of the problem). A column of $\mathbf{P}$ already gives a short nonzero lattice vector, so the purpose of this warm-up is only to explain how a good basis can supply the trapdoor information required by succinct SIS while preserving the connection between SIS solutions and short lattice vectors.

The argument has two parts. First, we construct an integer matrix $\mathbf{Y}$ with an efficient short-preimage procedure and a useful compatibility property with Λ . Second, we use $\mathbf{Y}$ to augment the Micciancio–Regev worst-case-to-average-case reduction for SIS [MR04]. No auxiliary columns $\mathbf{W}$ are needed in this warm-up: the published trapdoor will be for $\mathbf{A}$ itself.

Once the mechanism is clear, we will return to the full setting with $\ell > 1$ and a basis $\mathbf{P}$ of the superlattice Λ' . The same ideas will apply.

From a good basis to a matrix with short preimages. The matrix Y should have two properties. For every target $\mathbf{u} \in \mathbb{Z}_q^n$, the good basis should let us efficiently find a short integer vector $\mathbf{z}$ such that $Y\mathbf{z} \equiv \mathbf{u} \pmod{q}$. At the same time, the matrix LY/q should have small norm, which implies that it maps short integer vectors to short vectors in Λ/q . The first property will let us construct the succinct SIS trapdoor; the second will keep the lattice vector recovered from an SIS solution short.

To construct such a preimage, let $\tilde{\mathbf{u}} \in \mathbb{Z}^n$ be an integer representative of $\mathbf{u}$. It determines the vector $L\tilde{\mathbf{u}}/q \in \Lambda/q$. This vector is mapped to $\mathbf{u}$ under multiplication by qL^{-1} modulo q . Using the good basis P , Babai's nearest-plane algorithm finds a nearby point $\mathbf{u}' \in \Lambda$. Let $\mathbf{r}$ denote the residual vector, namely the difference between the input point $L\tilde{\mathbf{u}}/q$ and the lattice point $\mathbf{u}'$ returned by Babai's algorithm:

$$\mathbf{r} = \frac{1}{q}L\tilde{\mathbf{u}} - \mathbf{u}'.$$

Babai's guarantee gives $\|\mathbf{r}\| \leq \frac{\|P\|\sqrt{n}}{2}$. Moreover, since $\mathbf{u}' \in \Lambda$, we have $L^{-1}\mathbf{u}' \in \mathbb{Z}^n$, and hence

$$qL^{-1}\mathbf{r} = \tilde{\mathbf{u}} - qL^{-1}\mathbf{u}' \equiv \mathbf{u} \pmod{q}.$$

Thus, $\mathbf{r}$ is short and it is also mapped to $\mathbf{u}$ by qL^{-1} modulo q .

The problem is that the vector $\mathbf{r}$ is not necessarily integral, but the preimage-sampling procedure must return a short integer preimage. We therefore construct a fixed integer matrix Y and a short integer vector $\mathbf{z}$ such that

$$Y\mathbf{z} = qL^{-1}\mathbf{r},$$

which will ensure that $\mathbf{z}$ will be a preimage of the original target $\mathbf{u}$ under Y .

To do so, choose a scaling factor p with $1 \leq p < q$ (the reason for introducing this additional parameter p will become clear shortly) and set

$$\mathbf{K} = \left\lfloor \frac{q}{p}L^{-1} \right\rfloor, \quad \mathbf{z}_1 = \lfloor p\mathbf{r} \rfloor, \quad \mathbf{e} = qL^{-1}\mathbf{r} - \mathbf{K}\mathbf{z}_1.$$

Rounding is entrywise to a nearest integer. The vector $\mathbf{e}$ can be interpreted as the term introduced after using the product $\mathbf{K}\mathbf{z}_1$ of integer matrices to approximate the desired vector $qL^{-1}\mathbf{r}$. The vector $\mathbf{z}_1$ is short because $\mathbf{r}$ is short. Since each rounding changes an entry by at most $1/2$, we use the bound on $\mathbf{r}$ to choose $M \in \mathbb{N}$ such that every coordinate of $\mathbf{e}$ has absolute value strictly smaller than M . It follows that there exists $\mathbf{z}_2 \in \{-1, 0, 1\}^{n \lceil \log M \rceil}$ such that $\mathbf{e} = \mathbf{G}_{n,M}\mathbf{z}_2$. Consequently, the fixed integer matrix

$$Y = [\mathbf{K} \mid \mathbf{G}_{n,M}]$$

has the desired preimages:

$$Y \begin{bmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{bmatrix} = [\mathbf{K} \mid \mathbf{G}_{n,M}] \begin{bmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{bmatrix} = \mathbf{K}\mathbf{z}_1 + \mathbf{e} = qL^{-1}\mathbf{r} \equiv \mathbf{u} \pmod{q}.$$

Why this matrix is compatible with the lattice. The relevant norm we are interested in, as discussed in the beginning, is not that of Y itself, but that of

$$\frac{1}{q}LY = [LK/q \mid LG_{n,M}/q].$$

Here, since $\mathbf{K} = \left\lfloor \frac{q}{p} \mathbf{L}^{-1} \right\rfloor$, the first block $\frac{1}{q} \mathbf{L} \mathbf{K} = \frac{1}{p} \mathbf{I}_n + \frac{1}{q} \mathbf{L} \left(\left\lfloor \frac{q}{p} \mathbf{L}^{-1} \right\rfloor - \frac{q}{p} \mathbf{L}^{-1} \right)$ is approximately $\mathbf{I}_n/p$, so its norm decreases as p increases. The second block, which accounts for the correction term, has norm that increases with p (since the bound M on the norm of $\mathbf{e}$ increases with p), but decreases as q increases. It follows that, after choosing p so that the norm of the first block is sufficiently small, the norm of the second block can also be made sufficiently small by choosing q large enough. The precise parameters are discussed in Lemma 4.6.

We have therefore constructed an integer matrix $\mathbf{Y}$ with the two desired properties: using the good basis, we can find a short integer preimage under $\mathbf{Y}$ of every target in $\mathbb{Z}_q^n$. Moreover, $\mathbf{L} \mathbf{Y}/q$ has small norm, so multiplying it by a short integer vector produces a short vector in Λ/q . We now explain how these properties are used in the reduction.

Adding the trapdoor to the SIS reduction. We build on the Micciancio–Regev worst-case-to-average-case reduction for SIS [MR04]. Recall that the reduction works as follows: given the lattice Λ , sample independent short vectors from a discrete Gaussian of appropriate width over Λ/q and collect them as the columns of a private matrix $\mathbf{V}$. The corresponding public SIS matrix is

$$\mathbf{A}_0 = q \mathbf{L}^{-1} \mathbf{V} \bmod q.$$

If the Gaussian width used to sample $\mathbf{V}$ is sufficiently large, the matrix $\mathbf{A}_0$ is statistically close to uniform. More importantly, a relation $\mathbf{A}_0 \mathbf{x} \equiv 0 \bmod q$ implies

$$q \mathbf{L}^{-1} \mathbf{V} \mathbf{x} \equiv 0 \bmod q, \quad \text{and hence} \quad \mathbf{L}^{-1} \mathbf{V} \mathbf{x} \in \mathbb{Z}^n.$$

Therefore, $\mathbf{V} \mathbf{x} \in \Lambda$. If $\mathbf{x}$ is short, then $\mathbf{V} \mathbf{x}$ is also short because the columns of $\mathbf{V}$ are short Gaussian samples.

In our construction, to retain this connection while adding the trapdoor, sample a random short matrix $\mathbf{X}$ and define

$$\mathbf{A} = [\mathbf{A}_0 \mid \mathbf{A}_0 \mathbf{X} + \mathbf{Y}] \bmod q.$$

This has the same masking structure as the Micciancio–Peikert trapdoor construction [MP12]: our matrix $\mathbf{Y}$ replaces the gadget matrix used there. We next explain how to construct the trapdoor, why $\mathbf{A}$ is nearly uniform, and how to recover a short lattice vector from the solver’s answer.

Constructing the published trapdoor. Given a target $\mathbf{u}$, find a short $\mathbf{z}$ with $\mathbf{Y} \mathbf{z} \equiv \mathbf{u} \bmod q$. Then

$$\mathbf{A} \begin{bmatrix} -\mathbf{X} \mathbf{z} \\ \mathbf{z} \end{bmatrix} \equiv \mathbf{Y} \mathbf{z} \equiv \mathbf{u} \bmod q.$$

The first block cancels the random mask, and the preimage remains short because $\mathbf{X}$ has low norm. Apply this procedure to every column of the gadget matrix $\mathbf{G}_{n,q}$. Collecting the preimages gives a short matrix $\hat{\mathbf{T}}$ satisfying $\mathbf{A} \hat{\mathbf{T}} \equiv \mathbf{G}_{n,q}$.

Using standard lattice preimage-sampling techniques, $\hat{\mathbf{T}}$ can then be used to sample a trapdoor $\mathbf{T}$ with the correct distribution required by succinct SIS.

Why the public matrix is close to uniform. If the matrix $\mathbf{X}$ has enough entropy, the leftover hash lemma implies that $(\mathbf{A}_0, \mathbf{A}_0 \mathbf{X})$ is jointly close to uniform. Adding the fixed $\mathbf{Y}$ to the second block preserves this property, so $\mathbf{A}$ is statistically close to uniform too. Together with Gaussian trapdoor sampling, this gives the intended succinct SIS input distribution in the warm-up setting.

Recovering a short lattice vector. Give (A, T) to the succinct SIS solver. Suppose it returns a short nonzero $\mathbf{x} = (\mathbf{x}_0, \mathbf{x}_1)$ with $A\mathbf{x} \equiv 0 \pmod q$, split according to the two blocks of A . Using our private information V, X, Y , compute

$$\mathbf{v} = V(\mathbf{x}_0 + X\mathbf{x}_1) + \frac{1}{q}LY\mathbf{x}_1.$$

It follows that

$$qL^{-1}\mathbf{v} \equiv A_0(\mathbf{x}_0 + X\mathbf{x}_1) + Y\mathbf{x}_1 \equiv A\mathbf{x} \equiv 0 \pmod q,$$

so $\mathbf{v} \in \Lambda$ as before. The first summand in the definition of $\mathbf{v}$ is short because the Gaussian samples V , the sSIS solution $\mathbf{x}$, and the mask X are all low norm. The second is short by the compatibility property of Y . Finally, using standard techniques, one can show that the short lattice vector is nonzero with overwhelming probability, completing the warm-up reduction. We now adapt this construction to the full setting, where the given good basis is for the superlattice Λ' .

Using a basis of a superlattice. We now return to the actual input of the reduction. For arbitrary ℓ , write $\Lambda_\ell = \Lambda^{\oplus \ell}$ for the lattice consisting of ℓ copies of Λ , with basis $L_\ell = I_\ell \otimes L$. The short basis P we are given is a basis of a superlattice $\Lambda' \supseteq \Lambda_\ell$, which contains at most 2^k distinct cosets of Λ_ℓ . Our strategy is to retain the construction of A from the warm-up and use the additional block W to account for the contribution of these cosets when constructing the trapdoor.

To see what needs to change, consider an arbitrary target $\mathbf{u} \in \mathbb{Z}_q^{\ell n}$ with integer representative $\tilde{\mathbf{u}}$. As before, we apply Babai's algorithm to $L_\ell \tilde{\mathbf{u}}/q$, using the good basis P . The algorithm returns a point $\mathbf{u}' \in \Lambda'$ close to the input point. In particular, we have that

$$\mathbf{r} := \frac{1}{q}L_\ell \tilde{\mathbf{u}} - \mathbf{u}', \quad \|\mathbf{r}\| \leq \frac{\|P\|\sqrt{\ell n}}{2}.$$

As before, we refer to $\mathbf{r}$ as the residual. The problem is now that $\mathbf{u}'$ belongs to Λ' and need not belong to Λ_ℓ . Consider

$$qL_\ell^{-1}\mathbf{r} = \tilde{\mathbf{u}} - qL_\ell^{-1}\mathbf{u}'.$$

In the warm-up, the second term on the right-hand side belonged to $q\mathbb{Z}^{\ell n}$ and therefore vanished modulo q , leaving $qL_\ell^{-1}\mathbf{r}$ congruent to $\tilde{\mathbf{u}} = \mathbf{u} \pmod q$. Here this is no longer guaranteed. We must therefore account for this extra term before using the result from the warm-up.

Keeping only the quotient contribution. When constructing a short preimage of $\mathbf{u}$, the additional block W is used to account exactly for this extra contribution. The crucial observation is that it suffices to represent only its coset modulo Λ_ℓ . Indeed, replacing $\mathbf{u}'$ by $\mathbf{u}' + \mathbf{a}$ for any $\mathbf{a} \in \Lambda_\ell$ changes its contribution by

$$qL_\ell^{-1}(\mathbf{u}' + \mathbf{a}) - qL_\ell^{-1}\mathbf{u}' = qL_\ell^{-1}\mathbf{a} \in q\mathbb{Z}^{\ell n}.$$

Consequently, only the coset of $\mathbf{u}'$ modulo Λ_ℓ matters. These cosets form the finite quotient $\Gamma := \Lambda'/\Lambda_\ell$, and by the promise on the input, its size satisfies

$$|\Gamma| \leq 2^k.$$

This is where the promise on the quotient size becomes useful: the extra information we need to retain identifies one of at most 2^k cosets.

Encoding the quotient with short coefficients. The goal is to represent every coset in Λ'/Λ_ℓ using short integer linear combinations of a fixed collection of vectors. By the standard decomposition of finite abelian groups, there exist vectors $\mathbf{b}_1, \dots, \mathbf{b}_h \in \Lambda'$ together with integers $d_1, \dots, d_h \geq 2$, such that every coset has a unique linear combination of these vectors with coefficients $0 \leq y_j < d_j$. Write $\mathbf{B} = [\mathbf{b}_1 \mid \dots \mid \mathbf{b}_h]$, and

$$\mathbf{G}_\Gamma := \left[\text{diag}(\mathbf{g}_{d_1}^\top, \dots, \mathbf{g}_{d_h}^\top) \mid \mathbf{0} \right] \in \mathbb{Z}^{h \times 2k},$$

where enough zero columns are appended to make the total number of columns equal to $2k$.³ We can therefore express the vector returned by Babai's algorithm as

$$\mathbf{u}' = \mathbf{a} + \mathbf{B}\mathbf{y}, \quad \mathbf{a} \in \Lambda_\ell, \quad 0 \leq y_i < d_i \quad \forall i \in [h].$$

Since each y_j lies in the known range $0 \leq y_j < d_j$, taking their binary representation gives us a vector $\mathbf{z}_3$ such that

$$\mathbf{y} = \mathbf{G}_\Gamma \cdot \mathbf{z}_3, \quad \mathbf{z}_3 \in \{0, 1\}^{2k}.$$

We show how to handle the contribution of $\mathbf{B}\mathbf{G}_\Gamma\mathbf{z}_3$ next.

Making the quotient contribution integral. As in the warm-up, turning this geometric representation into an integer matrix requires an additional step. The columns of $\mathbf{B}\mathbf{G}_\Gamma$ need not belong to Λ_ℓ/q ; equivalently, $q\mathbf{L}_\ell^{-1}\mathbf{B}\mathbf{G}_\Gamma$ need not be integral. Therefore, we cannot use this matrix directly as part of the sSIS instance. We therefore apply Babai's nearest-plane algorithm, using the scaled basis $\mathbf{L}_\ell/q$, to round each column to a nearby point of Λ_ℓ/q , obtaining a matrix $\mathbf{B}'$ with each column in Λ_ℓ/q . This ensures that

$$\mathbf{C} := q\mathbf{L}_\ell^{-1}\mathbf{B}'$$

is an integer matrix that can be used to construct $\mathbf{W}$, and that $\|\mathbf{B}' - \mathbf{B}\mathbf{G}_\Gamma\|$ is small with a bound proportional to $1/q$.

Recall that $\mathbf{u}' = \mathbf{a} + \mathbf{B}\mathbf{G}_\Gamma\mathbf{z}_3$ for some $\mathbf{a} \in \Lambda_\ell$ and that $\mathbf{u}' = \frac{1}{q}\mathbf{L}_\ell\tilde{\mathbf{u}} - \mathbf{r}$, hence

$$\frac{1}{q}\mathbf{L}_\ell\tilde{\mathbf{u}} = \mathbf{r} + \mathbf{a} + \mathbf{B}\mathbf{G}_\Gamma\mathbf{z}_3.$$

When we replace $\mathbf{B}\mathbf{G}_\Gamma$ by $\mathbf{B}'$, we can account for the rounding error introduced and define

$$\mathbf{r}_{\text{corr}} := \mathbf{r} + (\mathbf{B}\mathbf{G}_\Gamma - \mathbf{B}')\mathbf{z}_3.$$

We refer to this vector as the corrected residual. Since the residual $\mathbf{r}$ is short, the vector $\mathbf{z}_3$ is binary, and the rounding error $\mathbf{B}\mathbf{G}_\Gamma - \mathbf{B}'$ can be made small by choosing q sufficiently large, $\mathbf{r}_{\text{corr}}$ is also short.

By the definition of the corrected residual,

$$\frac{1}{q}\mathbf{L}_\ell\tilde{\mathbf{u}} = \mathbf{r} + \mathbf{a} + \mathbf{B}\mathbf{G}_\Gamma\mathbf{z}_3 = \mathbf{r}_{\text{corr}} + \mathbf{a} + \mathbf{B}'\mathbf{z}_3.$$

³Notice that this padding is always possible since $\sum_{j=1}^h \lceil \log d_j \rceil \leq 2 \sum_{j=1}^h \log d_j = 2 \log |\Gamma| \leq 2k$.

This equality shows that $\mathbf{r}_{\text{corr}} \in \Lambda_\ell/q$ since every other term appearing in the above equation also belongs to it. Applying $q\mathbf{L}_\ell^{-1}$ therefore gives the integer relation

$$q\mathbf{L}_\ell^{-1}\mathbf{r}_{\text{corr}} + \mathbf{C}\mathbf{z}_3 \equiv \mathbf{u} \bmod q.$$

Thus, we have decomposed the target $\mathbf{u}$ into two integer contributions. The first comes from the short corrected residual and can be treated using ℓ copies of the warm-up construction. The second uses the binary vector $\mathbf{z}_3$ and will be incorporated into the additional matrix $\mathbf{W}$.

Combining with the warm-up. We now put these two contributions together to obtain short preimages for the matrix required by sSIS. We use the same matrix

$$\mathbf{A} = [\mathbf{A}_0 \mid \mathbf{A}_0\mathbf{X} + \mathbf{Y}] \bmod q,$$

as before. Applying the warm-up construction gives a short integer vector $\mathbf{t}$ such that

$$(\mathbf{I}_\ell \otimes \mathbf{A})\mathbf{t} \equiv q\mathbf{L}_\ell^{-1}\mathbf{r}_{\text{corr}} \bmod q.$$

To account for the quotient contribution, we hide $\mathbf{C}$ using another random low-norm matrix $\mathbf{S}$ and set

$$\mathbf{W} = \mathbf{C} + (\mathbf{I}_\ell \otimes \mathbf{A}_0)\mathbf{S} \bmod q.$$

We use $\mathbf{z}_3$ as the coefficients multiplying $\mathbf{W}$. This introduces the additional term $(\mathbf{I}_\ell \otimes \mathbf{A}_0)\mathbf{S}\mathbf{z}_3$, which we cancel as in the warm-up by splitting $\mathbf{S}\mathbf{z}_3$ into ℓ blocks, one for each copy of $\mathbf{A}_0$. In each copy of $\mathbf{A}$, subtract the corresponding block from the entries of $\mathbf{t}$ multiplying $\mathbf{A}_0$. Denote the resulting vector by $\mathbf{t}'$; all other entries remain unchanged. It follows that

$$[\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}] \begin{bmatrix} \mathbf{t}' \\ \mathbf{z}_3 \end{bmatrix} \equiv (\mathbf{I}_\ell \otimes \mathbf{A})\mathbf{t} + \mathbf{C}\mathbf{z}_3 \equiv q\mathbf{L}_\ell^{-1}\mathbf{r}_{\text{corr}} + \mathbf{C}\mathbf{z}_3 \equiv \mathbf{u} \bmod q.$$

This preimage remains short because $\mathbf{t}$, $\mathbf{S}$, and $\mathbf{z}_3$ have low norm. Moreover, $\mathbf{W}$ has only $2k$ columns.

Publishing the trapdoor and using the sSIS solution. We apply the same trapdoor construction and Gaussian resampling as in the warm-up, now to $[\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$. Together with the leftover hash argument for the masks $\mathbf{X}$ and $\mathbf{S}$, this gives the required distribution for the sSIS instance $(\mathbf{A}, \mathbf{W}, \mathbf{T})$. The solver receives this instance and returns a short nonzero vector $\mathbf{x} = (\mathbf{x}_0, \mathbf{x}_1)$ satisfying $\mathbf{A}\mathbf{x} \equiv 0 \bmod q$. Although $\mathbf{W}$ is used in the trapdoor construction, the relation returned by the solver is still with respect to $\mathbf{A}$ only. We can therefore recover a short lattice vector using the same procedure as in the warm-up.

It remains to explain how short the resulting lattice vector is. The bound from a single round may still depend on the initial basis $\mathbf{L}$. Our strategy is to use the recovered vectors to obtain a better basis for Gaussian sampling, then repeat the reduction with shorter Gaussian samples.

What controls the output length? The length of the recovered vector $\mathbf{v} \in \Lambda$ depends on the Gaussian width σ used to sample the columns of $\mathbf{V}$, and the bound γ on the norm of the solution provided by the sSIS solver. Two requirements determine how small we can choose the width σ to be. First, σ must always satisfy $\sigma \geq \sigma_0$, where σ_0 is an appropriate bound on the smoothing parameter of Λ . This ensures that $\mathbf{A}_0$ is statistically close to uniform. The value of σ_0 depends on the lattice, but not on the particular basis used to represent the lattice. Second, efficient Gaussian sampling requires a width that depends on the quality of the available basis.

In detail, we can choose

$$\sigma = \max \left\{ \sigma_0, \frac{\lambda \|\mathbf{L}\|}{q} \right\},$$

where λ is the security parameter, and $\mathbf{L}$ is the provided basis of Λ . Initially, we have no guarantee about the input basis $\mathbf{L}$, so a large value of $\|\mathbf{L}\|$ can force σ to be much larger than σ_0 .

The choices of p and q ensure, except with negligible probability,

$$\|\mathbf{v}\| \lesssim \gamma(\sigma + \sigma_0),$$

where we have suppressed some fixed factors depending on the matrix dimension in the above inequality to improve readability. The precise bound is given in Lemma 4.6.

This suggests how to improve the output bound: recover enough short vectors to construct a new basis with smaller norm, and then use this basis to sample with a smaller σ .

Iterating the improvement. Write $\mathbf{L}_0 = \mathbf{L}$, and let $\mathbf{L}_j$ be the sampling basis obtained after round j . In round j , use $\sigma_j = \max\{\sigma_0, \lambda \|\mathbf{L}_{j-1}\|/q\}$. The one-round analysis guarantees a non-negligible probability of obtaining a vector outside any fixed proper subspace. Repeating the one-round procedure polynomially many times with fresh randomness therefore yields n linearly independent short vectors with overwhelming probability. Together with the original basis, these vectors can be converted into a basis $\mathbf{L}_j$ of the full lattice. Consequently,

$$\|\mathbf{L}_j\| \lesssim \gamma(\sigma_j + \sigma_0) \lesssim 2\gamma\sigma_0 + \frac{\gamma\lambda}{q} \|\mathbf{L}_{j-1}\|.$$

Since the factors hidden by $\lesssim$ remain fixed across the iterations, we can choose q large enough that the contribution from $\|\mathbf{L}_{j-1}\|$ is at most $\frac{1}{4}\|\mathbf{L}_{j-1}\|$. Applying this inequality repeatedly gives

$$\|\mathbf{L}_j\| \lesssim \gamma\sigma_0 + 4^{-j}\|\mathbf{L}_0\|.$$

The contribution from the initial basis therefore decreases by a factor of four per round, while the remaining term stays fixed. This explains both the improvement and its limit: the bound approaches a constant multiple of $\gamma\sigma_0$. After logarithmically many rounds, the contribution from the initial basis becomes smaller than the fixed term $\gamma\sigma_0$. A final application of the reduction, using the improved sampling basis, then gives a nonzero vector satisfying

$$\|\mathbf{v}\| \lesssim \gamma\sigma_0.$$

The formal analysis of the complete reduction is given in Section 4.

2.2 Proof of Equivalence

We explain our equivalence proof for succinct and decomposed LWE. Throughout this subsection, set $m = \ell^2 n \lceil \log q \rceil$. Let σ_{dec} and χ_{dec} denote, respectively, the Gaussian widths of the matrix $\mathbf{R}$ and the error in the decomposed LWE instance. Similarly, let σ_{suc} and χ_{suc} denote the trapdoor and error widths of the succinct LWE instance. Once decomposed LWE is expressed as an instance of the general hinted LWE template (cf. Section 1.1), its relation to succinct LWE becomes particularly transparent. In both assumptions, the LWE challenge is associated with a matrix $\mathbf{A}_x \in \mathbb{Z}_q^{n \times m}$, $x \in \{\text{suc}, \text{dec}\}$, and the auxiliary information contains or implicitly determines a short gadget trapdoor $\mathbf{T}$ for the corresponding matrix $\mathbf{D}_\ell = [\mathbf{I}_\ell \otimes \mathbf{A}_x \mid \mathbf{W}]$, satisfying

$$\mathbf{D}_\ell \mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q.$$

For the decomposed input, we use the notation

$$\mathbf{W}^\star := [\mathbf{W}_1^\top \mid \cdots \mid \mathbf{W}_\ell^\top]^\top \in \mathbb{Z}_q^{\ell n \times k},$$

so the auxiliary matrix $\mathbf{W}$ in the common template above is $\mathbf{W}^\star$ in the reduction. For decomposed LWE, the matrices $\mathbf{A}_{\text{dec}}$ and $\mathbf{T}$ are efficiently obtained from $(\overline{\mathbf{W}}, \mathbf{R})$ as described in the introduction.

At the level of the underlying LWE challenge, the essential difference between the two assumptions is the distribution of the public matrix $\mathbf{A}_x$. In succinct LWE, $\mathbf{A}_{\text{suc}}$ is uniformly random, whereas in decomposed LWE the matrix $\mathbf{A}_{\text{dec}}$ has the structured form from Equation (2). We write $\mathbf{b}^\top$ for the response in the decomposed LWE instance. In the planted case,

$$\mathbf{b}^\top = \mathbf{s}^\top \mathbf{A}_{\text{dec}} + \mathbf{e}^\top \bmod q, \quad \mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m, \chi_{\text{dec}}},$$

whereas in the uniform case $\mathbf{b}$ is uniform over $\mathbb{Z}_q^m$.

The reduction must therefore transform $\mathbf{A}_{\text{dec}}$ into a matrix that is statistically close to uniform while simultaneously mapping the given decomposed LWE challenge to an LWE challenge for the transformed matrix. In the planted case, the resulting challenge must be a planted challenge for the transformed matrix; in the uniform case, it must remain statistically close to uniform. At the same time, the reduction must construct a valid short gadget trapdoor for the corresponding auxiliary matrix. This short trapdoor can then be used to sample a fresh trapdoor whose distribution is statistically close to the conditional Gaussian distribution required by succinct LWE.

A First Attempt. A naive first attempt is to sample an independent low-norm matrix $\mathbf{S} \in \mathbb{Z}^{m \times m}$ and set $\mathbf{A}^\star = \mathbf{A}_{\text{dec}} \cdot \mathbf{S}$. The structured leftover hash lemma of Afshar, Goyal, and Yadugiri [AGY26b, Lemma 3.6 and Claim A.3] shows that a transformation of this form makes the decomposed LWE matrix statistically close to uniform, even given its decomposition data.

To map the trapdoor, define

$$\bar{\mathbf{S}} = \text{diag}(\mathbf{I}_\ell \otimes \mathbf{S}, \mathbf{I}_k).$$

If $\mathbf{S}$ is invertible modulo q , then define

$$\tilde{\mathbf{D}}_\ell := \mathbf{D}_\ell \bar{\mathbf{S}} = [\mathbf{I}_\ell \otimes \mathbf{A}^\star \mid \mathbf{W}^\star],$$

and the natural candidate for the transformed trapdoor is $\bar{\mathbf{S}}^{-1} \mathbf{T}$.

The problem is that this approach gives no suitable bound on the inverse. Consequently, it does not guarantee that $\bar{\mathbf{S}}^{-1} \mathbf{T}$ is short enough to serve as a gadget trapdoor.

Rerandomization with a Short Inverse. The failure of the preceding attempt comes from sampling a short matrix S first and only afterward asking whether S^{-1} is also short. We instead restrict S to a structured family for which invertibility and the shortness of both S and S^{-1} hold by construction.

This simultaneous shortness requirement has appeared before in the lattice literature. We use the unitriangular rerandomization technique of Micciancio and Peikert [MP12]; a related transformation was subsequently used by Albrecht et al. [AFLN24, Theorem 3.5] to rerandomize instances of the Power-Ring-BASIS assumption [FMN24, Definition 3.2].

The key idea is to leave one block of the public matrix unchanged and use that block to mask all the remaining columns. Accordingly, write

$$\begin{aligned} \mathbf{A}_{\text{dec}} &= [\mathbf{A}_0 \mid \mathbf{A}_{\text{rest}}], \\ \mathbf{A}_0 &:= \mathbf{W}_1 \cdot \mathbf{R}_1 + \mathbf{G} \in \mathbb{Z}_q^{n \times n \lceil \log q \rceil}, \\ \mathbf{A}_{\text{rest}} &\in \mathbb{Z}_q^{n \times ((\ell^2 - 1)n \lceil \log q \rceil)}. \end{aligned}$$

For suitable parameters, every block $\mathbf{W}_i \mathbf{R}_j + \delta_{ij} \mathbf{G}$ is marginally statistically close to uniform, even jointly with $\overline{\mathbf{W}}$. This follows by applying the leftover hash lemma columnwise to $\mathbf{W}_i \mathbf{R}_j$; the fixed translation $\delta_{ij} \mathbf{G}$ does not affect the statistical distance. We single out the block $\mathbf{A}_0 = \mathbf{W}_1 \mathbf{R}_1 + \mathbf{G}$ and use it to mask all the remaining columns.

Independently, sample a short matrix

$$\mathbf{X} \sim \{-1, 0, 1\}^{n \lceil \log q \rceil \times ((\ell^2 - 1)n \lceil \log q \rceil)}$$

and define

$$\mathbf{S} := \begin{bmatrix} \mathbf{I}_{n \lceil \log q \rceil} & \mathbf{X} \\ \mathbf{0} & \mathbf{I}_{(\ell^2 - 1)n \lceil \log q \rceil} \end{bmatrix}, \quad \mathbf{S}^{-1} = \begin{bmatrix} \mathbf{I}_{n \lceil \log q \rceil} & -\mathbf{X} \\ \mathbf{0} & \mathbf{I}_{(\ell^2 - 1)n \lceil \log q \rceil} \end{bmatrix}.$$

Write $\mathbf{Q} := \mathbf{S}^{-1}$. Both $\mathbf{S}$ and $\mathbf{Q}$ have entries in $\{-1, 0, 1\}$. Then consider the rerandomized matrix

$$\mathbf{A}^\star := \mathbf{A}_{\text{dec}} \cdot \mathbf{S} = [\mathbf{A}_0 \mid \mathbf{A}_0 \cdot \mathbf{X} + \mathbf{A}_{\text{rest}}].$$

Although $\mathbf{A}_0$ is marginally close to uniform, it is correlated with $\mathbf{A}_{\text{rest}}$ through $\mathbf{W}_1$ and $\mathbf{R}_1$, and therefore cannot be replaced independently while leaving the remaining blocks fixed. Using conditional resampling followed by an application of the leftover hash lemma to $\mathbf{A}_0 \mathbf{X}$, Lemma 5.8 proves that

$$(\overline{\mathbf{W}}, \mathbf{A}_0, \mathbf{A}_0 \mathbf{X} + \mathbf{A}_{\text{rest}}) \approx_s (\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{U}_{\text{rest}}),$$

where $\mathbf{U}_0$ and $\mathbf{U}_{\text{rest}}$ are independent uniform matrices.

We next transport the decomposed LWE trapdoor through this transformation. Define $\overline{\mathbf{S}} := \text{diag}(\mathbf{I}_\ell \otimes \mathbf{S}, \mathbf{I}_k)$. Then

$$\widetilde{\mathbf{D}}_\ell := [\mathbf{I}_\ell \otimes \mathbf{A}^\star \mid \mathbf{W}^\star] = \mathbf{D}_\ell \overline{\mathbf{S}}.$$

Consequently, the matrix $\widehat{\mathbf{T}} := \overline{\mathbf{S}}^{-1} \mathbf{T}$ satisfies

$$\widetilde{\mathbf{D}}_\ell \widehat{\mathbf{T}} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q.$$

To see its norm bound directly, partition $\mathbf{Q} = [\mathbf{Q}_1 \mid \cdots \mid \mathbf{Q}_\ell]$, where each $\mathbf{Q}_i$ has m rows and $\ell n \lceil \log q \rceil$ columns. The explicit transformed trapdoor is

$$\hat{\mathbf{T}} = [\mathbf{Q}_1^\top \mid \cdots \mid \mathbf{Q}_\ell^\top \mid -\mathbf{R}^\top]^\top, \quad \|\hat{\mathbf{T}}\| \leq \max\{1, \|\mathbf{R}\|\}.$$

The proof aborts if $\|\mathbf{R}\| > 2\sigma_{\text{dec}} \sqrt{k\ell n \lceil \log q \rceil}$; this event has negligible probability. Thus, $\hat{\mathbf{T}}$ is a valid short gadget trapdoor for $\mathbf{D}_\ell$. It is not, however, distributed according to the conditional Gaussian distribution required by succinct LWE. The theorem's lower bound on σ_{suc} instead gives the Gaussian width required by the preimage sampler, with row dimension ℓn and column dimension $\ell m + k$. Writing SamplePre for this standard Gaussian preimage-sampling procedure, the reduction computes

$$\mathbf{T}^* \leftarrow \text{SamplePre}(\tilde{\mathbf{D}}_\ell, \hat{\mathbf{T}}, \mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\text{suc}})$$

which satisfies

$$\mathbf{T}^* \approx_s \tilde{\mathbf{D}}_\ell^{-1} (\mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\text{suc}}).$$

The freshly sampled matrix $\mathbf{T}^*$ is the trapdoor given to the succinct LWE distinguisher. Finally, the reduction applies the rerandomization algorithm of [KY16, Lemma 1] to $\mathbf{b}^\top$ using $\mathbf{S}$, which ensures that the modulus stays polynomial throughout. Denote the resulting vector by $\mathbf{c}^{*\top}$. In the planted case, $\mathbf{c}^{*\top}$ is statistically close to $\mathbf{s}^\top \mathbf{A}^* + \tilde{\mathbf{e}}^\top$, where $\tilde{\mathbf{e}} \sim \mathcal{D}_{\mathbb{Z}^m, \chi_{\text{suc}}}$. In the uniform case, $\mathbf{b}$ is uniform over $\mathbb{Z}_q^m$ and independent of $\mathbf{S}$. Since $\mathbf{S}$ is invertible modulo q , we obtain that $\mathbf{c}^{*\top} \approx_s \text{Unif}(\mathbb{Z}_q^m)$.

2.3 Search-to-Decision Reduction

Our search-to-decision reduction follows Regev's method of testing every possible value of each secret coordinate. We first recall this method for ordinary LWE and then apply the same transformation to decomposed LWE. Throughout this subsection, q is an odd prime and polynomially bounded in the security parameter λ , and all matrix and vector arithmetic is modulo q .

Revisiting Regev's Reduction. An ordinary LWE instance is a pair $(\mathbf{A}, \mathbf{b}^\top)$, where $\mathbf{A} \sim \mathbb{Z}_q^{n \times m}$ and $\mathbf{b}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$ for an unknown secret $\mathbf{s} \in \mathbb{Z}_q^n$ and an independent error vector $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m, \chi}$. For now, suppose that we have a decision algorithm that accepts LWE instances and rejects uniform ones, with negligible error for every fixed secret. Regev shows how to obtain such an algorithm from a distinguisher by randomizing the secret and repeating on fresh samples [Reg09, Lemmas 4.1 and 4.2].

Fix a coordinate $i \in [n]$ and a proposed guess $t \in \mathbb{Z}_q$ for s_i . Let $\mathbf{e}_i$ be the i -th unit vector in $\mathbb{Z}_q^n$. On a fresh collection of LWE samples $(\mathbf{A}, \mathbf{b}^\top)$ with secret $\mathbf{s}$, perform the following test.

1. Sample $\mathbf{r} \sim \mathbb{Z}_q^m$ uniformly and independently. Set $\mathbf{A}' := \mathbf{A} + \mathbf{e}_i \mathbf{r}^\top$. Thus, the matrix added to $\mathbf{A}$ has $\mathbf{r}^\top$ in row i and zeros in every other row.
2. Compute $\mathbf{b}'^\top := \mathbf{b}^\top + t \mathbf{r}^\top$.
3. Run the decision algorithm on $(\mathbf{A}', \mathbf{b}'^\top)$ and accept the guess t if the algorithm accepts.

The reason this tests s_i is the identity

$$\mathbf{b}'^\top = \mathbf{s}^\top \mathbf{A}' + \mathbf{e}^\top + (t - s_i) \mathbf{r}^\top.$$

The reduction computes the left-hand side using only the given response, the guess, and its own randomness. If $t = s_i$, the additional term vanishes and the output is exactly an LWE instance. If $t \neq s_i$, then $t - s_i$ is a nonzero field element, so multiplication by $t - s_i$ maps a uniform $\mathbf{r}$ to a uniform vector. Moreover, the map $(\mathbf{A}, \mathbf{r}) \mapsto (\mathbf{A} + \varepsilon_i \mathbf{r}^\top, \mathbf{r})$ is a bijection, so $\mathbf{r}$ remains uniform and independent of $\mathbf{A}'$. Consequently, $\mathbf{b}'$ is uniform even given $\mathbf{A}'$. Testing all $t \in \mathbb{Z}_q$ for every $i \in [n]$ therefore identifies the secret: return the unique accepted value for each coordinate, or report failure if any coordinate has no accepted value or more than one.

Adapting to Decomposed LWE. Let $\mathcal{A}$ distinguish decisional ℓ -decomposed LWE with advantage at least ε , and set

$$N := \left\lceil \frac{32(\lambda + 2) \ln 2}{\varepsilon^2} \right\rceil, \quad L := N\ell.$$

Restricting an L -decomposed input to N disjoint groups of ℓ block indices produces N independent uniform inputs in the uniform case and N conditionally independent planted inputs with a shared secret in the LWE case. The usual secret randomization and amplification then give a decision algorithm $\mathcal{B}$ that accepts L -decomposed LWE instances and rejects uniform instances with error at most $2^{-\lambda}$ for every fixed secret.

A search L -decomposed LWE instance gives $(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top)$ and determines the matrix

$$\mathbf{A} := \overline{\mathbf{W}}(\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}, \quad \mathbf{b}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top.$$

The matrix $\overline{\mathbf{W}} \sim \mathbb{Z}_q^{n \times Lk}$ is uniform, the entries of $\mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times Ln \lceil \log q \rceil}, \sigma}$ are independent discrete Gaussians, and $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{L^2 n \lceil \log q \rceil}, \chi}$ is independent of both matrices and the secret.

To mimic Regev's coordinate-guessing transformation while preserving the decomposed form of $\mathbf{A}$, we add a random row to $\overline{\mathbf{W}}$. For a coordinate $i \in [n]$ and a guess $t \in \mathbb{Z}_q$, perform the following test.

1. Independently sample uniform vectors $\mathbf{p} \sim \mathbb{Z}_q^{Lk}$ and $\mathbf{z} \sim \mathbb{Z}_q^n$. Set $\overline{\mathbf{V}} := \overline{\mathbf{W}} + \varepsilon_i \mathbf{p}^\top$ and compute its corresponding LWE matrix

$$\mathbf{A}' := \overline{\mathbf{V}}(\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}.$$

2. Compute the response

$$\mathbf{c}^\top := \mathbf{b}^\top + t \mathbf{p}^\top (\mathbf{I}_L \otimes \mathbf{R}) + \mathbf{z}^\top \mathbf{A}'.$$

The last term changes the secret from $\mathbf{s}$ to the independently uniform secret $\mathbf{s} + \mathbf{z}$.

3. Give $(\overline{\mathbf{V}}, \mathbf{R}, \mathbf{c}^\top)$ to $\mathcal{B}$, and accept t if $\mathcal{B}$ accepts.

Substituting for the original response gives exactly the same calculation as before:

$$\mathbf{c}^\top = (\mathbf{s} + \mathbf{z})^\top \mathbf{A}' + \mathbf{e}^\top + (t - s_i) \mathbf{p}^\top (\mathbf{I}_L \otimes \mathbf{R}).$$

In particular, if $t = s_i$, the additional term vanishes. The output then has exactly the decomposed LWE distribution: $\overline{\mathbf{V}}$ is uniform, $\mathbf{R}$ and $\mathbf{e}$ retain their Gaussian distributions, and the new secret $\mathbf{s} + \mathbf{z}$ is uniform and independent of these variables.

For a wrong guess, let $\mathbf{v} := (t - s_i)\mathbf{p}$, which is uniform because q is prime. The same bijection as before shows that $\mathbf{v}$ is independent of $\bar{\mathbf{V}}$, and it is also independent of $\mathbf{R}$, $\mathbf{e}$, and $\mathbf{s} + \mathbf{z}$. However, its product with the public matrix $\mathbf{I}_L \otimes \mathbf{R}$ need not be uniform. It remains to show that

$$(\mathbf{R}, \mathbf{v}^\top(\mathbf{I}_L \otimes \mathbf{R}) + \mathbf{e}^\top) \approx_c (\mathbf{R}, \mathbf{w}^\top),$$

where $\mathbf{w}$ is uniform in $\mathbb{Z}_q^{L^2 n \lceil \log q \rceil}$. Write $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_L)$ and split $\mathbf{e} = (\mathbf{e}_1, \dots, \mathbf{e}_L)$ into corresponding blocks. This follows from an L -step hybrid: at step $a \in [L]$, we replace $\mathbf{v}_a^\top \mathbf{R} + \mathbf{e}_a^\top$ with an independent uniform vector. The indistinguishability of this replacement follows from LWE with Gaussian matrix $\mathbf{R}$. Under the conditions of Theorem 6.1, the hardness of this variant follows from standard decisional LWE by Theorem 3.13.

Hence, for a wrong guess, the input to $\mathcal{B}$ is computationally indistinguishable from a uniform instance, whereas the correct guess produces an L -decomposed LWE instance. A union bound over the nq guesses shows that the entire secret is recovered except with negligible probability.

2.4 Quantum Reduction from Succinct SIS to Succinct LWE

We first recall how Regev's quantum method [Reg09], more explicitly presented in [CLZ22], uses a search LWE solver to solve SIS. Given a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ of full row rank, our goal is to find a short nonzero integer vector $\mathbf{x}$ with $\mathbf{A}\mathbf{x} \equiv \mathbf{0} \pmod{q}$. For this overview, we assume perfect secret recovery on the prepared superposition and omit negligible errors from Gaussian tails and state preparation; the formal proof in the technical sections handles a solver with any inverse-polynomial success probability. The reduction goes as follows:

1. Prepare a uniform superposition of $\mathbf{s} \in \mathbb{Z}_q^n$ and a Gaussian superposition of $\mathbf{e} \in \mathbb{Z}_q^m$.

$$q^{-n/2} \sum_{\substack{\mathbf{s} \in \mathbb{Z}_q^n \\ \mathbf{e} \in \mathbb{Z}_q^m}} f(\mathbf{e}) |\mathbf{A}^\top \mathbf{s} + \mathbf{e}\rangle |\mathbf{e}\rangle.$$

The state has nonnegative amplitudes $f(\mathbf{e}) \propto \exp(-\pi \|\mathbf{e}\|_2^2 / (2\chi^2))$, normalized so that $\sum_{\mathbf{e} \in \mathbb{Z}_q^m} f(\mathbf{e})^2 = 1$, and can be prepared efficiently using Grover–Rudolph [GR02]. After computing $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e}$, erase the initial secret register by reversibly solving $\mathbf{A}^\top \mathbf{s} = \mathbf{b} - \mathbf{e}$; full row rank makes this solution unique.

2. Run the LWE solver on $(\mathbf{A}, \mathbf{b})$ without measuring, preserving $\mathbf{b}$ and retaining all intermediate computations. Its answer is $\mathbf{s}$, so subtracting $\mathbf{b} - \mathbf{A}^\top \mathbf{s} = \mathbf{e}$ from the noise register resets that register to zero. Run the solver circuit in reverse to erase its answer and work registers. The remaining state is

$$q^{-n/2} \sum_{\substack{\mathbf{s} \in \mathbb{Z}_q^n \\ \mathbf{e} \in \mathbb{Z}_q^m}} f(\mathbf{e}) |\mathbf{A}^\top \mathbf{s} + \mathbf{e}\rangle.$$

3. Apply the quantum Fourier transform over $\mathbb{Z}_q^m$ and measure. The state before measurement is

$$q^{n/2} \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^m \\ \mathbf{A}\mathbf{x} \equiv \mathbf{0} \pmod{q}}} \hat{f}(\mathbf{x}) |\mathbf{x}\rangle, \quad \text{where } \hat{f}(\mathbf{x}) = q^{-m/2} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} f(\mathbf{e}) \exp(2\pi i \langle \mathbf{e}, \mathbf{x} \rangle / q).$$

Indeed, summing the Fourier phases over $\mathbf{s}$ gives q^n when $\mathbf{Ax} \equiv \mathbf{0} \pmod{q}$ and zero otherwise. Return the measured integer vector if it is nonzero and has Euclidean norm at most $q\sqrt{m}/(2\chi)$; otherwise, return $\perp$.

The distribution with probabilities $|\widehat{f}(\mathbf{x})|^2$ is negligibly close to $\mathcal{D}_{\mathbb{Z}^m, q/(2\chi)} \pmod{q}$ by Lemma 7.2. The measured state also restricts the support to $\mathbf{Ax} \equiv \mathbf{0} \pmod{q}$ and multiplies the probabilities by q^n . To control this restriction, average over the uniformly sampled matrix of full row rank. For every nonzero $\mathbf{x} \in \mathbb{Z}_q^m$,

$$\Pr_{\mathbf{A}}[\mathbf{Ax} = \mathbf{0}] = \frac{q^{m-n} - 1}{q^m - 1} \leq q^{-n}.$$

Together with the bounds on the Fourier tail and $q^n |\widehat{f}(\mathbf{0})|^2$ from Lemma 7.2, this makes the average probability of a zero or overly long output negligible.

For succinct SIS, receive $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ and give $(\mathbf{A}, \mathbf{b}, \mathbf{W}, \mathbf{T})$ to the search succinct LWE solver in Step 2. The auxiliary matrices stay fixed during the coherent computation, and the secret and error are prepared independently of them. The algebraic Fourier calculation holds for each fixed public tuple; the probability bound averages over its uniform marginal over full row rank matrices $\mathbf{A}$. In the formal reduction, an imperfect solver is run coherently and in reverse, and the error and work registers are tested for zero before the Fourier step. The overlap analysis in Theorem 7.1 gives success probability at least $\delta^2 - \text{negl}(\lambda)$ when the search succinct LWE solver succeeds with probability δ , with the SIS bound $\gamma \geq q\sqrt{m}/(2\chi)$.

3 Preliminaries

Notation. We write $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ and $[d] = \{1, \dots, d\}$, and use δ_{ij} for the Kronecker delta, equal to 1 when $i = j$ and to 0 otherwise. Bold lowercase letters denote column vectors (e.g., $\mathbf{r}$), and bold uppercase letters denote matrices (e.g., $\mathbf{R}$). Unless stated otherwise, $\log$ denotes $\log_2$.

For integral vectors and matrices, $\|\mathbf{r}\|$ and $\|\mathbf{R}\|$ denote the maximum absolute value of their entries. In the technical overview, we occasionally omit norm subscripts for readability: vector lengths are Euclidean, while matrix norms may denote either the maximum Euclidean column length or the maximum absolute entry. The formal sections specify the norm being used, with unsubscripted norms following the maximum-entry convention above.

For integers $a \geq 1, b \geq 2$, we set $\mathbf{g}_b^\top = (2^0, 2^1, \dots, 2^{\lceil \log b \rceil - 1})$ and $\mathbf{G}_{a,b} = \mathbf{I}_a \otimes \mathbf{g}_b^\top$. We write $\mathbf{G} = \mathbf{G}_{n,q}$ for the usual gadget matrix.

The security parameter is denoted by λ . A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if $|f(\lambda)| = 2^{-\omega(\log \lambda)}$, and polynomially bounded if $|f(\lambda)| = 2^{O(\log \lambda)} = \lambda^{O(1)}$. We write $\text{negl}(\lambda)$ for a negligible function.

We write $x \sim \mathcal{D}$ to mean that x is sampled according to the distribution $\mathcal{D}$; when the right-hand side is a finite set, the sampling is uniform. We denote the uniform distribution over a finite set Ω by $\text{Unif}(\Omega)$, and the discrete Gaussian over $\mathbb{Z}^d$ with parameter s by $\mathcal{D}_{\mathbb{Z}^d, s}$. For a distribution $\mathcal{X}$ over a finite set Ω , its min-entropy is

$$H_\infty(\mathcal{X}) := -\log \left(\max_{x \in \Omega} \Pr_{x' \sim \mathcal{X}}[x' = x] \right).$$

For two distributions $\mathcal{X}$ and $\mathcal{Y}$ over the same countable set Ω , their statistical distance is

$$\Delta(\mathcal{X}, \mathcal{Y}) := \frac{1}{2} \sum_{x \in \Omega} |\Pr[\mathcal{X} = x] - \Pr[\mathcal{Y} = x]|.$$

We use the same notation for random variables. We write $\mathcal{X} \approx_s \mathcal{Y}$ when $\Delta(\mathcal{X}, \mathcal{Y}) = \text{negl}(\lambda)$, in which case they are statistically close, and $\mathcal{X} \approx_c \mathcal{Y}$ if the corresponding distribution ensembles are computationally indistinguishable; namely, if for every probabilistic polynomial-time algorithm $\mathcal{A}$,

$$\left| \Pr[\mathcal{A}(1^\lambda, \mathcal{X}) = 1] - \Pr[\mathcal{A}(1^\lambda, \mathcal{Y}) = 1] \right| = \text{negl}(\lambda).$$

3.1 Statistical Distance Facts

Fact 3.1 (Data processing for statistical distance). *Let A and B be random variables over the same countable domain, and let K be any randomized map defined on that domain. Then*

$$\Delta(K(A), K(B)) \leq \Delta(A, B).$$

Lemma 3.2 (Leftover Hash Lemma [HILL99]). *Let q be prime, $\mathcal{X}$ be a distribution over $\mathbb{Z}_q^m$ and let the samples $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $\mathbf{x} \sim \mathcal{X}$, $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^n$ be mutually independent. Then*

$$\Delta((\mathbf{A}, \mathbf{A}\mathbf{x}), (\mathbf{A}, \mathbf{u})) \leq \frac{1}{2} \sqrt{q^n \cdot 2^{-H_\infty(\mathcal{X})}}.$$

Consequently, if $H_\infty(\mathcal{X}) \geq n \log q + 2\lambda - 2$, then $(\mathbf{A}, \mathbf{A}\mathbf{x}) \approx_s (\mathbf{A}, \mathbf{u})$, with statistical distance at most $2^{-\lambda}$. For an integer-valued source, apply the lemma to its reduction modulo q ; the entropy in the bound is the entropy of that reduced source.

3.2 Lattice Background

We denote by $\mathcal{D}_{\mathbb{Z}^m, \chi}$ the (centered) discrete Gaussian distribution over $\mathbb{Z}^m$ with parameter χ , i.e., the distribution over $\mathbb{Z}^m$ where for all $\mathbf{x}$, $\Pr[\mathbf{x}] \propto e^{-\pi \cdot (x_1^2 + \dots + x_m^2) / \chi^2}$.

We recall several useful facts from the literature that we will use in the technical part of the work.

Lemma 3.3 (Norm of Gaussian, adapted from [GPV08, Lemma 4.2, ePrint version]). *For all $\chi \geq 1$ and $k \geq 2$,*

$$\Pr[|x| \geq \chi \sqrt{k} \mid x \sim \mathcal{D}_{\mathbb{Z}, \chi}] \leq 2^{-k}.$$

Therefore, if $\mathbf{M} \sim \mathcal{D}_{\mathbb{Z}^{m \times m'}, \chi}$, taking $k = \lambda + \log(mm')$ plus a union bound gives

$$\|\mathbf{M}\| \leq \chi \sqrt{\lambda + \log(m \cdot m')}$$

except with probability at most $2^{-\lambda}$.

Lemma 3.4 (Noise rerandomization, adapted from [KY16, Lemma 1]). Let $m, t, q \in \mathbb{N}$, let $\chi = \omega(\sqrt{\log(m+t)})$, and let $\mathbf{V} \in \mathbb{Z}^{m \times t}$. Write $\|\mathbf{V}\|_{2 \rightarrow 2}$ for its spectral norm. For every $\chi' > 2\chi\|\mathbf{V}\|_{2 \rightarrow 2}$, there exists a probabilistic polynomial-time algorithm ReRand such that, for every $\mathbf{a} \in \mathbb{Z}_q^m$ and independent $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m, \chi}$, given only $\mathbf{a}^\top + \mathbf{e}^\top \bmod q$, it outputs

$$\text{ReRand}\left(\mathbf{V}, \mathbf{a}^\top + \mathbf{e}^\top, \chi, \frac{\chi'}{2\chi}\right) = \mathbf{a}^\top \mathbf{V} + (\mathbf{e}')^\top \bmod q,$$

where $\mathbf{e}'$ is statistically close to $\mathcal{D}_{\mathbb{Z}^t, \chi'}$. The algorithm multiplies its input by $\mathbf{V}$ and adds noise whose distribution depends only on $(\mathbf{V}, \chi, \chi')$. If $m, t = \text{poly}(\lambda)$ and $\chi \geq \sqrt{\lambda}$, its statistical error can be taken to be $2^{-\Omega(\lambda)}$.

For $t = m$, the sufficient condition $\chi' > 2\chi m \|\mathbf{V}\|$ follows from $\|\mathbf{V}\|_{2 \rightarrow 2} \leq m \|\mathbf{V}\|$, where the unsubscripted norm is the maximum absolute entry. If $\mathbf{V}$ is also invertible modulo q and $\mathbf{u} \sim \mathbb{Z}_q^m$ is uniform, then

$$\text{ReRand}\left(\mathbf{V}, \mathbf{u}^\top, \chi, \frac{\chi'}{2\chi}\right) \approx_s \text{Unif}(\mathbb{Z}_q^m).$$

Justification of the quantitative error bound. The proof of [KY16, Lemma 1, Appendix B.1] uses two Gaussian smoothing steps: adding a continuous Gaussian to the input noise, then rounding a continuous Gaussian to the integer lattice. For a nonzero integer matrix $\mathbf{V}$, $\chi'/(2\chi) > \|\mathbf{V}\|_{2 \rightarrow 2} \geq 1$, so both steps use smoothing widths at least $\chi/\sqrt{2}$. For $u \geq \chi/\sqrt{2}$ and $d \in \{m, t\}$, the nonzero terms in the Poisson summation bound satisfy

$$\sum_{\mathbf{z} \in \mathbb{Z}^d \setminus \{0\}} e^{-\pi u^2 \|\mathbf{z}\|_2^2} = \left(1 + 2 \sum_{j \geq 1} e^{-\pi u^2 j^2}\right)^d - 1 = O\left(de^{-\pi \chi^2/2}\right) = 2^{-\Omega(\lambda)}.$$

Thus each smoothing step has exponentially small statistical error, and data processing and the triangle inequality preserve this bound. For $\mathbf{V} = \mathbf{0}$, sample the target Gaussian directly. The Gaussian samplers and numerical approximations can be implemented with $2^{-\Omega(\lambda)}$ additional error using polynomial precision. $\square$

Min-Entropy of a Discrete Gaussian. The following fact follows from standard Gaussian estimates. All techniques below can be found in Lemmas 2.8–2.10 in [MR07].

Fact 3.5 (See [MR07]). Let $R \leftarrow \mathcal{D}_{\mathbb{Z}^k, \sigma}$ and let $\mathcal{R}$ denote the distribution $R \bmod q$ over $\mathbb{Z}_q^k$. If $q/\sigma = \omega(\sqrt{\log \lambda})$ and $k = \text{poly}(\lambda)$, then $H_\infty(\mathcal{R}) \geq k \log_2 \sigma - \text{negl}(\lambda)$.

Proof. For $\sigma > 0$ and a discrete set X in a Euclidean space, write $\rho_\sigma(X) := \sum_{\mathbf{x} \in X} e^{-\pi \|\mathbf{x}\|^2 / \sigma^2}$. For every $\mathbf{a} \in \mathbb{Z}_q^k$, since the centered Gaussian mass over a lattice is maximal [MR07, Lemma 2.9] (applied to $q\mathbb{Z}^k$),

$$\Pr[\mathcal{R} = \mathbf{a}] = \frac{\rho_\sigma(\mathbf{a} + q\mathbb{Z}^k)}{\rho_\sigma(\mathbb{Z}^k)} \leq \frac{\rho_\sigma(q\mathbb{Z}^k)}{\rho_\sigma(\mathbb{Z}^k)} = \left(\frac{\rho_{\sigma/q}(\mathbb{Z})}{\rho_\sigma(\mathbb{Z})}\right)^k.$$

By Poisson summation [MR07, Lemma 2.8], $\rho_\sigma(\mathbb{Z}) = \sigma \rho_{1/\sigma}(\mathbb{Z}) \geq \sigma$, while a geometric series bound gives $\rho_{\sigma/q}(\mathbb{Z}) \leq 1 + 3e^{-\pi(q/\sigma)^2} = 1 + \varepsilon(\lambda)$ where $\varepsilon(\lambda) = \lambda^{-\omega(1)}$ using $q/\sigma = \omega(\sqrt{\log \lambda})$. Hence

$$\max_{\mathbf{a} \in \mathbb{Z}_q^k} \Pr[\mathcal{R} = \mathbf{a}] \leq \sigma^{-k} (1 + \varepsilon(\lambda))^k.$$

By definition of min-entropy,

$$\begin{aligned}
H_\infty(\mathcal{R}) &= -\log_2 \max_{\mathbf{a} \in \mathbb{Z}_q^k} \Pr[\mathcal{R} = \mathbf{a}] \\
&\geq k \log_2 \sigma - k \log_2 (1 + \varepsilon(\lambda)) \\
&\geq k \log_2 \sigma - \frac{k \varepsilon(\lambda)}{\ln 2},
\end{aligned}$$

and since $k = \text{poly}(\lambda)$ and $\varepsilon(\lambda) = \lambda^{-\omega(1)}$, the loss term $k \varepsilon(\lambda) / \ln 2$ is $\text{negl}(\lambda)$, which proves the claim. $\square$

Trapdoor and Preimage Sampling

Our reduction will use gadget trapdoors [MP12], which build on a long sequence of works on constructing lattice trapdoors [Ajt96, GPV08, AP09, ABB10a, ABB10b, CHKP10]. We use the black-box formulation of [WWW25, Theorem 3.9, full version] throughout the manuscript.

Let $n, q \in \mathbb{N}$. Define the gadget vector and gadget matrix by

$$\mathbf{g}^\top := (1, 2, 4, \dots, 2^{\lceil \log q \rceil - 1}) \in \mathbb{Z}_q^{1 \times \lceil \log q \rceil}, \quad \mathbf{G} := \mathbf{I}_n \otimes \mathbf{g}^\top \in \mathbb{Z}_q^{n \times (n \cdot \lceil \log q \rceil)}.$$

For any $t \in \mathbb{N}$, the bit decomposition map $\mathbf{G}^{-1} : \mathbb{Z}_q^{n \times t} \rightarrow \{0, 1\}^{(n \cdot \lceil \log q \rceil) \times t}$ satisfies $\mathbf{G} \cdot \mathbf{G}^{-1}(\mathbf{V}) = \mathbf{V} \bmod q$ for every $\mathbf{V} \in \mathbb{Z}_q^{n \times t}$.

Let $m \in \mathbb{N}$, $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and consider a target vector $\mathbf{v} \in \mathbb{Z}_q^n$. We write $\mathbf{A}^{-1}(\mathbf{v}, \tau)$ for the discrete Gaussian distribution over $\mathbf{x} \in \mathbb{Z}^m$, with parameter τ , conditioned on $\mathbf{A} \cdot \mathbf{x} = \mathbf{v} \bmod q$. Equivalently,

$$\mathbf{x} \sim \mathbf{A}^{-1}(\mathbf{v}, \tau) \iff \mathbf{x} \sim \mathcal{D}_{\mathbb{Z}^m, \tau} \mid \mathbf{A} \cdot \mathbf{x} = \mathbf{v} \bmod q. \quad (3)$$

We extend $\mathbf{A}^{-1}$ to operate on matrices by applying $\mathbf{A}^{-1}$ columnwise.

Lemma 3.6 (Gadget trapdoors and preimage sampling, [WWW25, Theorem 3.9, full version]). *Let $n \geq 2$, $q \geq 2$, and $m \geq 3n \lceil \log q \rceil$ be integers. There are probabilistic polynomial-time algorithms (TrapGen, SamplePre) with the following properties.*

- TrapGen($1^n, q, m$) outputs $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{T} \in \mathbb{Z}^{m \times n \lceil \log q \rceil}$ such that $\mathbf{AT} = \mathbf{G} \bmod q$ and $\|\mathbf{T}\| = 1$. The marginal distribution of $\mathbf{A}$ is within distance 2^{-n} of uniform.
- For every $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, every $\mathbf{T} \in \mathbb{Z}^{m \times n \lceil \log q \rceil}$ satisfying $\mathbf{AT} = \mathbf{G} \bmod q$, every $\mathbf{v} \in \mathbb{Z}_q^n$, and every $s \geq m \|\mathbf{T}\| \log n$, the output $\mathbf{x} \leftarrow \text{SamplePre}(\mathbf{A}, \mathbf{T}, \mathbf{v}, s) \in \mathbb{Z}^m$ satisfies $\mathbf{Ax} = \mathbf{v} \bmod q$ and

$$\Delta(\text{SamplePre}(\mathbf{A}, \mathbf{T}, \mathbf{v}, s), \mathbf{A}^{-1}(\mathbf{v}, s)) \leq 2^{-n}. \quad (4)$$

Here $\|\mathbf{T}\| = \max_{i,j} |T_{i,j}|$ is the maximum absolute entry. The sampling guarantee holds for every such trapdoor; it need not have been generated by TrapGen.

The cited theorem is the formulation of the Micciancio–Peikert trapdoor tools used in [WWW25]. Its padded gadget convention is equivalent to the one above: append zero columns to $\mathbf{T}$ to obtain a trapdoor for $[\mathbf{G} \mid \mathbf{0}]$, without changing its norm or the preimage dimension. We retain the integer sample before reduction modulo q , as in the conditional Gaussian definition Eq. (3). The identity $\mathbf{AT} = \mathbf{G} \bmod q$ already ensures that every target has a preimage.

Matrix Targets and Retained Randomness. We apply SamplePre independently column-wise for a target $\mathbf{U} \in \mathbb{Z}_q^{n \times t}$. For a fixed width s , let Ξ be any retained state determining $\mathbf{A}, \mathbf{T}, \mathbf{U}$, such that the hypotheses of Lemma 3.6 hold for every state in its support. Let $\mathbf{X}$ be the sampler output using fresh coins, and, conditional on Ξ , sample $\mathbf{X}_{\text{can}} \sim \mathbf{A}^{-1}(\mathbf{U}, s)$ with independent columns. A hybrid over the t columns followed by averaging over Ξ gives

$$\Delta((\Xi, \mathbf{X}), (\Xi, \mathbf{X}_{\text{can}})) \leq \mathbb{E}_{\xi \sim \Xi}[t2^{-n}] = t2^{-n}. \quad (5)$$

The canonical conditional distribution depends only on $(\mathbf{A}, \mathbf{U}, s)$. In particular, this replacement removes any further dependence on the trapdoor construction while retaining all other state. For $t = \text{poly}(\lambda)$ and $n = \omega(\log \lambda)$, the total error is negligible in λ .

3.3 Computational Assumptions

Definition 3.7 (ℓ -Succinct LWE [Wee25]). Fix a security parameter λ , integers $n, m, \ell \geq 1, k \geq 0, q \geq 2$, and Gaussian widths $\chi, \sigma > 0$. Sample $\mathbf{A} \sim \mathbb{Z}_q^{n \times m}$ and $\mathbf{W} \sim \mathbb{Z}_q^{\ell n \times k}$ independently, and put

$$\mathbf{D} := [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}], \quad \mathbf{T} \sim \mathbf{D}^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma) \in \mathbb{Z}^{(\ell m + k) \times \ell n \lceil \log q \rceil},$$

with independent Gaussian columns as in Eq. (3). If a target has no preimage, both experiments below output a fixed failure symbol $\perp$. Independently sample $\mathbf{s} \sim \mathbb{Z}_q^n$, $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m, \chi}$, and $\mathbf{u} \sim \mathbb{Z}_q^m$, and set $\mathbf{b}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top \bmod q$. The decision problem is to distinguish

$$(\mathbf{A}, \mathbf{b}^\top, \mathbf{W}, \mathbf{T}) \quad \text{from} \quad (\mathbf{A}, \mathbf{u}^\top, \mathbf{W}, \mathbf{T}).$$

The decision ℓ -succinct LWE assumption states that these ensembles are computationally indistinguishable.

The search version is to recover $\mathbf{s}$ from $(\mathbf{A}, \mathbf{b}^\top, \mathbf{W}, \mathbf{T})$ in the planted experiment. For the search problem used in this paper, q is an odd prime, $m \geq n$, and the sampling of $\mathbf{A}$ is additionally conditioned on $\text{rank}(\mathbf{A}) = n$. We denote this problem by $\text{sLWE}_{n,m,q,\ell,k,s,\chi}$, where $s = \sigma$ is the Gaussian width used to sample $\mathbf{T}$.

For prime q , the probability that a uniform $n \times m$ matrix fails to have full row rank is at most q^{n-m} . Full row rank makes all gadget targets reachable, so conditioning changes the public sampling experiment by at most this distance. The loss is negligible whenever $(m - n) \log q = \omega(\log \lambda)$, as in the reductions using this comparison. This also bounds the probability of the failure symbol in the unconditioned decision experiment.

Definition 3.8 (ℓ -Decomposed LWE [AMR25]). Fix a security parameter λ , integers $n, \ell, k \geq 1, q \geq 2$, and Gaussian widths $\chi, \sigma > 0$. The gadget matrix $\mathbf{G}$ has width $n \lceil \log q \rceil$. Independently sample

$$\overline{\mathbf{W}} \sim \mathbb{Z}_q^{n \times \ell k}, \quad \mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times \ell n \lceil \log q \rceil}, \sigma},$$

and define the structured public matrix

$$\mathbf{A}_{\text{dec}} := \overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G} \in \mathbb{Z}_q^{n \times \ell^2 n \lceil \log q \rceil}.$$

Independently sample $\mathbf{s} \sim \mathbb{Z}_q^n$, $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{\ell^2 n \lceil \log q \rceil}, \chi}$, and $\mathbf{u} \sim \mathbb{Z}_q^{\ell^2 n \lceil \log q \rceil}$, and set $\mathbf{b}^\top = \mathbf{s}^\top \mathbf{A}_{\text{dec}} + \mathbf{e}^\top \bmod q$. The decision problem is to distinguish

$$(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top) \quad \text{from} \quad (\overline{\mathbf{W}}, \mathbf{R}, \mathbf{u}^\top).$$

The decision ℓ -decomposed LWE assumption states that these ensembles are computationally indistinguishable.

The search version is to recover $\mathbf{s}$ from the planted sample $(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top)$ with the same parameters (n, k, χ, σ, q) and decomposition degree ℓ .

The SIS problems below are search problems: the task is to find a short nonzero integer kernel vector. They use the public sampling experiments of the corresponding LWE problems, without an LWE challenge.

Definition 3.9 (ℓ -Succinct SIS [Wee24]). Let q be an odd prime, let $m \geq n$, and let $s, \gamma > 0$. Sample $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ as in the search version of Definition 3.7, with $\sigma = s$; in particular, $\mathbf{A}$ is uniform conditioned on full row rank. For later use, write

$$r_q := \lceil \log_2 q \rceil, \quad \mathbf{g}_q^\top := (1, 2, \dots, 2^{r_q-1}), \quad \mathbf{G}_{d,q} := \mathbf{I}_d \otimes \mathbf{g}_q^\top.$$

Thus $\mathbf{G}_{n,q} = \mathbf{G}$, and the public hint satisfies

$$\mathbf{D}\mathbf{T} = \mathbf{U} \bmod q, \quad \mathbf{D} = [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}], \quad \mathbf{U} = \mathbf{G}_{\ell n, q} = \mathbf{I}_\ell \otimes \mathbf{G}.$$

Given $(\mathbf{A}, \mathbf{W}, \mathbf{T})$, the goal is to output

$$\mathbf{0} \neq \mathbf{x} \in \mathbb{Z}^m, \quad \|\mathbf{x}\|_2 \leq \gamma, \quad \mathbf{A}\mathbf{x} = \mathbf{0} \bmod q.$$

We denote the problem by $\text{sSIS}_{n,m,q,\ell,k,s,\gamma}$.

Definition 3.10 (ℓ -Decomposed SIS [AMR26a]). Fix parameters n, k, q, ℓ, σ as in Definition 3.8 and a bound $\gamma > 0$. Sample $(\overline{\mathbf{W}}, \mathbf{R})$ and form $\mathbf{A}_{\text{dec}} \in \mathbb{Z}_q^{n \times \ell^2 n \lceil \log q \rceil}$ exactly as in that definition. Given $(\overline{\mathbf{W}}, \mathbf{R})$, the goal is to output

$$\mathbf{0} \neq \mathbf{x} \in \mathbb{Z}^{\ell^2 n \lceil \log q \rceil}, \quad \|\mathbf{x}\|_2 \leq \gamma, \quad \mathbf{A}_{\text{dec}} \mathbf{x} = \mathbf{0} \bmod q.$$

We denote the problem by $\text{dSIS}_{n,q,\ell,k,\sigma,\gamma}$.

In either case, the SIS assumption states that every probabilistic polynomial-time algorithm succeeds only with negligible probability. The definitions use Euclidean norm bounds. When a reduction estimates the maximum norm of a solution of dimension d , we use $\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|_2 \leq \sqrt{d} \|\mathbf{x}\|_\infty$ to translate between the two conventions, with the corresponding change in the solution bound. We also recall the standard learning with errors (LWE) assumption.

Definition 3.11 (Decisional learning with errors). Let $r, M, q \in \mathbb{N}$ and let $\chi > 0$. The (r, M, χ, q) -decisional LWE problem is to distinguish

$$(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \quad \text{from} \quad (\mathbf{A}, \mathbf{u}^\top),$$

where $\mathbf{A} \sim \mathbb{Z}_q^{r \times M}$, $\mathbf{s} \sim \mathbb{Z}_q^r$, $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^M, \chi}$, and $\mathbf{u} \sim \mathbb{Z}_q^M$ are mutually independent.

Definition 3.12 (Decisional nonuniform learning with errors). *Let $K, M, q \in \mathbb{N}$ and let $\sigma, \chi > 0$. The (K, M, σ, χ, q) -decisional NLWE problem is to distinguish*

$$(\mathbf{R}, \mathbf{r}^\top \mathbf{R} + \mathbf{e}^\top) \quad \text{from} \quad (\mathbf{R}, \mathbf{u}^\top),$$

where $\mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{K \times M}, \sigma}$, $\mathbf{r} \sim \mathbb{Z}_q^K$, $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^M, \chi}$, and $\mathbf{u} \sim \mathbb{Z}_q^M$ are mutually independent.

We use the following specialization of the nonuniform LWE reduction.

Theorem 3.13 ([BLMR13, Corollary 4.6]). *Let $r, K, q \in \mathbb{N}$ with $K \geq 6r \log q$, and let $\sigma = \Omega(\sqrt{r \log q})$. For every polynomially bounded number of samples M , the (K, M, σ, χ, q) -decisional NLWE problem is at least as hard as the (r, M, χ, q) -decisional LWE problem.*

3.4 Additional Lattice Tools

Definition 3.14 (Lattice). *A full rank lattice of dimension n consists of any set*

$$\Lambda = \{\mathbf{B}\mathbf{x} | \mathbf{x} \in \mathbb{Z}^n\}$$

where $\mathbf{B} \in \text{GL}_n(\mathbb{R})$ is an invertible matrix of $\mathbb{R}^{n \times n}$. We say that $\mathbf{B}$ is a basis for Λ . The determinant (or volume) of the lattice Λ is $|\det \mathbf{B}|$.

Notice that $\mathbf{B}_0, \mathbf{B}_1 \in \text{GL}_n(\mathbb{R})$ are bases for the same lattice Λ if and only if there exists a matrix $\mathbf{U} \in \mathbb{Z}^{n \times n}$ such that $\mathbf{B}_0 = \mathbf{B}_1 \mathbf{U}$ and $\det \mathbf{U} \in \{\pm 1\}$. A $\mathbf{U}$ with these properties is called unimodular. Notice that this observation proves that the determinant of the lattice Λ is independent of the basis we consider. If $\mathbf{S} \in \mathbb{R}^{n \times m}$ is such that $\Lambda = \mathbf{S}\mathbb{Z}^m$, we say that $\mathbf{S}$ generates Λ . This does not mean that $\mathbf{S}$ is a basis, as $\mathbf{S}$ may not be invertible (in fact, $\mathbf{S}$ may not even be a square matrix). It turns out that a set of generators can be efficiently turned into a basis (this is done by deriving the Hermite normal form of $\mathbf{S}$).

Lemma 3.15. *There exists a deterministic polynomial time algorithm that, given as input a matrix $\mathbf{S} \in \mathbb{Q}^{n \times m}$ that generates a lattice Λ , outputs a basis $\mathbf{B} \in \mathbb{Q}^{n \times n}$ for Λ .*

All norms in this section are labeled explicitly. For a vector we use the Euclidean norm $\|\cdot\|_2$. For a matrix we use the spectral norm $\|\cdot\|_{2 \rightarrow 2}$, the Frobenius norm $\|\cdot\|_F$, and

$$\|\mathbf{B}\|_{2, \infty} := \max_j \|\mathbf{b}_j\|_2$$

for the maximum Euclidean column norm. For a basis $\mathbf{B}$, let $\tilde{\mathbf{B}}$ denote its Gram–Schmidt orthogonalization and use the same $\|\cdot\|_{2, \infty}$ convention for its Gram–Schmidt vectors. We call an ordered basis $\mathbf{B} = [\mathbf{b}_1 | \cdots | \mathbf{b}_d]$ *size-reduced* when its Gram–Schmidt coefficients satisfy $|\mu_{j,i}| \leq 1/2$ for every $i < j$, where $\mathbf{b}_j = \tilde{\mathbf{b}}_j + \sum_{i < j} \mu_{j,i} \tilde{\mathbf{b}}_i$. In particular,

$$\|\mathbf{B}\|_{2 \rightarrow 2} \leq \|\mathbf{B}\|_F \leq d \|\tilde{\mathbf{B}}\|_{2, \infty}.$$

Definition 3.16 (q -ary lattices). Let $q > 0$ be a positive integer. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix. We define the lattices

$$\begin{aligned}\Lambda(\mathbf{A}) &:= \{\mathbf{v} \in \mathbb{Z}^n \mid \exists \mathbf{x} \in \mathbb{Z}^m : \mathbf{v} \equiv \mathbf{A}\mathbf{x} \bmod q\} \\ \Lambda^\perp(\mathbf{A}) &:= \{\mathbf{v} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{v} \equiv \mathbf{0} \bmod q\}\end{aligned}$$

Definition 3.17 (Dual lattice). Let Λ be a lattice over $\mathbb{R}^n$. The dual lattice of Λ is

$$\Lambda^* := \{\mathbf{x} \in \mathbb{R}^n \mid \forall \mathbf{y} \in \Lambda : \mathbf{x}^\top \mathbf{y} \in \mathbb{Z}\}.$$

Lemma 3.18. Let Λ be a lattice. Then, $(\Lambda^*)^* = \Lambda$. Moreover, if $\Lambda = \Lambda(\mathbf{A})$ for a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we have that $\Lambda^* = \frac{1}{q}\Lambda^\perp(\mathbf{A}^\top)$.

For a full rank lattice $\Lambda \subseteq \mathbb{R}^d$, a coset $\mathbf{c} + \Lambda$, and $\sigma > 0$, let $\mathcal{D}_{\mathbf{c}+\Lambda, \sigma}$ denote the spherical discrete Gaussian whose mass at $\mathbf{x}$ is proportional to $\exp(-\pi\|\mathbf{x}\|_2^2/\sigma^2)$. We use $\eta_\varepsilon(\Lambda)$ for the smoothing parameter.

Lemma 3.19 (Smoothing bound [Pei08, Lemma 3.5, author's manuscript]). For every full rank lattice $\Lambda \subseteq \mathbb{R}^n$ and $\varepsilon > 0$,

$$\lambda_1^\infty(\Lambda^*) \cdot \eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1 + 1/\varepsilon))}{\pi}},$$

where $\|\mathbf{x}\|_\infty = \max_i |x_i|$, $\lambda_1^\infty(\Lambda^*) = \min_{\mathbf{x} \in \Lambda^* \setminus \{0\}} \|\mathbf{x}\|_\infty$.

Lemma 3.20 (Gaussian residues and point probabilities). Let $\Lambda_2 \subseteq \Lambda_1$ be full rank lattices and let $0 < \varepsilon \leq 1/2$ and $\sigma \geq \eta_{\varepsilon/4}(\Lambda_2)$. If $\mathbf{v} \sim \mathcal{D}_{\Lambda_1, \sigma}$, then $\mathbf{v} \bmod \Lambda_2$ is within statistical distance $\varepsilon/2$ of uniform on Λ_1/Λ_2 . For a full rank lattice $\Lambda \subseteq \mathbb{R}^d$, if $\sigma \geq 2\eta_\varepsilon(\Lambda)$, then for every coset $\mathbf{c} + \Lambda$ and every point $\mathbf{x}$,

$$\Pr_{\mathbf{v} \sim \mathcal{D}_{\mathbf{c}+\Lambda, \sigma}} [\mathbf{v} = \mathbf{x}] \leq 2^{-d} \frac{1 + \varepsilon}{1 - \varepsilon}.$$

Finally, if $0 < \varepsilon \leq 1/3$ and $\sigma \geq \eta_\varepsilon(\Lambda)$, then a sample $\mathbf{v} \sim \mathcal{D}_{\mathbf{c}+\Lambda, \sigma}$ satisfies $\|\mathbf{v}\|_2 \leq \sigma\sqrt{d}$ except with probability at most 2^{-d+1} . This tail bound does not require the stronger width hypothesis of the point probability bound.

The residue statement and the point probability bound are standard consequences of smoothing; see, e.g., [MR04, GPV08, PR06]. The use of $\eta_{\varepsilon/4}$ in the residue statement absorbs the normalization loss between coset masses and yields the displayed $\varepsilon/2$ distance for $\varepsilon \leq 1/2$.

Proof of the tail bound. Write $\rho_\sigma(X) = \sum_{\mathbf{x} \in X} e^{-\pi\|\mathbf{x}\|_2^2/\sigma^2}$. Comparing the Gaussian masses at widths σ and $\sqrt{2}\sigma$ gives

$$\Pr[\|\mathbf{v}\|_2 > \sigma\sqrt{d}] \leq e^{-\pi d/2} \frac{\rho_{\sqrt{2}\sigma}(\mathbf{c} + \Lambda)}{\rho_\sigma(\mathbf{c} + \Lambda)} \leq \frac{1 + \varepsilon}{1 - \varepsilon} 2^{d/2} e^{-\pi d/2} \leq 2^{-d+1}.$$

The second inequality uses smoothing at both widths, and the last uses $\varepsilon \leq 1/3$ and $\pi \geq 3 \ln 2$. $\square$

Lemma 3.21 (Gaussian mass on an affine subspace). *For a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$, a coset $\mathbf{u} + \mathcal{L}$, a proper affine subspace H , and $\sigma \geq 2\eta_\varepsilon(\mathcal{L})$, where $0 < \varepsilon \leq 1/16$,*

$$\Pr_{\mathbf{v} \sim \mathcal{D}_{\mathbf{u} + \mathcal{L}, \sigma}} [\mathbf{v} \in H] \leq \frac{1 + \varepsilon}{2(1 - \varepsilon)}. \quad (6)$$

Proof. Write $\rho_\sigma(X) = \sum_{\mathbf{x} \in X} e^{-\pi \|\mathbf{x}\|_2^2 / \sigma^2}$. Since $\eta_\varepsilon(2\mathcal{L}) = 2\eta_\varepsilon(\mathcal{L})$, smoothing for $2\mathcal{L}$ and $\mathcal{L}$ gives

$$\rho_\sigma(\mathbf{a} + 2\mathcal{L}) \leq \frac{(1 + \varepsilon)\sigma^n}{2^n \det \mathcal{L}}, \quad \rho_\sigma(\mathbf{u} + \mathcal{L}) \geq \frac{(1 - \varepsilon)\sigma^n}{\det \mathcal{L}}.$$

If the intersection $(\mathbf{u} + \mathcal{L}) \cap H$ is nonempty, it is a translate of the primitive subgroup $\mathcal{L} \cap \text{dir}(H)$, whose rank is at most $n - 1$. This subgroup extends to a basis of $\mathcal{L}$, so the intersection meets at most 2^{n-1} residue classes modulo $2\mathcal{L}$. Summing the corresponding coset masses proves the claim. $\square$

Lemma 3.22 (Efficient Gaussian sampling). *Given a basis $\mathbf{L}$ of Λ , samples from $\mathcal{D}_{\frac{1}{q}\Lambda, \sigma}$ can be generated in probabilistic polynomial time, up to statistical error negligible in λ , whenever*

$$\sigma \geq \frac{\|\tilde{\mathbf{L}}\|_{2, \infty}}{q} \cdot \omega(\sqrt{\log(\lambda + n)}).$$

Proof. Apply the standard basis-dependent discrete Gaussian sampler [GPV08] to the basis $\mathbf{L}/q$ of Λ/q , with truncation and numerical precision chosen for negligible error in λ . $\square$

Lemma 3.23 (Joint leftover hashing with sign masks). *Let q be an odd prime, let $\mathbf{A} \sim \mathbb{Z}_q^{n \times m_0}$, and independently sample matrices $\mathbf{R}_i \sim \{-1, 1\}^{m_0 \times c_i}$ for $i \in [r]$. If $C = \sum_i c_i$, then*

$$\Delta((\mathbf{A}, \mathbf{A}\mathbf{R}_1, \dots, \mathbf{A}\mathbf{R}_r), (\mathbf{A}, \mathbf{U}_1, \dots, \mathbf{U}_r)) \leq \frac{C}{2} \sqrt{q^n 2^{-m_0}},$$

where the $\mathbf{U}_i$ are mutually independent and uniform. In particular, $m_0 \geq n \log_2 q + 2\lambda + 2 \log_2 C$ makes the distance at most $2^{-\lambda}$.

Proof. For odd q , two distinct sign vectors remain distinct modulo q . The usual leftover hash argument therefore applies to each column. Replace the C product columns one at a time by independent uniform columns, retaining $\mathbf{A}$ throughout, and sum the statistical distance bounds. This gives the displayed bound. Notice that this is a joint statement over random $\mathbf{A}$, not a statement conditioned on every fixed value of $\mathbf{A}$. $\square$

Lemma 3.24 (Nearest plane [Bab85]). *There is a deterministic polynomial-time algorithm which, given a basis $\mathbf{P} \in \mathbb{Q}^{d \times d}$ of a lattice Λ' and a target $\mathbf{u} \in \mathbb{Q}^d$, returns $\mathbf{u}' \in \Lambda'$ satisfying*

$$\|\mathbf{u} - \mathbf{u}'\|_2 \leq \frac{1}{2} \left(\sum_{i=1}^d \|\tilde{\mathbf{p}}_i\|_2^2 \right)^{1/2} \leq \frac{\sqrt{d}}{2} \|\mathbf{P}\|_{2, \infty}.$$

Lemma 3.25 (Basis conversion). *Given a basis of a rank- d rational lattice $\mathcal{L}$ and d independent vectors in $\mathcal{L}$ of Euclidean norm at most H , one can compute in polynomial time a basis $\mathbf{B}$ of $\mathcal{L}$ such that*

$$\|\tilde{\mathbf{B}}\|_{2,\infty} \leq H, \quad \|\mathbf{B}\|_{2,\infty} \leq \frac{\sqrt{d+3}}{2}H.$$

The same conclusion holds given a generating set of maximum column norm H .

Proof. Write the given independent vectors as the columns of $\mathbf{S}$ and express them in the supplied lattice basis $\mathbf{B}_0$: $\mathbf{S} = \mathbf{B}_0\mathbf{C}$ for a nonsingular integer matrix $\mathbf{C}$. Compute a unimodular $\mathbf{U}$ such that $\mathbf{UC}$ is upper triangular, using a row Hermite normal form. The basis $\mathbf{B} = \mathbf{B}_0\mathbf{U}^{-1}$ then satisfies $\text{span}(\mathbf{b}_1, \dots, \mathbf{b}_i) = \text{span}(\mathbf{s}_1, \dots, \mathbf{s}_i)$ for every i . The orthogonal projection of $\mathbf{s}_i$ onto the complement of the preceding span is a nonzero integer multiple of $\tilde{\mathbf{b}}_i$. Hence $\|\tilde{\mathbf{b}}_i\|_2 \leq H$. For $j = 2, \dots, d$ and $i = j - 1, \dots, 1$, perform the integer column operation $\mathbf{b}_j \leftarrow \mathbf{b}_j - \lfloor \mu_{j,i} \rfloor \mathbf{b}_i$, where $\mu_{j,i}$ is the current Gram–Schmidt coefficient. These operations preserve the lattice and its Gram–Schmidt vectors, ensure $|\mu_{j,i}| \leq 1/2$, and give $\|\mathbf{b}_i\|_2^2 \leq H^2 + (i-1)H^2/4$. The normal form algorithm computes its transformation matrices with polynomial bit complexity [KB79]; inversion of a unimodular matrix and these integer column operations also have polynomial bit complexity. For a generating set, first compute a basis of the generated lattice after clearing denominators and select d independent generating columns. $\square$

Binary Gadgets. For $d \geq 2$, put $r_d := \lceil \log_2 d \rceil$ and

$$\mathbf{g}_d^\top := (1, 2, \dots, 2^{r_d-1}).$$

For invariant factors $d_1, \dots, d_h$, define

$$\mathbf{G}_\Gamma := \text{diag}(\mathbf{g}_{d_1}^\top, \dots, \mathbf{g}_{d_h}^\top) \in \mathbb{Z}^{h \times k_\Gamma}, \quad k_\Gamma := \sum_{j=1}^h r_{d_j}.$$

If y_j is represented by its canonical integer in $[0, d_j)$, ordinary binary expansion gives $\mathbf{z}_3 \in \{0, 1\}^{k_\Gamma}$ with $\mathbf{G}_\Gamma \mathbf{z}_3 = \mathbf{y}$ over $\mathbb{Z}$.

For the signed residual gadget, let $M \in \mathbb{N}$ with $M \geq 1$ and set

$$r_M := \lceil \log_2(M+1) \rceil, \quad \mathbf{G}_{n,M} := \mathbf{I}_n \otimes \mathbf{g}_{M+1}^\top.$$

For every $\mathbf{a} \in \mathbb{Z}^n$ with $\|\mathbf{a}\|_\infty \leq M$, define $\text{sdec}_M(\mathbf{a})$ by taking the binary expansion of $|a_i|$ and multiplying its digits by $\text{sgn}(a_i)$. Then

$$\text{sdec}_M(\mathbf{a}) \in \{-1, 0, 1\}^{nr_M}, \quad \mathbf{G}_{n,M} \text{sdec}_M(\mathbf{a}) = \mathbf{a}$$

over the integers. The use of $M+1$ makes the representation valid also at the inclusive endpoint $|a_i| = M$.

Preimage Lattices and Short Generators. For $\mathbf{D} \in \mathbb{Z}_q^{d \times N}$ and $\mathbf{U} \in \mathbb{Z}_q^{d \times t}$, write

$$\Lambda_{\mathbf{U}}^{\perp}(\mathbf{D}) := \{\mathbf{Z} \in \mathbb{Z}^{N \times t} : \mathbf{D}\mathbf{Z} = \mathbf{U} \bmod q\}.$$

The Gaussian on this set is sampled independently columnwise and is exactly the distribution $\mathbf{D}^{-1}(\mathbf{U}, s)$ from Eq. (3). Write $\Lambda^{\perp}(\mathbf{D}) = \{\mathbf{z} \in \mathbb{Z}^N : \mathbf{D}\mathbf{z} = \mathbf{0} \bmod q\}$ and $\mathbf{G}_{d,q} = \mathbf{I}_d \otimes \mathbf{g}_q^{\top}$.

Lemma 3.26 (Short generators from a gadget trapdoor). *Let q be an odd prime, let $\mathbf{D} \in \mathbb{Z}_q^{d \times N}$, and let $\hat{\mathbf{T}} \in \mathbb{Z}^{N \times c}$ satisfy $\mathbf{D}\hat{\mathbf{T}} = \mathbf{G}_{d,q} \bmod q$, where $c = d \lceil \log_2 q \rceil$. One can compute a generating set of $\Lambda^{\perp}(\mathbf{D})$ whose columns have Euclidean norm at most $1 + \left\| \hat{\mathbf{T}} \right\|_{2 \rightarrow 2} \sqrt{c + 5}$.*

Proof. Write $\mathbf{G} = \mathbf{G}_{d,q}$ and let $\mathbf{F} \in \{0, 1\}^{c \times N}$ be the columnwise binary expansion of the canonical representatives of $\mathbf{D}$, so $\mathbf{G}\mathbf{F} = \mathbf{D} \bmod q$. There is a basis $\mathbf{S}_G$ of $\Lambda^{\perp}(\mathbf{G})$ with maximum column norm at most $\sqrt{\lceil \log_2 q \rceil + 5}$: in each gadget block use the carry vectors $2\mathbf{e}_i - \mathbf{e}_{i+1}$ and the binary expansion of q . Their determinant has absolute value q per block, equal to that kernel's index.

The columns of

$$[\mathbf{I}_N - \hat{\mathbf{T}}\mathbf{F} \mid \hat{\mathbf{T}}\mathbf{S}_G] \quad (7)$$

generate $\Lambda^{\perp}(\mathbf{D})$. They belong to this lattice, and for every $\mathbf{z} \in \Lambda^{\perp}(\mathbf{D})$, $\mathbf{F}\mathbf{z} \in \Lambda^{\perp}(\mathbf{G})$ and

$$\mathbf{z} = (\mathbf{I}_N - \hat{\mathbf{T}}\mathbf{F})\mathbf{z} + \hat{\mathbf{T}}\mathbf{F}\mathbf{z}.$$

Each generating column has the claimed norm, since columns of $\mathbf{F}$ have norm at most $\sqrt{c}$. This proves the algebraic claim without imposing the sampler's dimension or width conditions. $\square$

3.5 Quotient Decomposition and Smith Normal Form

Definition 3.27 (Quotient coordinates). *Let $\mathbf{L}, \mathbf{P} \in \mathbb{Q}^{d \times d}$ be nonsingular, with $\Lambda = \mathbf{L}\mathbb{Z}^d \subseteq \Lambda' = \mathbf{P}\mathbb{Z}^d$, and put $\Gamma := \Lambda' / \Lambda$. The matrix $\mathbf{H} := \mathbf{P}^{-1}\mathbf{L}$ is a nonsingular integer matrix. Its Smith normal form is*

$$\mathbf{E}\mathbf{H}\mathbf{F} = \text{diag}(e_1, \dots, e_d), \quad 1 \leq e_1 \mid e_2 \mid \dots \mid e_d, \quad (8)$$

where $\mathbf{E}, \mathbf{F} \in \mathbb{Z}^{d \times d}$ are unimodular. The positive integers e_i are the invariant factors. Let $i_1 < \dots < i_h$ be precisely the indices for which $e_{i_j} > 1$, and put $d_j := e_{i_j}$. Define the matrix of quotient representatives by

$$\mathbf{B} := [\mathbf{P}\mathbf{E}^{-1}\mathbf{e}_{i_1} \mid \dots \mid \mathbf{P}\mathbf{E}^{-1}\mathbf{e}_{i_h}] \in \mathbb{Q}^{d \times h}, \quad (9)$$

where $\mathbf{e}_i$ is the i th standard basis vector. For $\mathbf{u}' \in \Lambda'$, its quotient coordinates $\mathbf{y} = (y_1, \dots, y_h)^{\top}$ are defined by taking y_j to be the canonical representative in $\{0, \dots, d_j - 1\}$ of $(\mathbf{E}\mathbf{P}^{-1}\mathbf{u}')_{i_j} \bmod d_j$. The associated binary gadget $\mathbf{G}_{\Gamma}$ has width $k_{\Gamma} := \sum_{j=1}^h \lceil \log_2 d_j \rceil$.

Lemma 3.28 (Quotient decomposition). *The data in Definition 3.27 can be computed in polynomial time from the rational bases $\mathbf{L}, \mathbf{P}$ and the rational vector $\mathbf{u}' \in \Lambda'$. They satisfy*

$$\Gamma \cong \bigoplus_{j=1}^h \mathbb{Z}/d_j\mathbb{Z}, \quad |\Gamma| = |\det \mathbf{H}| = \prod_{j=1}^h d_j, \quad \mathbf{u}' - \mathbf{B}\mathbf{y} \in \Lambda.$$

One can also compute $\mathbf{z}_3 \in \{0, 1\}^{k_{\Gamma}}$ with $\mathbf{G}_{\Gamma}\mathbf{z}_3 = \mathbf{y}$ over the integers. If $h > 0$, the exact width satisfies $0 \leq k_{\Gamma} - \log_2 |\Gamma| < h$; it need not equal $\lceil \log_2 |\Gamma| \rceil$. If $h = 0$, then $\Lambda' = \Lambda$ and $k_{\Gamma} = 0$; the coordinate vectors and generator and gadget matrices are empty, with $\mathbf{B}\mathbf{y} = \mathbf{0}$.

Proof. The Smith normal form and its transformation matrices can be computed with polynomial bit complexity [KB79]. Multiplication by $\mathbf{P}^{-1}$ identifies Γ with $\mathbb{Z}^d/\mathbf{H}\mathbb{Z}^d$. Since $\mathbf{F}\mathbb{Z}^d = \mathbb{Z}^d$, multiplication by $\mathbf{E}$ identifies this group with $\bigoplus_{i=1}^d \mathbb{Z}/e_i\mathbb{Z}$. Omitting the trivial factors gives the asserted decomposition and cardinality.

For each j , applying $\mathbf{E}\mathbf{P}^{-1}$ to column j of $\mathbf{B}$ gives $\mathbf{e}_{i_j}$. Thus every coordinate of $\mathbf{E}\mathbf{P}^{-1}(\mathbf{u}' - \mathbf{B}\mathbf{y})$ is divisible by its corresponding e_i , including the coordinates with $e_i = 1$. Consequently this vector belongs to $\mathbf{E}\mathbf{H}\mathbb{Z}^d$, which proves $\mathbf{u}' - \mathbf{B}\mathbf{y} \in \Lambda$.

Binary expansion of each $y_j \in \{0, \dots, d_j - 1\}$ gives the claimed $\mathbf{z}_3$ in polynomial time. Finally, $\log_2 |\Gamma| = \sum_j \log_2 d_j$, and each difference $\lceil \log_2 d_j \rceil - \log_2 d_j$ lies in $[0, 1)$. Summing proves the width bound. The empty case follows by taking empty sums to be zero and empty products to be one. $\square$

4 A Worst-Case to Average-Case Reduction for Succinct SIS

We prove the reduction described in the technical overview. A succinct SIS solver first gives one short vector of Λ that lies outside any fixed proper subspace with probability at least an inverse polynomial in λ . Repeating this step gives n independent vectors and then a shorter basis of the full lattice. Iterating the basis improvement proves the worst-case reduction.

All dimensions and input bit lengths are polynomially bounded in the security parameter λ , and $n = \omega(\log \lambda)$. Supplied scalar parameters are rational, with binary numerator/denominator encodings of polynomial length. We use a supplied inverse-polynomial lower bound on the solver's success to set the maximum number of trials in each solver repetition. Gaussian sampling, numerical approximation, and failure to obtain full row rank within λ attempts contribute a negligible error per invocation. The proofs first analyze exact initial Gaussian samples and then include this error.

4.1 The Worst-Case Problem

Definition 4.1 (Superlattice SVP). *Fix a positive integer ℓ . An instance consists of a basis $\mathbf{L} \in \mathbb{Q}^{n \times n}$, where*

$$\Lambda := \mathbf{L}\mathbb{Z}^n, \quad \mathbf{L}_\ell := \mathbf{I}_\ell \otimes \mathbf{L}, \quad \Lambda_\ell := \Lambda^{\oplus \ell} = \mathbf{L}_\ell \mathbb{Z}^{\ell n},$$

together with a basis $\mathbf{P} \in \mathbb{Q}^{\ell n \times \ell n}$ of a superlattice

$$\Lambda_\ell \subseteq \Lambda' \subseteq \mathbb{Q}^{\ell n}.$$

Write $\mathbf{P} = [\mathbf{p}_1 | \dots | \mathbf{p}_{\ell n}]$. The promises are

$$\max_{i \in [\ell n]} \|\mathbf{p}_i\|_2 \leq \beta, \quad |\Lambda' / \Lambda_\ell| \leq 2^k.$$

The goal is to output a nonzero $\mathbf{v} \in \Lambda$ with $\|\mathbf{v}\|_2 \leq \gamma'$. We denote this problem by $\text{SuperSVP}_{n, \ell, k, \beta, \gamma'}$.

4.2 The Reduction and Its Parameters

Fix $\lambda \geq 4$, integers $n, \ell, k \geq 1$, bounds $\beta, \gamma \geq 1$, and put $\varepsilon = 2^{-\lambda}$. Taking $\beta \geq 1$ is without loss of generality, since a smaller promise bound can be enlarged to 1. Throughout this section, $\|\mathbf{M}\|$

denotes the spectral norm of a matrix $\mathbf{M}$. In particular, $\|\tilde{\mathbf{L}}\|_{2,\infty} \leq \|\mathbf{L}\|$, since each Gram–Schmidt vector is an orthogonal projection of a column of $\mathbf{L}$.

Given an instance $(\mathbf{L}, \mathbf{P})$ from Definition 4.1, let $d_1, \dots, d_h$ be the nontrivial invariant factors of $\Gamma := \Lambda' / \Lambda_\ell$ and put

$$k_\Gamma := \sum_{j=1}^h \lceil \log_2 d_j \rceil.$$

For the trivial quotient, $h = k_\Gamma = 0$. The reduction is also given a parameter $\sigma_0 \in \mathbb{Q}_{>0}$, satisfying

$$\sigma_0 \geq 2\eta_{2^{-(\lambda+2)}}(\Lambda).$$

All parameter choices below use the original input basis $\mathbf{L}$ and remain fixed throughout the reduction.

Table 1 summarizes the parameter roles and dependencies. In particular, the relation between the security parameter and the lower bound on the Gaussian sampling width is

$$\lambda \mapsto \varepsilon/4 = 2^{-(\lambda+2)}, \quad \sigma_0 \geq 2\eta_{2^{-(\lambda+2)}}(\Lambda).$$

Thus λ alone does not determine σ_0 : the lower bound also depends on the scale and geometry of Λ . For a fixed lattice, increasing λ cannot decrease the required lower bound.

The cardinality promise gives $k_\Gamma \leq 2k$, as shown in the proof of Lemma 4.5. The construction uses the exact hint width k_Γ ; the bound $2k$ is used only in the common parameter bounds. Alternatively, one may pad the quotient gadget and its coefficient vector with zeros to width $2k$, and sample the masking matrix with $2k$ columns. The resulting matrix $\mathbf{W}$ then has $2k$ columns, and the same proof gives a reduction to $\text{sSIS}_{n,m,q,\ell,2k,s,\gamma}$.

Choose an integer p and an odd prime q satisfying

$$p \geq \max \left\{ 1, \frac{6\|\mathbf{L}\|\|\mathbf{L}^{-1}\|\sqrt{\ell n}}{\sigma_0} \right\}, \quad (10)$$

$$q \geq 17 \max \left\{ 1, k\|\mathbf{L}\|, \frac{p\beta\sqrt{\ell} n^{3/2}\|\mathbf{L}\|}{\sigma_0} \right\}. \quad (11)$$

Define the common residual bound and the matrix dimensions by

$$r_M := \left\lceil \log_2 \left(\frac{q\|\mathbf{L}^{-1}\|\sqrt{\ell n}}{p} + 2p\beta\sqrt{\ell} n^{3/2} \right) \right\rceil, \quad M := 2^{r_M} - 1, \quad (12)$$

$$m_0 := \lceil 3n\lceil \log_2 q \rceil + 4\lambda + 2\log_2(2nr_M + 2\ell k) \rceil, \quad (13)$$

$$m := m_0 + n + nr_M, \quad \gamma_{\det} := 2m_0 n \sqrt{r_M}. \quad (14)$$

Finally, fix the trapdoor width by requiring

$$s \geq (\ell m + 2k) (2p\beta n \sqrt{\ell} + nr_M + 2k) \log(\ell n). \quad (15)$$

These choices determine the target $\text{sSIS}_{n,m,q,\ell,k_\Gamma,s,\gamma}$ for the given input instance. The Gaussian width σ used inside the reduction may change from one round to the next; the public parameters above remain fixed.

Parameter	Meaning	Relations and dependencies
λ, ε	Security and statistical accuracy	$\varepsilon = 2^{-\lambda}$; dimensions are polynomial in λ , and $n = \omega(\log \lambda)$.
n, ℓ	Lattice dimension and number of copies	$\Lambda \subseteq \mathbb{R}^n$, $\Lambda_\ell = \Lambda^{\oplus \ell}$; the SIS matrix $\mathbf{A}$ has n rows.
k, β	Size and quality of the hint	$ \Lambda'/\Lambda_\ell \leq 2^k$ and $\max_i \ \mathbf{p}_i\ _2 \leq \beta$.
k_Γ	Number of SIS hint columns	$k_\Gamma = \sum_{j=1}^h \lceil \log_2 d_j \rceil \leq 2k$, where $d_1, \dots, d_h$ are the nontrivial invariant factors of $\Gamma = \Lambda'/\Lambda_\ell$; $k_\Gamma = 0$ for the trivial quotient. The matrix $\mathbf{W}$ has k_Γ columns.
σ_0	Supplied upper bound on the smoothing parameter	$\sigma_0 \geq 2\eta_{2^{-(\lambda+2)}}(\Lambda)$: depends on the input lattice and smoothing accuracy.
γ, δ	Solver guarantees	$\ \mathbf{x}\ _2 \leq \gamma$ with success at least $\delta \geq 1/\text{poly}(\lambda)$; δ controls repetition.
p	Integer rounding scale	Depends on $\mathbf{L}, \mathbf{L}^{-1}, n, \ell, \sigma_0$ through Equation (10); scales the residual before rounding to an integer vector in Lemma 4.5.
q	Odd prime modulus	Depends on $\mathbf{L}, n, \ell, k, \beta, p, \sigma_0$ through Equation (11); iteration also requires $q \geq 4\gamma\gamma_{\det}\lambda$.
r_M, M	Residual bit length and magnitude bound	$M = 2^{r_M} - 1$; determined by $p, q, \mathbf{L}^{-1}, n, \ell, \beta$ in Equation (12).
m_0, m	Initial Gaussian columns and total SIS width	m_0 depends on $n, q, \lambda, r_M, \ell, k$ in Equation (13); $m = m_0 + n + nr_M$.
$\gamma_{\det}$	Coefficient of the sampling width in the norm bound	$\gamma_{\det} = 2m_0n\sqrt{r_M}$; appears in $\ \mathbf{Q}\ \leq \gamma_{\det}\sigma + \sigma_0$ from Lemma 4.6, except with negligible probability; $\mathbf{Q}$ is the matrix used to compute $\mathbf{v} = \mathbf{Q}\mathbf{x}$.
s	Published trapdoor Gaussian width	Lower bound from $\ell, m, k, p, \beta, n, r_M$ in Equation (15); used to sample $\mathbf{T}$.
σ_j	Gaussian width in round j	$\sigma_j = \max\{\sigma_0, (\lambda/q)\ \widetilde{\mathbf{L}_{j-1}}\ _{2,\infty}\}$ for $j \geq 1$; samples columns from $\mathcal{D}_{\Lambda/q, \sigma_j}$.
γ'	Final lattice vector length bound	$\gamma' = 4\gamma\gamma_{\det}\sigma_0$; an absolute Euclidean norm bound.

Table 1: Parameters of the worst-case to average-case reduction. All parameters above σ_j remain fixed; only the sampling basis and its width change between rounds.

Theorem 4.2 (Worst-case to average-case reduction). *Let $(\mathbf{L}, \mathbf{P})$ be any instance of Definition 4.1, let σ_0 be supplied as above, and let the parameters satisfy Equations (10) to (15), and suppose additionally that*

$$q \geq 4\gamma\gamma_{\det}\lambda. \quad (16)$$

If there is a probabilistic polynomial-time algorithm solving $\text{sSIS}_{n,m,q,\ell,k_\Gamma,s,\gamma}$ with probability at least $\delta \geq 1/\text{poly}(\lambda)$, then there is a probabilistic polynomial-time algorithm that, on the given instance, outputs a valid solution to $\text{SuperSVP}_{n,\ell,k,\beta,\gamma'}$, where

$$\gamma' := 4\gamma\gamma_{\det}\sigma_0.$$

Its success probability is at least $1 - \text{negl}(\lambda)$.

Proof. On input $(\mathbf{L}, \mathbf{P})$, run the following algorithm with the chosen succinct SIS parameters $(n, m, q, \ell, k_\Gamma, s, \gamma)$ from Table 1 and the supplied success lower bound δ .

1. Set $\mathbf{L}_0 = \mathbf{L}$ and

$$J := \left\lceil \log_2 \max \left\{ 1, \frac{\|\widetilde{\mathbf{L}}\|_{2,\infty}}{2\gamma\gamma_{\det}\sigma_0} \right\} \right\rceil + 1.$$

2. For $j = 1, \dots, J - 1$, set

$$\sigma_j := \max \left\{ \sigma_0, \frac{\lambda}{q} \|\widetilde{\mathbf{L}_{j-1}}\|_{2,\infty} \right\}.$$

Run Theorem 4.3 with input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{j-1}, \sigma_j)$ and fresh randomness. Abort if the call fails; otherwise let $\mathbf{L}_j$ be its returned basis of Λ .

3. Set $\sigma_J = \max\{\sigma_0, (\lambda/q) \|\widetilde{\mathbf{L}_{J-1}}\|_{2,\infty}\}$. Run Theorem 4.4 with input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{J-1}, \sigma_J)$ and fresh randomness at most $\lceil 2\lambda \ln(2)/\delta \rceil$ times. Output the first vector $\mathbf{v}$ satisfying $\mathbf{L}^{-1}\mathbf{v} \in \mathbb{Z}^n$ and $0 < \|\mathbf{v}\|_2 \leq \gamma(\gamma_{\det}\sigma_J + \sigma_0)$; abort if no trial succeeds.

For each successful basis call $j \in [J - 1]$, Theorem 4.3 gives

$$\begin{aligned} \|\widetilde{\mathbf{L}_j}\|_{2,\infty} &\leq \gamma(\gamma_{\det}\sigma_j + \sigma_0) \\ &\leq 2\gamma\gamma_{\det}\sigma_0 + \frac{\gamma\gamma_{\det}\lambda}{q} \|\widetilde{\mathbf{L}_{j-1}}\|_{2,\infty} \\ &\leq 2\gamma\gamma_{\det}\sigma_0 + \frac{1}{4} \|\widetilde{\mathbf{L}_{j-1}}\|_{2,\infty}. \end{aligned}$$

Here we used $\gamma_{\det} \geq 1$ and Equation (16). The same upper bound, with $j = J$, applies to $\|\mathbf{v}\|_2$ in the final step. Iterating therefore gives

$$\|\mathbf{v}\|_2 \leq \frac{8}{3}\gamma\gamma_{\det}\sigma_0 + 4^{-J} \|\widetilde{\mathbf{L}}\|_{2,\infty} < 4\gamma\gamma_{\det}\sigma_0 = \gamma'.$$

The strict inequality follows from the choice of J , which gives $4^{-J} \|\widetilde{\mathbf{L}}\|_{2,\infty} \leq \gamma\gamma_{\det}\sigma_0/2$. The final step checks that $\mathbf{v} \in \Lambda$ and that $\mathbf{v}$ is nonzero.

Let b be λ plus the total binary encoding length of $(\mathbf{L}, \mathbf{P})$, the fixed numerical parameters (including σ_0), and the supplied bound δ . Since $\|\mathbf{L}\| \leq n2^b$ and $\sigma_0 \geq 2^{-b}$, the definition gives $J = O(b)$. The preceding bounds and the column norm guarantee of Theorem 4.3 bound all entries of $\mathbf{L}_j$ in magnitude by $2^{O(b)}$. In lowest terms, their denominators divide the product of the denominators in $\mathbf{L}$, which is at most 2^b . Thus every intermediate basis has polynomial bit length in b . There are at most $J - 1$ calls to Theorem 4.3 and $\lceil 2\lambda \ln(2)/\delta \rceil$ calls to Theorem 4.4 in the final step. Their total running time is $\text{poly}(b, 1/\delta) = \text{poly}(\lambda)$ under the standing assumptions.

Conditional on preceding successful calls, each basis call fails with probability $\text{negl}(\lambda)$. Conditional on their success, each final trial succeeds with probability at least $\delta - \text{negl}(\lambda) \geq \delta/2$ for sufficiently large λ . The probability that all final trials fail is therefore at most $(1 - \delta/2)^{\lceil 2\lambda \ln(2)/\delta \rceil} \leq 2^{-\lambda}$. The total failure probability is at most $(J - 1)\text{negl}(\lambda) + 2^{-\lambda} = \text{negl}(\lambda)$. $\square$

4.3 From One Vector to a Short Basis

Theorem 4.3 (A short basis). *Fix an instance $(\mathbf{L}, \mathbf{P})$ and choose σ_0 and the parameters as in Section 4.2; in particular, they satisfy Equations (10) to (15). Suppose a probabilistic polynomial-time algorithm solves $\text{sSIS}_{n,m,q,\ell,k_\Gamma,s,\gamma}$ with probability at least $\delta \geq 1/\text{poly}(\lambda)$. Given a basis $\mathbf{L}_{\text{samp}} \in \mathbb{Q}^{n \times n}$ of Λ , and a width*

$$\sigma \geq \max \left\{ \sigma_0, \frac{\lambda}{q} \left\| \widetilde{\mathbf{L}_{\text{samp}}} \right\|_{2,\infty} \right\}, \quad (17)$$

put $H := \gamma(\gamma_{\text{det}}\sigma + \sigma_0)$. There is a probabilistic polynomial-time algorithm producing a basis $\mathbf{B} \in \mathbb{Q}^{n \times n}$ of Λ such that

$$\left\| \widetilde{\mathbf{B}} \right\|_{2,\infty} \leq H, \quad \|\mathbf{B}\|_{2,\infty} \leq \frac{\sqrt{n+3}}{2} H.$$

The algorithm succeeds with probability at least $1 - \text{negl}(\lambda)$.

Proof. On input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{samp}}, \sigma)$, run the following algorithm with the chosen succinct SIS parameters $(n, m, q, \ell, k_\Gamma, s, \gamma)$ from Table 1 and the supplied success lower bound δ .

1. Initialize an empty list of retained vectors, whose span is $\{0\}$.
2. For $i = 1, \dots, n$, make at most $\lceil (4/\delta)(\lambda \ln 2 + \ln n) \rceil$ calls to Theorem 4.4, each with input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{samp}}, \sigma)$ and fresh randomness. Retain the first output $\mathbf{v}$ satisfying $\mathbf{L}^{-1}\mathbf{v} \in \mathbb{Z}^n$, $\|\mathbf{v}\|_2 \leq H$, and $\mathbf{v} \notin \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_{i-1})$. Name this vector $\mathbf{v}_i$ and proceed to the next stage. If no such vector is found within the stated number of attempts, output $\perp$.
3. Apply Lemma 3.25 with the original basis $\mathbf{L}$ and the retained vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ to obtain a basis $\mathbf{B}$ of Λ . Output $\mathbf{B}$.

If the algorithm reaches the final step, its checks ensure that $\mathbf{v}_1, \dots, \mathbf{v}_n$ are linearly independent vectors of Λ of norm at most H . With dimension n , Lemma 3.25 therefore gives

$$\left\| \widetilde{\mathbf{B}} \right\|_{2,\infty} \leq H, \quad \|\mathbf{B}\|_{2,\infty} \leq \frac{\sqrt{n+3}}{2} H.$$

Let b be λ plus the total binary encoding length of $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{samp}}, \sigma)$, the fixed numerical parameters, and the supplied bound δ . There are at most $n \lceil (4/\delta)(\lambda \ln 2 + \ln n) \rceil$ calls to Theorem 4.4, all on the same input. Each call has running time and output bit length polynomial in b . The membership, norm, and independence checks and the basis conversion have polynomial bit complexity. Thus the total running time is $\text{poly}(b, 1/\delta) = \text{poly}(\lambda)$ under the standing assumptions.

Conditional on all preceding calls, the span of the retained vectors is fixed and proper until the list has length n . The next call uses independent fresh randomness, so Theorem 4.4 gives acceptance probability at least $\delta/4$ for sufficiently large λ . Conditional on reaching a stage, its failure probability is therefore at most

$$(1 - \delta/4)^{\lceil (4/\delta)(\lambda \ln 2 + \ln n) \rceil} \leq \frac{2^{-\lambda}}{n}.$$

A union bound over the n stages gives total failure probability at most $2^{-\lambda}$. □

Theorem 4.4 (One short vector). *Under the parameter and solver hypotheses of Theorem 4.3, there is a probabilistic polynomial-time algorithm that, on input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{samp}}, \sigma)$ as in that theorem, outputs $\perp$ or a nonzero $\mathbf{v} \in \Lambda$ satisfying*

$$\|\mathbf{v}\|_2 \leq \gamma(\gamma_{\text{det}}\sigma + \sigma_0).$$

It outputs a vector with probability at least $\delta - \text{negl}(\lambda)$. Moreover, for every fixed proper linear subspace $\mathcal{H} \subsetneq \mathbb{R}^n$, the probability of outputting a vector in $\Lambda \setminus \mathcal{H}$ is at least $\delta/4$ for all sufficiently large λ . The algorithm makes at most one call to the succinct SIS solver.

Proof. On input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{samp}}, \sigma)$, run the following algorithm with the chosen succinct SIS parameters $(n, m, q, \ell, k_{\Gamma}, s, \gamma)$ from Table 1.

1. Run the algorithm that constructs a succinct SIS instance in Lemma 4.5 with input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{samp}}, \sigma)$ and fresh randomness. Let $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ be its public instance and $\mathbf{Q}$ the matrix it retains.
2. Run the succinct SIS solver on $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ with independent fresh randomness. Output $\perp$ unless its output $\mathbf{x} \in \mathbb{Z}^m$ satisfies $\mathbf{A}\mathbf{x} = 0 \pmod q$ and $0 < \|\mathbf{x}\|_2 \leq \gamma$.
3. Compute $\mathbf{v} = \mathbf{Q}\mathbf{x}$. Output $\mathbf{v}$ if $0 < \|\mathbf{v}\|_2 \leq \gamma(\gamma_{\text{det}}\sigma + \sigma_0)$, and $\perp$ otherwise.

The algorithm that constructs a succinct SIS instance in Lemma 4.5 guarantees that $q\mathbf{L}^{-1}\mathbf{Q}$ is integral and congruent to $\mathbf{A}$ modulo q . Thus every accepted solver output satisfies $\mathbf{L}^{-1}\mathbf{Q}\mathbf{x} \in \mathbb{Z}^n$, so $\mathbf{v} \in \Lambda$. The last step checks that $\mathbf{v}$ is nonzero and satisfies the norm bound.

Let b be λ plus the total binary encoding length of $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{samp}}, \sigma)$ and the fixed numerical parameters. The algorithm that constructs a succinct SIS instance in Lemma 4.5 has running time and output bit length polynomial in b . There is at most one solver call, and the remaining matrix operations and checks have polynomial bit complexity. Hence the total running time is $\text{poly}(b) = \text{poly}(\lambda)$ under the standing assumptions.

Let $\mathcal{E}$ be the event that the solver returns a valid solution. By Lemma 4.5, the public triple is negligibly close to the succinct SIS input distribution, so $\Pr[\mathcal{E}] \geq \delta - \text{negl}(\lambda)$. By Lemmas 4.6 and 4.7, the probability that $\mathcal{E}$ occurs but $\mathbf{Q}\mathbf{x}$ is zero or has norm greater than $\gamma(\gamma_{\text{det}}\sigma + \sigma_0)$ is negligible. Thus the algorithm succeeds with probability at least $\delta - \text{negl}(\lambda)$. For any fixed proper subspace $\mathcal{H}$, Lemma 4.8 bounds $\Pr[\mathcal{E} \text{ and } \mathbf{Q}\mathbf{x} \notin \mathcal{H}]$. Subtracting the probability that $\mathcal{E}$ occurs and $\|\mathbf{Q}\mathbf{x}\|_2 > \gamma(\gamma_{\text{det}}\sigma + \sigma_0)$ shows that the probability of outputting a vector outside $\mathcal{H}$ that satisfies the norm bound is at least $\frac{1}{3}\Pr[\mathcal{E}] - \text{negl}(\lambda) \geq \delta/3 - \text{negl}(\lambda)$. For sufficiently large λ , this is at least $\delta/4$ because $\delta \geq 1/\text{poly}(\lambda)$. Since $0 \in \mathcal{H}$, any vector outside $\mathcal{H}$ is nonzero. $\square$

4.4 Converting the Hinted Instance to Succinct SIS

We now prove the two ingredients used by the one-vector theorem. The first constructs $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ while retaining $\mathbf{Q}$ with columns in Λ/q and $q\mathbf{L}^{-1}\mathbf{Q} = \mathbf{A} \pmod q$. The second analyzes $\mathbf{v} = \mathbf{Q}\mathbf{x}$, using the conditional distribution of the Gaussian columns after the solver's input and output have been fixed.

Lemma 4.5 (Construction of a succinct SIS instance). *Let the parameters, the instance $(\mathbf{L}, \mathbf{P})$, the sampling basis $\mathbf{L}_{\text{samp}}$, and σ satisfy the hypotheses of Theorem 4.3, without assuming a solver. There*

is a probabilistic polynomial-time algorithm that produces the public triple

$$\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \quad \mathbf{W} \in \mathbb{Z}_q^{\ell n \times k_\Gamma}, \quad \mathbf{T} \in \mathbb{Z}^{(\ell m + k_\Gamma) \times \ell n r_q},$$

where $r_q = \lceil \log_2 q \rceil$, and retains $\mathbf{Q} \in (\frac{1}{q}\Lambda)^m$ satisfying $q\mathbf{L}^{-1}\mathbf{Q} = \mathbf{A} \bmod q$. The public triple $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ is within negligible statistical distance of the input distribution of $\text{sSIS}_{n,m,q,\ell,k_\Gamma,s,\gamma}$.

Proof. By Lemma 3.22 and Equation (17), independently sample $\mathbf{v}_i \sim \mathcal{D}_{\Lambda/q,\sigma}$ for $i \in [m_0]$, using $\mathbf{L}_{\text{samp}}/q$ as the sampling basis, and set

$$\mathbf{V} := [\mathbf{v}_1 | \cdots | \mathbf{v}_{m_0}], \quad \mathbf{A}_0 := q\mathbf{L}^{-1}\mathbf{V} \bmod q. \quad (18)$$

By Section 3.5, compute its nontrivial invariant factors $d_1, \dots, d_h \geq 2$ and quotient generators $\mathbf{B} \in \mathbb{Q}^{\ell n \times h}$. The associated binary gadget $\mathbf{G}_\Gamma$ has h rows and width

$$k_\Gamma = \sum_{j=1}^h \lceil \log_2 d_j \rceil \leq 2 \sum_{j=1}^h \log_2 d_j = 2 \log_2 |\Gamma| \leq 2k, \quad (19)$$

where $\lceil \log_2 d_j \rceil \leq 2 \log_2 d_j$ since $d_j \geq 2$. For the trivial quotient, $h = k_\Gamma = 0$ and the same bound holds. Round each column of $\mathbf{B}\mathbf{G}_\Gamma$ to Λ_ℓ/q by nearest plane with the original basis $\mathbf{L}_\ell/q$, obtaining

$$\mathbf{B}' \in (\frac{1}{q}\Lambda_\ell)^{k_\Gamma}, \quad \|\mathbf{B}\mathbf{G}_\Gamma - \mathbf{B}'\|_{2,\infty} \leq \frac{\sqrt{\ell n}}{2q} \|\tilde{\mathbf{L}}\|_{2,\infty}. \quad (20)$$

Put $\mathbf{C} := q\mathbf{L}_\ell^{-1}\mathbf{B}' \in \mathbb{Z}^{\ell n \times k_\Gamma}$.

Encoding a target. For a target $\mathbf{u} \in \mathbb{Z}_q^{\ell n}$, let $\tilde{\mathbf{u}} \in \{0, \dots, q-1\}^{\ell n}$ be its canonical representative. Nearest plane with $\mathbf{P}$ gives $\mathbf{u}' \in \Lambda'$ such that $\mathbf{r} = \mathbf{L}_\ell \tilde{\mathbf{u}}/q - \mathbf{u}'$ has norm at most $\sqrt{\ell n}\beta/2$. Compute the quotient coordinate $\mathbf{y}$ of $\mathbf{u}'$ and its binary encoding $\mathbf{z}_3 \in \{0, 1\}^{k_\Gamma}$, so that $\mathbf{G}_\Gamma \mathbf{z}_3 = \mathbf{y}$ and $\mathbf{u}' - \mathbf{B}\mathbf{y} \in \Lambda_\ell$. Define

$$\mathbf{r}_{\text{corr}} := \mathbf{r} + (\mathbf{B}\mathbf{G}_\Gamma - \mathbf{B}')\mathbf{z}_3. \quad (21)$$

Writing $\mathbf{a} = \mathbf{u}' - \mathbf{B}\mathbf{y} \in \Lambda_\ell$ gives $\mathbf{r}_{\text{corr}} = \mathbf{L}_\ell \tilde{\mathbf{u}}/q - \mathbf{a} - \mathbf{B}'\mathbf{z}_3$. Consequently,

$$\mathbf{r}_{\text{corr}} \in \frac{1}{q}\Lambda_\ell, \quad \mathbf{C}\mathbf{z}_3 + q\mathbf{L}_\ell^{-1}\mathbf{r}_{\text{corr}} = \mathbf{u} \bmod q. \quad (22)$$

The nearest-plane bounds, $k_\Gamma \leq 2k$, $\|\tilde{\mathbf{L}}\|_{2,\infty} \leq \|\mathbf{L}\|$, $\beta \geq 1$, and Equation (11) give directly

$$\|\mathbf{r}_{\text{corr}}\|_2 \leq \frac{\sqrt{\ell n}}{2} \left(\beta + \frac{k_\Gamma}{q} \|\tilde{\mathbf{L}}\|_{2,\infty} \right) \leq \frac{19}{34} \beta \sqrt{\ell n} < \beta \sqrt{\ell n}.$$

Set

$$\mathbf{K} := \left\lfloor \frac{q}{p} \mathbf{L}^{-1} \right\rfloor, \quad \Delta_K := \mathbf{K} - \frac{q}{p} \mathbf{L}^{-1}, \quad \|\Delta_K\| \leq \|\Delta_K\|_F \leq \frac{n}{2}. \quad (23)$$

For the corrected residual, compute $\mathbf{z}_1 = \lfloor p\mathbf{r}_{\text{corr}} \rfloor \in \mathbb{Z}^{\ell n}$ and $\mathbf{e}_z = \mathbf{z}_1 - p\mathbf{r}_{\text{corr}}$. Then $\|\mathbf{e}_z\|_2 \leq \sqrt{\ell n}/2$ and $\|\mathbf{z}_1\|_2 \leq 2p\beta\sqrt{\ell n}$, since $\|\mathbf{r}_{\text{corr}}\|_2 < \beta\sqrt{\ell n}$ and $p\beta \geq 1$. The remaining error is the integer vector

$$\begin{aligned} \bar{\mathbf{z}} &:= q\mathbf{L}_\ell^{-1}\mathbf{r}_{\text{corr}} - (\mathbf{I}_\ell \otimes \mathbf{K})\mathbf{z}_1 \\ &= -\frac{q}{p}\mathbf{L}_\ell^{-1}\mathbf{e}_z - (\mathbf{I}_\ell \otimes \Delta_K)\mathbf{z}_1. \end{aligned} \quad (24)$$

Its norm is at most $q\|\mathbf{L}^{-1}\|\sqrt{\ell n}/(2p) + p\beta\sqrt{\ell} n^{3/2} \leq 2^{r_M-1} \leq M$. The signed binary decomposition from Section 3.4 therefore gives $\mathbf{z}_2 \in \{-1, 0, 1\}^{\ell nr_M}$ such that

$$q\mathbf{L}_\ell^{-1}\mathbf{r}_{\text{corr}} = (\mathbf{I}_\ell \otimes \mathbf{K})\mathbf{z}_1 + (\mathbf{I}_\ell \otimes \mathbf{G}_{n,M})\mathbf{z}_2. \quad (25)$$

The bound M is computed from the original input basis and remains unchanged throughout basis improvement.

Publishing the masked encoders. Independently of $\mathbf{V}$, sample

$$\mathbf{R}_1 \sim \{-1, 1\}^{m_0 \times n}, \quad \mathbf{R}_2 \sim \{-1, 1\}^{m_0 \times nr_M}, \quad \mathbf{S} \sim \{-1, 1\}^{\ell m_0 \times k_\Gamma}.$$

Publish the matrices

$$\mathbf{A} := [\mathbf{A}_0 \mid \mathbf{A}_0\mathbf{R}_1 + \mathbf{K} \mid \mathbf{A}_0\mathbf{R}_2 + \mathbf{G}_{n,M}], \quad (26)$$

$$\mathbf{W} := \mathbf{C} + (\mathbf{I}_\ell \otimes \mathbf{A}_0)\mathbf{S} \quad (27)$$

over $\mathbb{Z}_q$. Retain

$$\mathbf{Q} := [\mathbf{V} \mid \mathbf{V}\mathbf{R}_1 + \frac{1}{q}\mathbf{L}\mathbf{K} \mid \mathbf{V}\mathbf{R}_2 + \frac{1}{q}\mathbf{L}\mathbf{G}_{n,M}]. \quad (28)$$

All its columns lie in Λ/q , and

$$q\mathbf{L}^{-1}\mathbf{Q} = \mathbf{A} \bmod q. \quad (29)$$

If $\text{rank}(\mathbf{A}) < n$, discard the sampled state and restart with fresh randomness, beginning with the Gaussian columns. If all λ attempts fail, return fixed all-zero matrices of the dimensions in the statement. Otherwise, the tuple $(\mathbf{V}, \mathbf{R}_1, \mathbf{R}_2, \mathbf{S})$ is distributed as in a single attempt of the construction conditioned on $\text{rank}(\mathbf{A}) = n$.

Producing the Gaussian trapdoor. Let $\mathbf{D} = [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$. For each column $\mathbf{u}$ of $\mathbf{U} = \mathbf{G}_{\ell n, q}$, compute $\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3$ as above and set

$$\mathbf{z}_0 := -(\mathbf{I}_\ell \otimes \mathbf{R}_1)\mathbf{z}_1 - (\mathbf{I}_\ell \otimes \mathbf{R}_2)\mathbf{z}_2 - \mathbf{S}\mathbf{z}_3.$$

Let Π_ℓ be the column-permutation matrix characterized by

$$(\mathbf{I}_\ell \otimes \mathbf{A})\Pi_\ell = [\mathbf{I}_\ell \otimes \mathbf{A}_0 \mid \mathbf{I}_\ell \otimes (\mathbf{A}_0\mathbf{R}_1 + \mathbf{K}) \mid \mathbf{I}_\ell \otimes (\mathbf{A}_0\mathbf{R}_2 + \mathbf{G}_{n,M})].$$

Thus Π_ℓ places the blocks of $\mathbf{z}_0, \mathbf{z}_1, \mathbf{z}_2$ belonging to each copy of Λ consecutively, in the order expected by $\mathbf{I}_\ell \otimes \mathbf{A}$. For the target column $\mathbf{u}$, construct the preimage

$$\mathbf{z} := \begin{pmatrix} \Pi_\ell(\mathbf{z}_0^\top, \mathbf{z}_1^\top, \mathbf{z}_2^\top)^\top \\ \mathbf{z}_3 \end{pmatrix} \in \mathbb{Z}^{\ell m + k_\Gamma}. \quad (30)$$

The three masked products cancel in $\mathbf{D}\mathbf{z}$, leaving

$$\mathbf{D}\mathbf{z} = (\mathbf{I}_\ell \otimes \mathbf{K})\mathbf{z}_1 + (\mathbf{I}_\ell \otimes \mathbf{G}_{n,M})\mathbf{z}_2 + \mathbf{C}\mathbf{z}_3 = \mathbf{u} \bmod q$$

by Equations (22) and (25).

Every entry of $\mathbf{z}_0$ has absolute value at most $2p\beta n\sqrt{\ell} + nr_M + k_\Gamma$: a row of $\mathbf{R}_1$ has norm $\sqrt{n}$, the corresponding block of $\mathbf{z}_2$ has nr_M ternary entries, and $\mathbf{z}_3$ has $k_\Gamma \leq 2k$ binary entries by

Equation (19). This also bounds all other entries of $\mathbf{z}$. Collecting the preimages $\mathbf{z}$ constructed in Equation (30), one for each column of $\mathbf{U}$, gives $\widehat{\mathbf{T}} \in \mathbb{Z}^{(\ell m + k_\Gamma) \times \ell n r_q}$ satisfying

$$\mathbf{D}\widehat{\mathbf{T}} = \mathbf{G}_{\ell n, q} \bmod q.$$

Every entry of $\widehat{\mathbf{T}}$ has absolute value at most $2p\beta n\sqrt{\ell} + nr_M + k_\Gamma$. Since $\ell m + k_\Gamma \geq \ell m \geq 3\ell n \lceil \log_2 q \rceil$ and $k_\Gamma \leq 2k$, Equation (15) implies

$$s \geq (\ell m + k_\Gamma) (2p\beta n\sqrt{\ell} + nr_M + k_\Gamma) \log(\ell n).$$

Thus Lemma 3.6 applies with row dimension ℓn and column dimension $\ell m + k_\Gamma$. Apply $\text{SamplePre}(\mathbf{D}, \widehat{\mathbf{T}}, \mathbf{U}, s)$ independently columnwise to obtain the published trapdoor $\mathbf{T}$.

Let $\text{st} := (\mathbf{V}, \mathbf{R}_1, \mathbf{R}_2, \mathbf{S})$ denote the retained tuple. For every fixed value of st with $\text{rank}(\mathbf{A}) = n$, the trapdoor $\widehat{\mathbf{T}}$ constructed from these preimages satisfies $\mathbf{D}\widehat{\mathbf{T}} = \mathbf{U} \bmod q$ and the entry bound above. For comparison, define $\mathbf{T}_{\text{ideal}}$ by sampling its columns independently from $\mathbf{D}^{-1}(\mathbf{U}, s)$ with fresh coins. Thus Equation (5) gives

$$\Delta((\text{st}, \mathbf{T}), (\text{st}, \mathbf{T}_{\text{ideal}})) \leq \ell n r_q 2^{-\ell n}.$$

Set $\delta_{\text{samp}} := \ell n r_q 2^{-\ell n}$. This calculation uses exact initial samples. The conditional distribution of $\mathbf{T}_{\text{ideal}}$ depends on st only through $(\mathbf{A}, \mathbf{W})$.

The public distribution. We work with exact Gaussian samples, without conditioning on $\text{rank}(\mathbf{A}) = n$. Write $\mathbf{S}^{(1)}, \dots, \mathbf{S}^{(\ell)}$ for the $m_0 \times k_\Gamma$ block rows of $\mathbf{S}$. The following experiments output a pair of public matrices; changes from the preceding experiment are highlighted in blue.

$\text{Expt}_{\text{real}}$ is the distribution of $(\mathbf{A}, \mathbf{W})$ defined in Equations (26) and (27).

Expt_1 samples the masks as in the real experiment, samples $\mathbf{U}_0 \sim \mathbb{Z}_q^{n \times m_0}$ independently, and outputs

$$\left(\left[\begin{array}{c|c|c} \mathbf{U}_0 & \mathbf{U}_0 \mathbf{R}_1 + \mathbf{K} & \mathbf{U}_0 \mathbf{R}_2 + \mathbf{G}_{n, M} \end{array} \right], \mathbf{C} + \begin{bmatrix} \mathbf{U}_0 \mathbf{S}^{(1)} \\ \vdots \\ \mathbf{U}_0 \mathbf{S}^{(\ell)} \end{bmatrix} \right).$$

Expt_2 samples $\mathbf{U}_0$ and $\mathbf{S}$ as in Expt_1 , independently samples

$$\mathbf{U}_1 \sim \mathbb{Z}_q^{n \times n}, \quad \mathbf{U}_2 \sim \mathbb{Z}_q^{n \times nr_M},$$

and outputs

$$\left(\left[\begin{array}{c|c|c} \mathbf{U}_0 & \mathbf{U}_1 + \mathbf{K} & \mathbf{U}_2 + \mathbf{G}_{n, M} \end{array} \right], \mathbf{C} + \begin{bmatrix} \mathbf{U}_0 \mathbf{S}^{(1)} \\ \vdots \\ \mathbf{U}_0 \mathbf{S}^{(\ell)} \end{bmatrix} \right).$$

$\text{Expt}_{\text{rand}}$ samples $\mathbf{U}_0, \mathbf{U}_1, \mathbf{U}_2$ as above and samples $\mathbf{U}_W^{(j)} \sim \mathbb{Z}_q^{n \times k_\Gamma}$ for every $j \in [\ell]$. It outputs

$$\left(\left[\begin{array}{c|c|c} \mathbf{U}_0 & \mathbf{U}_1 + \mathbf{K} & \mathbf{U}_2 + \mathbf{G}_{n, M} \end{array} \right], \mathbf{C} + \begin{bmatrix} \mathbf{U}_W^{(1)} \\ \vdots \\ \mathbf{U}_W^{(\ell)} \end{bmatrix} \right).$$

We compare consecutive experiments as follows.

1. By Lemma 3.20, each column of $\mathbf{A}_0$ is within distance $\varepsilon/2$ of uniform because $\sigma \geq \sigma_0 \geq \eta_{\varepsilon/4}(\mathbf{A})$. Replacing its m_0 columns one at a time by independent uniform columns and summing the statistical-distance bounds gives a total distance of at most $m_0\varepsilon/2$. Sampling $\mathbf{R}_1, \mathbf{R}_2, \mathbf{S}$ independently and computing $(\mathbf{A}, \mathbf{W})$ by the same formulas in both distributions cannot increase this distance. Thus

$$\Delta(\text{Expt}_{\text{real}}, \text{Expt}_1) \leq \frac{m_0\varepsilon}{2}.$$

2. Apply Lemma 3.23 to $(\mathbf{U}_0, \mathbf{U}_0\mathbf{R}_1, \mathbf{U}_0\mathbf{R}_2)$, whose two masked matrices have $n + nr_M$ columns in total. In both distributions, sample $\mathbf{S}$ independently, compute $\mathbf{W} = \mathbf{C} + (\mathbf{I}_\ell \otimes \mathbf{U}_0)\mathbf{S}$, and add $\mathbf{K}$ and $\mathbf{G}_{n,M}$ to the respective blocks of the first matrix. These identical operations cannot increase statistical distance, so

$$\Delta(\text{Expt}_1, \text{Expt}_2) \leq \frac{n + nr_M}{2} \sqrt{q^n 2^{-m_0}}.$$

3. Apply the same lemma to $(\mathbf{U}_0, \mathbf{U}_0\mathbf{S}^{(1)}, \dots, \mathbf{U}_0\mathbf{S}^{(\ell)})$. The products $\mathbf{U}_0\mathbf{S}^{(j)}$ have ℓk_Γ columns in total. In both distributions, append independent uniform matrices $\mathbf{U}_1, \mathbf{U}_2$, form $[\mathbf{U}_0 \mid \mathbf{U}_1 + \mathbf{K} \mid \mathbf{U}_2 + \mathbf{G}_{n,M}]$, and add $\mathbf{C}$ to the stacked blocks of the second matrix. These identical operations cannot increase statistical distance, so

$$\Delta(\text{Expt}_2, \text{Expt}_{\text{rand}}) \leq \frac{\ell k_\Gamma}{2} \sqrt{q^n 2^{-m_0}}.$$

The two matrices in $\text{Expt}_{\text{rand}}$ are independent and uniform. Since $n + nr_M + \ell k_\Gamma \leq 2nr_M + 2\ell k$, the triangle inequality gives

$$\Delta(\text{Expt}_{\text{real}}, \text{Expt}_{\text{rand}}) \leq \delta_{\text{pub}}, \quad \delta_{\text{pub}} := \frac{m_0\varepsilon}{2} + (nr_M + \ell k) \sqrt{q^n 2^{-m_0}}.$$

Equivalently,

$$\Delta((\mathbf{A}, \mathbf{W}), \text{Unif}(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{\ell n \times k_\Gamma})) \leq \delta_{\text{pub}}.$$

A uniform $n \times m$ matrix has rank less than n with probability at most q^{n-m} . Hence $\text{rank}(\mathbf{A}) = n$ has probability at least $\alpha := 1 - q^{n-m} - \delta_{\text{pub}}$, and conditioning both distributions on $\text{rank}(\mathbf{A}) = n$ gives distance at most $2\delta_{\text{pub}}/\alpha$. Set $\delta_{\text{rank}} := 2\delta_{\text{pub}}/\alpha$. In both distributions conditioned on $\text{rank}(\mathbf{A}) = n$, append a matrix whose columns are sampled independently from $\mathbf{D}^{-1}(\mathbf{U}, s)$, where $\mathbf{D} = [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$. Using the same conditional distribution given $(\mathbf{A}, \mathbf{W})$ does not increase statistical distance, so the public triples are within statistical distance δ_{rank} . This comparison concerns only the public triples. The Gaussian columns of $\mathbf{V}$ and the matrices $\mathbf{R}_1, \mathbf{R}_2, \mathbf{S}$ retain their original joint distribution conditioned on $\text{rank}(\mathbf{A}) = n$ for the analysis of $\mathbf{Q}\mathbf{x}$.

The parameter choices imply that δ_{pub} , δ_{rank} , and δ_{samp} are negligible and that $\alpha \geq 1/2$ for all sufficiently large λ . Finally, quotient decomposition, nearest plane, and the integer encoding operations have polynomial bit complexity. The probability that $\text{rank}(\mathbf{A}) = n$ in one attempt is at least $1/2$ for sufficiently large λ , so the probability that all λ attempts fail is at most $2^{-\lambda}$. Approximate Gaussian sampling and numerical approximation contribute only negligible statistical error. Together with the negligible probability of using the all-zero fallback, this proves the claim. $\square$

4.5 Lifting a Succinct SIS Solution

Throughout this subsection, use the parameters and the algorithm that constructs a succinct SIS instance in Lemma 4.5. The parameters are summarized in Table 1. For $\mathbf{x} = (\mathbf{x}_0^\top, \mathbf{x}_1^\top, \mathbf{x}_2^\top)^\top \in \mathbb{Z}^{m_0} \times \mathbb{Z}^n \times \mathbb{Z}^{nr_M}$, define

$$\mathbf{v} := \mathbf{Q}\mathbf{x} = \mathbf{V}(\mathbf{x}_0 + \mathbf{R}_1\mathbf{x}_1 + \mathbf{R}_2\mathbf{x}_2) + \frac{1}{q}\mathbf{L}\mathbf{K}\mathbf{x}_1 + \frac{1}{q}\mathbf{L}\mathbf{G}_{n,M}\mathbf{x}_2. \quad (31)$$

If the construction returns its all-zero matrices, define $\mathbf{v} = \mathbf{Q}\mathbf{x} = 0$ directly. Whenever $\mathbf{A}\mathbf{x} = 0 \pmod q$, Equation (29) gives $q\mathbf{L}^{-1}\mathbf{Q}\mathbf{x} \in q\mathbb{Z}^n$, so $\mathbf{v} \in \Lambda$.

Fix any algorithm that, given only $(\mathbf{A}, \mathbf{W}, \mathbf{T})$, uses independent randomness to produce an answer. Let $\mathcal{E}$ be the event that this answer is a vector $\mathbf{x} \in \mathbb{Z}^m$ satisfying $\mathbf{A}\mathbf{x} = 0 \pmod q$ and $0 < \|\mathbf{x}\|_2 \leq \gamma$. On $\neg\mathcal{E}$, use $\mathbf{x} = 0$ when evaluating $\mathbf{Q}\mathbf{x}$.

Comparison Experiments. The following experiments output $(\text{st}, \mathbf{A}, \mathbf{W}, \mathbf{Q}, \mathbf{T})$, where $\text{st} = (\mathbf{V}, \mathbf{R}_1, \mathbf{R}_2, \mathbf{S})$ is the tuple retained by the construction. Each experiment then runs the fixed algorithm on its public triple with independent randomness and records $\mathbf{x}$ and $\mathcal{E}$ as above. Changes are highlighted in blue.

Expt₀ runs the algorithm that constructs a succinct SIS instance in Lemma 4.5, with input $(\mathbf{L}, \mathbf{P}, \mathbf{L}_{\text{sam}}, \sigma)$, including its maximum of λ attempts. If all attempts fail, it returns the prescribed all-zero matrices and sets $\text{st} = \perp$.

Expt₁ uses **exact initial samples from $\mathcal{D}_{\Lambda/q, \sigma}$ in a single attempt of the construction, conditioned on $\text{rank}(\mathbf{A}) = n$** . It computes $\mathbf{A}, \mathbf{W}, \mathbf{Q}$ by the same formulas and produces $\mathbf{T}$ using SamplePre as in Lemma 4.5.

Expt₂ samples $(\text{st}, \mathbf{A}, \mathbf{W}, \mathbf{Q})$ as in Expt₁, sets $\mathbf{D} = [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$ and replaces $\mathbf{T}$ by $\mathbf{T}_{\text{ideal}} \sim \mathbf{D}^{-1}(\mathbf{G}_{\ell n, q}, \mathbf{s})$, sampling the columns independently with fresh randomness.

We compare consecutive experiments, retaining all the listed matrices.

1. Replacing the approximate initial Gaussian samples by exact samples changes the joint distribution by $\text{negl}(\lambda)$ over at most λ attempts. With exact samples, the first successful attempt has the single-attempt distribution conditioned on $\text{rank}(\mathbf{A}) = n$. The probability that all attempts fail is at most $2^{-\lambda}$ by Lemma 4.5. Including the numerical approximation errors therefore gives $\Delta(\text{Expt}_0, \text{Expt}_1) \leq \text{negl}(\lambda)$.
2. The joint sampling bound in the proof of Lemma 4.5 gives

$$\Delta(\text{Expt}_1, \text{Expt}_2) \leq \ell n \lceil \log_2 q \rceil 2^{-\ell n} = \text{negl}(\lambda).$$

The matrices $\mathbf{A}, \mathbf{W}, \mathbf{Q}$ are functions of st , so the bound for $(\text{st}, \mathbf{T})$ also holds jointly with these matrices.

Running the same algorithm with independent randomness and evaluating $\mathbf{x}$ and $\mathcal{E}$ cannot increase either distance. Write $\Pr_j$ for probability in Expt_j. For every event $\mathcal{F}$ determined by the retained matrices and algorithm's answer, the triangle inequality gives

$$|\Pr_0[\mathcal{F}] - \Pr_2[\mathcal{F}]| \leq \text{negl}(\lambda). \quad (32)$$

Unsubscripted probabilities in the following statements refer to Expt₀.

Lemma 4.6 (Norm of the lattice vector). *Except with negligible probability, simultaneously for every $\mathbf{x} \in \mathbb{R}^m$ with $\|\mathbf{x}\|_2 \leq \gamma$,*

$$\|\mathbf{Q}\mathbf{x}\|_2 \leq \gamma(\gamma_{\det}\sigma + \sigma_0).$$

Proof. We first bound $\|\mathbf{Q}\|$ in Expt_2 . Before conditioning on $\text{rank}(\mathbf{A}) = n$, Lemma 3.20 gives $\|\mathbf{V}\| \leq \|\mathbf{V}\|_F \leq \sigma\sqrt{m_0n}$ except with probability m_02^{-n+1} . For every choice of $\mathbf{R}_1, \mathbf{R}_2$,

$$\|[\mathbf{I}_{m_0} \mid \mathbf{R}_1 \mid \mathbf{R}_2]\| \leq 2\sqrt{m_0nr_M}.$$

Whenever this bound on $\|\mathbf{V}\|$ holds, $\|\mathbf{V}[\mathbf{I}_{m_0} \mid \mathbf{R}_1 \mid \mathbf{R}_2]\| \leq 2m_0n\sqrt{r_M}\sigma = \gamma_{\det}\sigma$.

Using $\mathbf{L}\mathbf{K}/q = \mathbf{I}_n/p + \mathbf{L}\Delta_K/q$ and $\|\mathbf{G}_{n,M}\| < 2^{r_M}$ in Equation (28) gives

$$\begin{aligned} \|\mathbf{Q}\| &\leq \gamma_{\det}\sigma + \frac{1}{p} + \frac{n\|\mathbf{L}\|}{2q} + \frac{2^{r_M}\|\mathbf{L}\|}{q} \\ &< \gamma_{\det}\sigma + \frac{1 + 2\|\mathbf{L}\|\|\mathbf{L}^{-1}\|\sqrt{\ell n}}{p} + \frac{\|\mathbf{L}\|(n + 8p\beta\sqrt{\ell}n^{3/2})}{2q} \\ &\leq \gamma_{\det}\sigma + \frac{3\|\mathbf{L}\|\|\mathbf{L}^{-1}\|\sqrt{\ell n}}{p} + \frac{9p\beta\sqrt{\ell}n^{3/2}\|\mathbf{L}\|}{2q} \\ &\leq \gamma_{\det}\sigma + \frac{\sigma_0}{2} + \frac{9\sigma_0}{34} < \gamma_{\det}\sigma + \sigma_0. \end{aligned}$$

The second line uses the definition of r_M . The third uses $\|\mathbf{L}\|\|\mathbf{L}^{-1}\|\sqrt{\ell n} \geq 1$ and $p\beta\sqrt{\ell}n^{3/2} \geq n$. The fourth uses Equations (10) and (11).

The event $\text{rank}(\mathbf{A}) = n$ has probability at least $1/2$ for sufficiently large λ . Conditioning on it increases the probability that $\|\mathbf{V}\| > \sigma\sqrt{m_0n}$ by at most a factor of two, giving $2m_02^{-n+1} = \text{negl}(\lambda)$ because $n = \omega(\log \lambda)$ and m_0 is polynomially bounded. By Equation (32), $\|\mathbf{Q}\| \leq \gamma_{\det}\sigma + \sigma_0$ also holds in Expt_0 except with negligible probability. Multiplying by $\|\mathbf{x}\|_2 \leq \gamma$ proves the simultaneous bound. $\square$

Lemma 4.7 (Probability that the lattice vector is zero). *For the algorithm and event $\mathcal{E}$ defined above,*

$$\Pr[\mathcal{E} \text{ and } \mathbf{Q}\mathbf{x} = 0] \leq \text{negl}(\lambda).$$

Proof. In Expt_2 , we first bound the probability that $\mathcal{E}$ occurs and $\mathbf{x}_0 + \mathbf{R}_1\mathbf{x}_1 + \mathbf{R}_2\mathbf{x}_2 = 0$. If $\mathbf{x}_1 = \mathbf{x}_2 = 0$, this coefficient vector is $\mathbf{x}_0$, which is nonzero on $\mathcal{E}$. Otherwise, some coordinate of $\mathbf{x}_1$ or $\mathbf{x}_2$ is nonzero.

Fix a column position j of $\mathbf{R}_1$ and consider the event that $\mathcal{E}$ occurs, $x_{1,j} \neq 0$, and the coefficient vector is zero. Before conditioning on $\text{rank}(\mathbf{A}) = n$, fix $\mathbf{A}_0, \mathbf{S}$, all columns of $\mathbf{R}_2$, all columns of $\mathbf{R}_1$ other than $(\mathbf{R}_1)_{*,j}$, the randomness used to sample $\mathbf{T}_{\text{ideal}}$, and the algorithm's randomness. On inputs with $\text{rank}(\mathbf{A}) < n$, declare $\mathcal{E}$ false and set $\mathbf{x} = 0$. Otherwise, the solver's input $(\mathbf{A}, \mathbf{W}, \mathbf{T}_{\text{ideal}})$ depends on the uniform sign column $(\mathbf{R}_1)_{*,j}$ only through $\mathbf{A}_0(\mathbf{R}_1)_{*,j}$. For each fixed value of this product whose output satisfies $\mathcal{E}$ and $x_{1,j} \neq 0$, the equation

$$\mathbf{x}_0 + \sum_{t \neq j} (\mathbf{R}_1)_{*,t} x_{1,t} + \mathbf{R}_2 \mathbf{x}_2 + (\mathbf{R}_1)_{*,j} x_{1,j} = 0$$

has at most one solution for $(\mathbf{R}_1)_{*,j}$. There are at most q^n products, and each sign column has probability 2^{-m_0} , so the event has probability at most $q^n 2^{-m_0}$. The same argument applies to each column of $\mathbf{R}_1$ and $\mathbf{R}_2$. A union bound over their $n + nr_M \leq 2nr_M$ columns, followed by conditioning on $\text{rank}(\mathbf{A}) = n$, therefore gives

$$\Pr_2[\mathcal{E} \text{ and } \mathbf{x}_0 + \mathbf{R}_1 \mathbf{x}_1 + \mathbf{R}_2 \mathbf{x}_2 = 0] \leq 4nr_M q^n 2^{-m_0} = \text{negl}(\lambda). \quad (33)$$

Here conditioning costs at most a factor of two, and negligibility follows from the choice of m_0 in Equation (13).

Next, on $\mathcal{E}$ with $\mathbf{x}_0 + \mathbf{R}_1 \mathbf{x}_1 + \mathbf{R}_2 \mathbf{x}_2 \neq 0$, choose the first index i for which $(\mathbf{x}_0 + \mathbf{R}_1 \mathbf{x}_1 + \mathbf{R}_2 \mathbf{x}_2)_i \neq 0$. Condition on $\mathbf{A}_0, \mathbf{R}_1, \mathbf{R}_2, \mathbf{S}, \mathbf{T}_{\text{ideal}}$, the algorithm's randomness and output, and all Gaussian columns except $\mathbf{v}_i$. These values determine $\mathcal{E}$, the coefficient vector, and the index i . The conditional distribution of $\mathbf{T}_{\text{ideal}}$ given $(\mathbf{V}, \mathbf{R}_1, \mathbf{R}_2, \mathbf{S})$ depends only on $(\mathbf{A}, \mathbf{W})$. These public matrices and the event $\text{rank}(\mathbf{A}) = n$ depend on $\mathbf{V}$ only through $\mathbf{A}_0$. Thus the fixed variables restrict $\mathbf{v}_i$ only to its residue class $\mathbf{a}_i + \Lambda$, and its conditional distribution is

$$\Pr_2[\mathbf{v}_i = \mathbf{w} \mid \text{the fixed variables}] = \frac{e^{-\pi \|\mathbf{w}\|_2^2 / \sigma^2} \mathbf{1}[\mathbf{w} \in \mathbf{a}_i + \Lambda]}{\rho_\sigma(\mathbf{a}_i + \Lambda)}. \quad (34)$$

Here $\rho_\sigma(X) = \sum_{\mathbf{w} \in X} e^{-\pi \|\mathbf{w}\|_2^2 / \sigma^2}$.

Under this conditioning, Equation (31) expresses $\mathbf{Q}\mathbf{x} = 0$ as an affine equation in $\mathbf{v}_i$ whose coefficient $(\mathbf{x}_0 + \mathbf{R}_1 \mathbf{x}_1 + \mathbf{R}_2 \mathbf{x}_2)_i$ is nonzero. It therefore forces $\mathbf{v}_i$ to at most one point. Since $\sigma \geq \sigma_0 \geq 2\eta_\varepsilon(\Lambda)$, Lemma 3.20 bounds this conditional probability by $2^{-n}(1 + \varepsilon)/(1 - \varepsilon)$. Averaging and adding Equation (33) gives

$$\Pr_2[\mathcal{E} \text{ and } \mathbf{Q}\mathbf{x} = 0] \leq 2^{-n} \frac{1 + \varepsilon}{1 - \varepsilon} + \text{negl}(\lambda) = \text{negl}(\lambda).$$

The last equality uses $n = \omega(\log \lambda)$. Applying Equation (32) to this event proves the claim in Expt_0 . $\square$

Lemma 4.8 (Probability of a lattice vector outside a fixed subspace). *For the algorithm and event $\mathcal{E}$ defined above and every fixed proper linear subspace $H \subsetneq \mathbb{R}^n$,*

$$\Pr[\mathcal{E} \text{ and } \mathbf{Q}\mathbf{x} \notin H] \geq \frac{1}{3} \Pr[\mathcal{E}] - \text{negl}(\lambda).$$

Proof. In Expt_2 , condition on the variables specified in the proof of Lemma 4.7, on $\mathcal{E}$ with $\mathbf{x}_0 + \mathbf{R}_1 \mathbf{x}_1 + \mathbf{R}_2 \mathbf{x}_2 \neq 0$. By Equation (31), the condition $\mathbf{Q}\mathbf{x} \in H$ restricts $\mathbf{v}_i$ to a proper affine subspace because $(\mathbf{x}_0 + \mathbf{R}_1 \mathbf{x}_1 + \mathbf{R}_2 \mathbf{x}_2)_i \neq 0$. Applying Lemma 3.21 to the conditional distribution in Equation (34) gives

$$\Pr_2[\mathbf{Q}\mathbf{x} \notin H \mid \text{the fixed variables}] \geq 1 - \frac{1 + \varepsilon}{2(1 - \varepsilon)} = \frac{1 - 3\varepsilon}{2(1 - \varepsilon)} > \frac{1}{3},$$

since $\varepsilon = 2^{-\lambda} \leq 1/16$.

Average over the fixed variables on the stated event and use Equation (33). This gives

$$\begin{aligned}\Pr_2[\mathcal{E} \text{ and } \mathbf{Q}\mathbf{x} \notin H] &\geq \frac{1}{3}(\Pr_2[\mathcal{E}] - \Pr_2[\mathcal{E} \text{ and } \mathbf{x}_0 + \mathbf{R}_1\mathbf{x}_1 + \mathbf{R}_2\mathbf{x}_2 = 0]) \\ &\geq \frac{1}{3}\Pr_2[\mathcal{E}] - \text{negl}(\lambda).\end{aligned}$$

Finally, apply Equation (32) to both $\mathcal{E} \cap \{\mathbf{Q}\mathbf{x} \notin H\}$ and $\mathcal{E}$:

$$\begin{aligned}\Pr_0[\mathcal{E} \text{ and } \mathbf{Q}\mathbf{x} \notin H] &\geq \Pr_2[\mathcal{E} \text{ and } \mathbf{Q}\mathbf{x} \notin H] - \text{negl}(\lambda) \\ &\geq \frac{1}{3}\Pr_2[\mathcal{E}] - \text{negl}(\lambda) \\ &\geq \frac{1}{3}\Pr_0[\mathcal{E}] - \text{negl}(\lambda).\end{aligned}$$

□

4.6 The Opposite Implication

In this subsection, we reduce average-case succinct SIS to worst-case SuperSVP with the parameter changes stated below. We first bound the smoothing parameter of the lattice associated with the public hint.

Lemma 4.9. *Let q be an odd prime, let $n, m, k, \ell \geq 1$ be polynomially bounded in λ , and suppose $m \geq 2n \log q + 2\lambda$. For every $\epsilon > 0$, with probability at least $1 - 2^{-\lambda}$ over $\mathbf{A} \sim \mathbb{Z}_q^{n \times m}$, the following holds simultaneously for every $\mathbf{W} \in \mathbb{Z}_q^{\ell n \times k}$:*

$$\eta_\epsilon(\Lambda^\perp(\mathbf{D})) \leq 4\sqrt{\frac{\ln(2(\ell m + k)(1 + 1/\epsilon))}{\pi}}, \quad \mathbf{D} = [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}].$$

Proof. For every fixed $\mathbf{x} \in \mathbb{Z}_q^n \setminus \{0\}$, the vector $\mathbf{A}^\top \mathbf{x}$ is uniform over $\mathbb{Z}_q^m$. The interval $(-q/4, q/4)$ contains $2\lceil q/4 \rceil - 1$ integers, at most $2q/3$ for every odd prime q . Hence

$$\Pr[\exists \mathbf{y} \in \mathbb{Z}^m : \|\mathbf{A}^\top \mathbf{x} + q\mathbf{y}\|_\infty < q/4] \leq (2/3)^m.$$

Taking a union bound over the $q^n - 1$ nonzero residue classes gives an exceptional probability at most $q^n(2/3)^m < 2^{n \log q - m/2} \leq 2^{-\lambda}$. Outside this event, $\mathbf{A}$ has full row rank over $\mathbb{Z}_q$ and $\lambda_1^\infty(\Lambda(\mathbf{A}^\top)) \geq q/4$. Indeed, a nonzero kernel element would yield the zero vector in the preceding event, while a nonzero vector in $q\mathbb{Z}^m$ has norm at least q .

Write $\mathbf{W}$ as vertical blocks $\mathbf{W}_1, \dots, \mathbf{W}_\ell \in \mathbb{Z}_q^{n \times k}$. Every vector of $\Lambda(\mathbf{D}^\top)$ has the form

$$\begin{bmatrix} \mathbf{A}^\top \mathbf{x}_1 + q\mathbf{y}_1 \\ \vdots \\ \mathbf{A}^\top \mathbf{x}_\ell + q\mathbf{y}_\ell \\ \sum_{i=1}^\ell \mathbf{W}_i^\top \mathbf{x}_i + q\mathbf{y} \end{bmatrix}, \quad \mathbf{x}_i \in \mathbb{Z}^n, \quad \mathbf{y}_i \in \mathbb{Z}^m, \quad \mathbf{y} \in \mathbb{Z}^k.$$

If some $\mathbf{x}_i$ is nonzero modulo q , its corresponding block has maximum norm at least $q/4$. Otherwise, every coordinate is a multiple of q , so any nonzero vector has maximum norm at least q . Consequently, $\lambda_1^\infty(\Lambda(\mathbf{D}^\top)) \geq q/4$, uniformly over $\mathbf{W}$. The duality identity $(\Lambda^\perp(\mathbf{D}))^* = q^{-1}\Lambda(\mathbf{D}^\top)$ and Lemma 3.19 now give the claimed bound. □

Theorem 4.10 (The hardness of succinct SIS implies that of superlattice SVP). *Let q be an odd prime, let $n, m, k, k', \ell \geq 1$ be $\text{poly}(\lambda)$ -bounded integers where $m\ell + k = \Omega(\lambda)$, and $\gamma, \beta, \sigma > 0$ reals such that*

- (i) $\beta \geq \frac{\sqrt{m\ell+3}}{2} \cdot \left(1 + \sigma \cdot n\ell \lceil \log q \rceil \cdot \sqrt{m\ell + k}\right) \cdot \sqrt{n\ell \lceil \log q \rceil + 5}$
- (ii) $m \geq 2n \log q + 2\lambda$
- (iii) $\sigma \geq 2\sqrt{\lambda + \log(m\ell + k) + 1}$
- (iv) $k' \geq k \log q$

Then, every probabilistic polynomial-time solver for superlattice SVP with parameters $(m, \ell, k', \beta, \gamma)$ yields a probabilistic polynomial-time solver for $\text{sSIS}_{n,m,q,\ell,k,\sigma,\gamma}$ with only a $2^{-\Omega(\lambda)}$ additive loss in success probability.

Proof. We show a reduction from succinct SIS to the worst-case hardness of superlattice SVP. Suppose that the challenger provided us with $(\mathbf{A}, \mathbf{W}, \mathbf{T}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{n\ell \times k} \times \mathbb{Z}^{(\ell m + k) \times n\ell \lceil \log q \rceil}$ such that

$$\underbrace{[\mathbf{I}_\ell \otimes \mathbf{A} | \mathbf{W}]}_{=: \mathbf{D}} \cdot \mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q, \quad \|\mathbf{T}\|_{2,\infty} \leq \sigma \sqrt{m\ell + k}.$$

By the tail bound in Lemma 3.20, this event fails with probability at most $n\ell \lceil \log q \rceil 2^{-m\ell - k + 1} + 2^{1-\lambda}$. The second term allows for conditioning $\mathbf{A}$ on full row rank: Lemma 4.9 fails with unconditional probability at most $2^{-\lambda}$, and full row rank has probability at least $1/2$. Indeed, the smoothing parameter $\eta_{2^{-\lambda}}(\Lambda^\perp(\mathbf{D})) \leq 2\sqrt{\lambda + \log(m\ell + k) + 1} \leq \sigma$ with conditional probability at least $1 - 2^{1-\lambda}$ (see Lemma 4.9).

Let $\pi : \mathbb{R}^{m\ell+k} \rightarrow \mathbb{R}^{m\ell}$ be the projection onto the first ℓm coordinates of the input. Define the lattices $\Lambda' := \pi(\Lambda^\perp(\mathbf{D}))$, $\Lambda := \Lambda^\perp(\mathbf{A})$. Notice that $\pi(\Lambda^\perp(\mathbf{D}))$ is a lattice as the identity

$$\forall \mathbf{a}, \mathbf{b} \in \Lambda^\perp(\mathbf{D}) : \quad \pi(\mathbf{a} + \mathbf{b}) = \pi(\mathbf{a}) + \pi(\mathbf{b})$$

proves that it is an additive subgroup of $\mathbb{Z}^{m\ell}$.

The reduction proceeds by using the Micciancio–Peikert trapdoor $\mathbf{T}$ to generate a set $\mathbf{S}$ of generators for $\Lambda^\perp(\mathbf{D})$ having Euclidean norm at most $(1 + \sigma \cdot n\ell \lceil \log q \rceil \cdot \sqrt{m\ell + k}) \cdot \sqrt{n\ell \lceil \log q \rceil + 5}$ (this can be done efficiently by Lemma 3.26). Observe that $\pi(\mathbf{S})$ generates Λ' . Indeed, for every $\mathbf{v} \in \Lambda'$, there exists a vector $\mathbf{w} \in \Lambda^\perp(\mathbf{D})$ such that $\pi(\mathbf{w}) = \mathbf{v}$. Moreover, there exists an integer vector $\mathbf{c}$ such that $\mathbf{w} = \mathbf{S}\mathbf{c}$. From this we derive that $\mathbf{v} = \pi(\mathbf{w}) = \pi(\mathbf{S})\mathbf{c}$. Notice also that $\|\pi(\mathbf{S})\|_{2,\infty} \leq \|\mathbf{S}\|_{2,\infty}$.

Next, the reduction computes high-norm bases $\mathbf{L}$ and $\mathbf{L}'$ for $\Lambda = \Lambda^\perp(\mathbf{A})$ and $\Lambda' = \pi(\Lambda^\perp(\mathbf{D}))$, respectively. To derive $\mathbf{L}$, it is sufficient to apply Lemma 3.15 on $[q\mathbf{I}_m | \mathbf{R}]$ where the columns of $\mathbf{R}$ generate the kernel of $\psi_{\mathbf{A}} : \mathbb{Z}_q^m \rightarrow \mathbb{Z}_q^n$, $\psi_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\mathbf{x} \bmod q$ ($\mathbf{R}$ can be efficiently derived using Gaussian elimination over $\mathbb{Z}_q$). For $\mathbf{L}'$ it is sufficient to apply Lemma 3.15 on $\pi(\mathbf{S})$.

Finally, the reduction uses Lemma 3.25 on $\mathbf{L}'$ and $\pi(\mathbf{S})$ to derive a basis $\mathbf{P}$ for Λ' such that

$$\|\mathbf{P}\|_{2,\infty} \leq \frac{\sqrt{m\ell + 3}}{2} \cdot \left(1 + \sigma \cdot n\ell \lceil \log q \rceil \cdot \sqrt{m\ell + k}\right) \cdot \sqrt{n\ell \lceil \log q \rceil + 5} \leq \beta.$$

The reduction outputs the vector $\mathbf{x}$ produced by the superlattice SVP solver on input $(\mathbf{L}, \mathbf{P})$.

Observe that if the solver is successful, $\mathbf{x}$ is a nonzero vector of $\Lambda = \Lambda^\perp(\mathbf{A})$ such that $\|\mathbf{x}\|_2 \leq \gamma$. This means that $\mathbf{x} \in \mathbb{Z}^m$ and $\mathbf{A}\mathbf{x} \equiv \mathbf{0} \pmod{q}$.

It remains to show that $(\mathbf{L}, \mathbf{P})$ is a valid instance of $\text{SuperSVP}_{m, \ell, k', \beta, \gamma}$. The matrix $\mathbf{L}$ is a basis of Λ , while $\mathbf{P}$ is a basis of Λ' such that $\|\mathbf{P}\|_{2, \infty} \leq \beta$. It also holds that $\Lambda' \supseteq \Lambda^{\oplus \ell}$. Indeed, for any $\mathbf{v} \in \Lambda = \Lambda^\perp(\mathbf{A})$ and vector $\mathbf{e}_i$ in the canonical basis of $\mathbb{R}^\ell$, we have $\mathbf{e}_i \otimes \mathbf{v} = \pi(\mathbf{w})$ for

$$\mathbf{w} = \begin{bmatrix} \mathbf{e}_i \otimes \mathbf{v} \\ \mathbf{0} \end{bmatrix} \in \mathbb{Z}^{m\ell+k}.$$

Notice that $\mathbf{D}\mathbf{w} = \mathbf{e}_i \otimes \mathbf{A}\mathbf{v} \equiv \mathbf{0} \pmod{q}$, so $\mathbf{w} \in \Lambda^\perp(\mathbf{D})$. This proves that $\mathbf{e}_i \otimes \mathbf{v} \in \Lambda'$. Finally, let Λ_ℓ be the lattice generated by $\mathbf{I}_\ell \otimes \mathbf{L}$, define $\phi : \Lambda' \rightarrow \mathbb{Z}_q^{n\ell}$ to be the map that sends $\mathbf{v} \mapsto (\mathbf{I}_\ell \otimes \mathbf{A})\mathbf{v} \pmod{q}$. Notice that $\ker \phi = \Lambda_\ell$. We also have that $\text{Im} \phi \subseteq \mathbf{W}\mathbb{Z}_q^k$. Indeed, every $\mathbf{v} \in \Lambda'$ coincides with $\pi(\mathbf{w})$ for some $\mathbf{w} \in \Lambda^\perp(\mathbf{D}) \subseteq \mathbb{Z}^{m\ell+k}$. Now, if we denote the last k entries of $\mathbf{w}$ by $\mathbf{y} \in \mathbb{Z}^k$, we have that

$$\mathbf{0} = \mathbf{D}\mathbf{w} = (\mathbf{I}_\ell \otimes \mathbf{A})\mathbf{v} + \mathbf{W}\mathbf{y} = \phi(\mathbf{v}) + \mathbf{W}\mathbf{y} \pmod{q}.$$

This gives $\phi(\mathbf{v}) = -\mathbf{W}\mathbf{y} \in \mathbf{W}\mathbb{Z}_q^k$. By the homomorphism theorem, we obtain that

$$\Lambda' / \Lambda_\ell = \Lambda' / \ker \phi \cong \phi(\Lambda') \subseteq \mathbf{W}\mathbb{Z}_q^k.$$

We conclude that $|\Lambda' / \Lambda_\ell| \leq |\mathbf{W}\mathbb{Z}_q^k| \leq q^k = 2^{k \log q} \leq 2^{k'}$. □

5 Proof of Equivalence

In this section we show that decomposed LWE implies succinct LWE. A preliminary version of this result appeared in [AAC⁺26]. Our main theorem is the following.

Theorem 5.1 (Decomposed LWE implies Succinct LWE). *Let $\lambda \in \mathbb{N}$ be the security parameter and let*

$$\begin{aligned} n &= n(\lambda), & k &= k(\lambda), & \ell &= \ell(\lambda) \in \mathbb{N}, \\ \sigma_{\text{dec}} &= \sigma_{\text{dec}}(\lambda), & \sigma_{\text{suc}} &= \sigma_{\text{suc}}(\lambda), \\ \chi_{\text{dec}} &= \chi_{\text{dec}}(\lambda), & \chi_{\text{suc}} &= \chi_{\text{suc}}(\lambda) \in \mathbb{R}^+, \\ q &= q(\lambda) \text{ odd prime.} \end{aligned}$$

*Assume that $n, k, \ell, \log q, \sigma_{\text{dec}}, \sigma_{\text{suc}}, \chi_{\text{dec}}$, and χ_{suc} are efficiently computable and polynomially bounded in λ . Let $\text{Sample}(1^n)$ be any efficiently sampleable distribution over $\mathbb{Z}_q^n$. For either problem below, the qualifier *Sample* means that the secret is sampled from $\text{Sample}(1^n)$ instead of uniformly over $\mathbb{Z}_q^n$. Define $m := \ell^2 n \lceil \log q \rceil$ and suppose that*

$$\begin{aligned} n &= \omega(\log \lambda), & k &\geq 2n \lceil \log q \rceil, \\ \frac{q}{\sigma_{\text{dec}}} &= \omega(\sqrt{\log \lambda}), & \chi_{\text{dec}} &= \omega(\sqrt{\log \lambda}), \\ \sigma_{\text{dec}} &\geq 2, & \chi_{\text{suc}} &\geq 3m \chi_{\text{dec}}, \\ \sigma_{\text{suc}} &\geq 2(\ell^3 n \lceil \log q \rceil + k) \log(\ell n) \\ &\quad \cdot \sigma_{\text{dec}} \sqrt{k \ell n \lceil \log q \rceil}. \end{aligned}$$

Then every probabilistic polynomial-time distinguisher for (ℓ, Sample) -succinct LWE with parameters $(n, m, k, \chi_{\text{suc}}, \sigma_{\text{suc}}, q)$ yields a probabilistic polynomial-time distinguisher for (ℓ, Sample) -decomposed LWE with parameters $(n, k, \chi_{\text{dec}}, \sigma_{\text{dec}}, q)$, with only a negligible loss in distinguishing advantage.

We now give a proof of Theorem 5.1 by combining Lemmas 3.4, 3.6 and 5.8. We first describe the reduction and establish three statistical properties of its output.

Let $\mathcal{A}$ be a probabilistic polynomial-time distinguisher for (ℓ, Sample) -succinct LWE. We construct a probabilistic polynomial-time distinguisher $\mathcal{B}^{\mathcal{A}}$ for (ℓ, Sample) -decomposed LWE.

Write the decomposed LWE input as $(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top)$, where

$$\overline{\mathbf{W}} = [\mathbf{W}_1 \mid \cdots \mid \mathbf{W}_\ell] \in \mathbb{Z}_q^{n \times \ell k}, \quad \mathbf{R} = [\mathbf{R}_1 \mid \cdots \mid \mathbf{R}_\ell] \in \mathbb{Z}^{k \times \ell n \lceil \log q \rceil}.$$

The vector $\mathbf{b}^\top$ is either uniform over $\mathbb{Z}_q^m$ or is of the form $\mathbf{s}^\top \mathbf{A}_{\text{dec}} + \mathbf{e}^\top \bmod q$, where $\mathbf{s} \sim \text{Sample}(1^n)$ and $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m, \chi_{\text{dec}}}$.

The Reduction. On input $(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top)$, the distinguisher $\mathcal{B}^{\mathcal{A}}$ proceeds as follows.

1. If $\|\mathbf{R}\| > 2\sigma_{\text{dec}} \sqrt{k\ell n \lceil \log q \rceil}$, return an independent uniform bit and terminate.
2. Compute

$$\mathbf{A}_{\text{dec}} := \overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G} \in \mathbb{Z}_q^{n \times m}$$

and write $\mathbf{A}_{\text{dec}} = [\mathbf{A}_0 \mid \mathbf{A}_{\text{rest}}]$, where $\mathbf{A}_0 = \mathbf{W}_1 \mathbf{R}_1 + \mathbf{G}$.

3. Sample $\mathbf{X} \sim \{-1, 0, 1\}^{n \lceil \log q \rceil \times (m - n \lceil \log q \rceil)}$ and set

$$\mathbf{S} := \begin{bmatrix} \mathbf{I}_{n \lceil \log q \rceil} & \mathbf{X} \\ \mathbf{0} & \mathbf{I}_{m - n \lceil \log q \rceil} \end{bmatrix}, \quad \mathbf{Q} := \mathbf{S}^{-1} = \begin{bmatrix} \mathbf{I}_{n \lceil \log q \rceil} & -\mathbf{X} \\ \mathbf{0} & \mathbf{I}_{m - n \lceil \log q \rceil} \end{bmatrix}.$$

Define

$$\mathbf{A}^\star := \mathbf{A}_{\text{dec}} \mathbf{S}, \quad \mathbf{W}^\star := \begin{bmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_\ell \end{bmatrix}.$$

4. Partition $\mathbf{Q} = [\mathbf{Q}_1 \mid \cdots \mid \mathbf{Q}_\ell]$, where $\mathbf{Q}_i \in \mathbb{Z}^{m \times \ell n \lceil \log q \rceil}$, and define

$$\widehat{\mathbf{T}} := \begin{bmatrix} \mathbf{Q}_1 \\ \vdots \\ \mathbf{Q}_\ell \\ -\mathbf{R} \end{bmatrix}.$$

Let

$$\widetilde{\mathbf{D}}_\ell := [\mathbf{I}_\ell \otimes \mathbf{A}^\star \mid \mathbf{W}^\star].$$

Sample

$$\mathbf{T}^\star \leftarrow \text{SamplePre}(\widetilde{\mathbf{D}}_\ell, \widehat{\mathbf{T}}, \mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\text{suc}}).$$

5. Compute

$$\mathbf{c}^{\star\top} \leftarrow \text{ReRand} \left(\mathbf{S}, \mathbf{b}^\top, \chi_{\text{dec}}, \frac{\chi_{\text{suc}}}{2\chi_{\text{dec}}} \right).$$

6. Give $(\mathbf{A}^\star, \mathbf{c}^{\star\top}, \mathbf{W}^\star, \mathbf{T}^\star)$ to $\mathcal{A}$ and return its output.

We analyze the distribution given to $\mathcal{A}$.

The Abort Event. Let $\mathcal{E}$ denote the event that Step 1 does not terminate. Applying Lemma 3.3 to every entry of $\mathbf{R}$ with tail parameter $4k\ell n \lceil \log q \rceil$ and taking a union bound gives

$$\Pr[\neg \mathcal{E}] \leq k\ell n \lceil \log q \rceil 2^{-4k\ell n \lceil \log q \rceil} = \text{negl}(\lambda),$$

where the last equality follows from $n \lceil \log q \rceil = \omega(\log \lambda)$.

We now analyze the execution conditioned on $\mathcal{E}$. Removing this conditioning changes any distribution below by at most $\Pr[\neg \mathcal{E}]$. In particular, unconditional statistical statements remain valid by the triangle inequality, with only a negligible additional loss.

Distribution of the Trapdoor. Since $\mathbf{A}^\star \mathbf{Q} = \mathbf{A}_{\text{dec}}$, for every $i \in [\ell]$ we have

$$\mathbf{A}^\star \mathbf{Q}_i = [\mathbf{W}_i \mathbf{R}_1 + \delta_{i1} \mathbf{G} \mid \cdots \mid \mathbf{W}_i \mathbf{R}_\ell + \delta_{i\ell} \mathbf{G}].$$

Consequently, the i -th block row of $\tilde{\mathbf{D}}_\ell \hat{\mathbf{T}}$ is

$$\mathbf{A}^\star \mathbf{Q}_i - \mathbf{W}_i \mathbf{R} = [\delta_{i1} \mathbf{G} \mid \cdots \mid \delta_{i\ell} \mathbf{G}]$$

and therefore

$$\tilde{\mathbf{D}}_\ell \hat{\mathbf{T}} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q.$$

Thus, every column of $\mathbf{I}_\ell \otimes \mathbf{G}$ lies in the column span of $\tilde{\mathbf{D}}_\ell$. Since $\mathbf{I}_\ell \otimes \mathbf{G}$ has full row rank, this identity also shows that $\tilde{\mathbf{D}}_\ell$ has full row rank.

Every entry of $\mathbf{Q}$ belongs to $\{-1, 0, 1\}$. On $\mathcal{E}$,

$$\|\hat{\mathbf{T}}\| \leq \max\{1, \|\mathbf{R}\|\} \leq 2\sigma_{\text{dec}} \sqrt{k\ell n \lceil \log q \rceil}.$$

Since $m = \ell^2 n \lceil \log q \rceil$, the assumed lower bound on σ_{suc} gives

$$(\ell m + k) \|\hat{\mathbf{T}}\| \log(\ell n) \leq 2(\ell^3 n \lceil \log q \rceil + k) \sigma_{\text{dec}} \sqrt{k\ell n \lceil \log q \rceil} \log(\ell n) \leq \sigma_{\text{suc}}. \quad (35)$$

Moreover,

$$\ell m + k \geq (\ell^3 + 2)n \lceil \log q \rceil \geq 3\ell n \lceil \log q \rceil,$$

where the last inequality follows from $\ell^3 + 2 - 3\ell = (\ell - 1)^2(\ell + 2) \geq 0$.

Thus, the call to SamplePre in Step 4 satisfies all the hypotheses of Lemma 3.6, with row dimension ℓn and column dimension $\ell m + k$.

The Hybrid Experiments. All challenge vectors $\mathbf{c}^\star$ below are reduced modulo q . As before, the changes from each experiment to the preceding one are highlighted in blue.

1. For $x \in \{\text{pl}, \text{unif}\}$, let $\text{Hyb}_{0,x}$ denote the distribution of the view

$$(\mathbf{A}^\star, \mathbf{c}^{\star\top}, \mathbf{W}^\star, \mathbf{T}^\star)$$

as defined in the reduction $\mathcal{B}^\mathcal{A}$ and conditioned on the non-abort event $\mathcal{E}$. We use $x = \text{pl}$ to denote the distribution when the decomposed LWE challenge vector $\mathbf{b}$ is planted and $x = \text{unif}$ when it is uniform.

2. For $x \in \{\text{pl}, \text{unif}\}$, define $\text{Hyb}_{1,x}$ by modifying $\text{Hyb}_{0,x}$ such that the fourth coordinate of $\text{Hyb}_{0,x}$ is instead freshly sampled as

$$\mathbf{T} \sim [\mathbf{I}_\ell \otimes \mathbf{A}^\star \mid \mathbf{W}^\star]^{-1} (\mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\text{suc}}).$$

That is, this hybrid samples $\mathbf{T}$ from the exact Gaussian preimage distribution, rather than using SamplePre as was done in the reduction, so that

$$\text{Hyb}_{1,x} := (\mathbf{A}^\star, \mathbf{c}^{\star\top}, \mathbf{W}^\star, \mathbf{T}).$$

3. For $x \in \{\text{pl}, \text{unif}\}$, we define $\text{Hyb}_{2,x}$ by modifying the distribution of the challenge vector $\mathbf{c}$ that lies in the second coordinate of $\text{Hyb}_{1,x}$. In the prior hybrid, the challenge vector was sampled by ReRand . The new hybrid samples it freshly according to the succinct LWE distribution. Namely, in the hybrid $\text{Hyb}_{2,x}$, we still have

$$\mathbf{T} \sim [\mathbf{I}_\ell \otimes \mathbf{A}^\star \mid \mathbf{W}^\star]^{-1} (\mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\text{suc}})$$

but now we freshly sample $\mathbf{s} \sim \text{Sample}(1^n)$, $\tilde{\mathbf{e}} \sim \mathcal{D}_{\mathbb{Z}^m, \chi_{\text{suc}}}$, and $\mathbf{u} \sim \mathbb{Z}_q^m$ independently of $\mathbf{A}^\star$ and $\mathbf{W}^\star$ and modify the challenge vector as follows:

$$\begin{aligned} \text{Hyb}_{2,\text{pl}} &:= (\mathbf{A}^\star, \mathbf{s}^\top \mathbf{A}^\star + \tilde{\mathbf{e}}^\top, \mathbf{W}^\star, \mathbf{T}), \\ \text{Hyb}_{2,\text{unif}} &:= (\mathbf{A}^\star, \mathbf{u}^\top, \mathbf{W}^\star, \mathbf{T}). \end{aligned}$$

4. For $x \in \{\text{pl}, \text{unif}\}$, we define $\text{Hyb}_{3,x}$ by modifying the distribution of the previously structured matrix $\mathbf{A}^\star$ to a uniform random matrix. Namely, sample $\mathbf{s} \sim \text{Sample}(1^n)$, $\tilde{\mathbf{e}} \sim \mathcal{D}_{\mathbb{Z}^m, \chi_{\text{suc}}}$, $\mathbf{u} \sim \mathbb{Z}_q^m$, $\mathbf{U}_A \sim \mathbb{Z}_q^{n \times m}$, and $\mathbf{U}_W \sim \mathbb{Z}_q^{\ell n \times k}$ mutually independently, and then sample

$$\mathbf{T} \sim [\mathbf{I}_\ell \otimes \mathbf{U}_A \mid \mathbf{U}_W]^{-1} (\mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\text{suc}}).$$

Define

$$\begin{aligned} \text{Hyb}_{3,\text{pl}} &:= (\mathbf{U}_A, \mathbf{s}^\top \mathbf{U}_A + \tilde{\mathbf{e}}^\top, \mathbf{U}_W, \mathbf{T}), \\ \text{Hyb}_{3,\text{unif}} &:= (\mathbf{U}_A, \mathbf{u}^\top, \mathbf{U}_W, \mathbf{T}). \end{aligned}$$

If the conditional Gaussian distribution defining $\mathbf{T}$ is empty, the corresponding experiment outputs $\perp$.

Lemma 5.2 (Trapdoor hybrid). *For every $x \in \{\text{pl}, \text{unif}\}$, we have $\text{Hyb}_{0,x} \approx_s \text{Hyb}_{1,x}$.*

Proof. Let ρ_{RR} denote the random coins used by ReRand, and referring to the same random variables defined above in the reduction, define

$$\Omega_{\text{pl}} = (\overline{\mathbf{W}}, \mathbf{R}, \mathbf{s}, \mathbf{e}, \mathbf{X}, \rho_{\text{RR}}), \quad \Omega_{\text{unif}} = (\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}, \mathbf{X}, \rho_{\text{RR}}).$$

Fix $x \in \{\text{pl}, \text{unif}\}$ and condition on $\mathcal{E}$ and on any value $\Omega_x = \omega$. Under this conditioning, the first three coordinates of the hybrid, as well as $\tilde{\mathbf{D}}_\ell$ and $\hat{\mathbf{T}}$, are fixed. The only remaining randomness in the fourth coordinate of $\text{Hyb}_{0,x}$ consists of the coins of SamplePre.

Since Equation (35) holds, the preimage distribution guarantee of Lemma 3.6, followed by a hybrid over the $\ell n \lceil \log q \rceil$ target columns, gives

$$\Delta((\text{Hyb}_{0,x} \mid \Omega_x = \omega, \mathcal{E}), (\text{Hyb}_{1,x} \mid \Omega_x = \omega, \mathcal{E})) \leq \ell n \lceil \log q \rceil 2^{-\ell n}.$$

Both hybrids use the same distribution of $\Omega_x \mid \mathcal{E}$. Therefore, by convexity of statistical distance,

$$\begin{aligned} \Delta(\text{Hyb}_{0,x}, \text{Hyb}_{1,x}) &\leq \mathbb{E}_{\omega \sim \Omega_x \mid \mathcal{E}} [\Delta((\text{Hyb}_{0,x} \mid \omega, \mathcal{E}), (\text{Hyb}_{1,x} \mid \omega, \mathcal{E}))] \\ &\leq \ell n \lceil \log q \rceil 2^{-\ell n} \\ &= \text{negl}(\lambda). \end{aligned}$$

□

Lemma 5.3 (Challenge hybrid). *For every $x \in \{\text{pl}, \text{unif}\}$, we have $\text{Hyb}_{1,x} \approx_s \text{Hyb}_{2,x}$.*

Proof. Since $\|\mathbf{S}\| = 1$ and $\chi_{\text{suc}} \geq 3m\chi_{\text{dec}}$, we have $\chi_{\text{suc}} > 2\chi_{\text{dec}}m\|\mathbf{S}\|$. Moreover, $m = \text{poly}(\lambda)$ and $\chi_{\text{dec}} = \omega(\sqrt{\log \lambda})$ imply $\chi_{\text{dec}} = \omega(\sqrt{\log m})$. Therefore, Lemma 3.4 applies with $\mathbf{V} = \mathbf{S}$.

Let

$$Z_{\text{pl}} = (\overline{\mathbf{W}}, \mathbf{R}, \mathbf{X}, \mathbf{s}, \mathbf{T}), \quad Z_{\text{unif}} = (\overline{\mathbf{W}}, \mathbf{R}, \mathbf{X}, \mathbf{T}).$$

In the planted branch, we may use the same $\mathbf{s}$ in both hybrids, since it is distributed according to $\text{Sample}(1^n)$ and is independent of $(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{X}, \mathbf{T})$. Conditioned on $Z_{\text{pl}} = z$ and $\mathcal{E}$, the matrices $\mathbf{S}$, $\mathbf{A}_{\text{dec}}$, and $\mathbf{A}^*$ are fixed, whereas the original error and the coins of ReRand remain independent. Hence Lemma 3.4 gives

$$\mathbf{c}^{\star\top} \approx_s \mathbf{s}^\top \mathbf{A}_{\text{dec}} \mathbf{S} + \tilde{\mathbf{e}}^\top = \mathbf{s}^\top \mathbf{A}^* + \tilde{\mathbf{e}}^\top \text{ mod } q.$$

In the uniform branch, conditioned on $Z_{\text{unif}} = z$ and $\mathcal{E}$, the matrix $\mathbf{S}$ is fixed and invertible modulo q , while $\mathbf{b}$ remains uniform and independent. The uniform input conclusion of Lemma 3.4 therefore gives $\mathbf{c}^{\star\top} \approx_s \mathbf{u}^\top$.

Let $\epsilon_{\text{RR}}(\lambda)$ denote the negligible statistical error in Lemma 3.4. In either branch, for every z in the support of $Z_x \mid \mathcal{E}$, the preceding pointwise comparison gives

$$\Delta(\text{Hyb}_{1,x} \mid Z_x = z, \mathcal{E}, \text{Hyb}_{2,x} \mid Z_x = z, \mathcal{E}) \leq \epsilon_{\text{RR}}(\lambda).$$

Since $Z_x \mid \mathcal{E}$ has the same distribution in both hybrids, averaging over z yields $\Delta(\text{Hyb}_{1,x}, \text{Hyb}_{2,x}) \leq \epsilon_{\text{RR}}(\lambda)$, proving the claim. □

Lemma 5.4 (Public matrix hybrid). *For every $x \in \{\text{pl}, \text{unif}\}$, we have $\text{Hyb}_{2,x} \approx_s \text{Hyb}_{3,x}$.*

Proof. For a pair of matrices $(\mathbf{A}, \mathbf{W})$, define $\mathbf{D}(\mathbf{A}, \mathbf{W}) := [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$. By Lemma 5.8 and since $\overline{\mathbf{W}}$ is the deterministic block-transpose of $\mathbf{W}$, we have $(\mathbf{A}^\star, \mathbf{W}^\star) \approx_s (\mathbf{U}_\mathbf{A}, \mathbf{U}_\mathbf{W})$. Moreover,

$$\Delta((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}, (\mathbf{A}^\star, \mathbf{W}^\star)) \leq \Pr[\neg \mathcal{E}] = \text{negl}(\lambda).$$

Hence $((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}) \approx_s (\mathbf{U}_\mathbf{A}, \mathbf{U}_\mathbf{W})$.

Then, define the randomized map K_{pl} as follows. On input $(\mathbf{A}, \mathbf{W})$, if $\mathbf{D}(\mathbf{A}, \mathbf{W})$ does not have full row rank, output $\perp$. Otherwise, using fresh randomness, sample $\mathbf{s} \sim \text{Sample}(1^n)$, $\tilde{\mathbf{e}} \sim \mathcal{D}_{\mathbb{Z}^m, \chi_{\text{suc}}}$, and

$$\mathbf{T} \sim \mathbf{D}(\mathbf{A}, \mathbf{W})^{-1} (\mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\text{suc}}),$$

and output $K_{\text{pl}}(\mathbf{A}, \mathbf{W}) = (\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \tilde{\mathbf{e}}^\top, \mathbf{W}, \mathbf{T})$.

Define K_{unif} analogously, except that it samples $\mathbf{u} \sim \mathbb{Z}_q^m$ and outputs $K_{\text{unif}}(\mathbf{A}, \mathbf{W}) = (\mathbf{A}, \mathbf{u}^\top, \mathbf{W}, \mathbf{T})$. Let $K_x((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E})$ denote the distribution obtained by sampling $(\mathbf{A}^\star, \mathbf{W}^\star)$ according to its distribution conditioned on $\mathcal{E}$ and then applying K_x using fresh randomness.

When the input is $(\mathbf{A}^\star, \mathbf{W}^\star)$, we have $\mathbf{D}(\mathbf{A}^\star, \mathbf{W}^\star) \hat{\mathbf{T}} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q$. Since $\mathbf{I}_\ell \otimes \mathbf{G}$ has full row rank, $\mathbf{D}(\mathbf{A}^\star, \mathbf{W}^\star)$ has full row rank. Hence $K_x((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}) \equiv \text{Hyb}_{2,x}$.

When the input is $(\mathbf{U}_\mathbf{A}, \mathbf{U}_\mathbf{W})$,

$$\Pr[\text{rank}(\mathbf{U}_\mathbf{A}) < n] \leq \frac{q^n - 1}{q - 1} q^{-m} < \frac{q^{n-m}}{q - 1} = \text{negl}(\lambda).$$

Whenever $\mathbf{U}_\mathbf{A}$ has full row rank, so does $\mathbf{D}(\mathbf{U}_\mathbf{A}, \mathbf{U}_\mathbf{W}) = [\mathbf{I}_\ell \otimes \mathbf{U}_\mathbf{A} \mid \mathbf{U}_\mathbf{W}]$. It follows that $K_x(\mathbf{U}_\mathbf{A}, \mathbf{U}_\mathbf{W}) \approx_s \text{Hyb}_{3,x}$, where the only difference is the negligible probability of outputting $\perp$.

By Fact 3.1,

$$K_x((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}) \approx_s K_x(\mathbf{U}_\mathbf{A}, \mathbf{U}_\mathbf{W}).$$

Moreover, $K_x((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}) \equiv \text{Hyb}_{2,x}$. Consequently,

$$\text{Hyb}_{2,x} \equiv K_x((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}) \approx_s K_x(\mathbf{U}_\mathbf{A}, \mathbf{U}_\mathbf{W}) \approx_s \text{Hyb}_{3,x}.$$

□

Proof of Theorem 5.1. Let $\mathcal{A}$ and $\mathcal{B}^\mathcal{A}$ be as above. Combining Lemmas 5.2 to 5.4 and using the triangle inequality gives

$$\text{Hyb}_{0,\text{pl}} \approx_s \text{Hyb}_{3,\text{pl}}, \quad \text{Hyb}_{0,\text{unif}} \approx_s \text{Hyb}_{3,\text{unif}}.$$

The final experiments are, up to the negligible failure probability above, the planted and uniform distributions of (ℓ, Sample) -succinct LWE. Finally, removing the conditioning on $\mathcal{E}$ incurs only negligible loss. Consequently, $\mathcal{B}^\mathcal{A}$ has the distinguishing advantage of $\mathcal{A}$, up to a negligible loss. □

Remark 5.5 (SIS analog). Decomposed SIS was introduced in [AMR26a] as the natural search analog of decomposed LWE.

Our proof of Theorem 5.1 also shows that hardness of decomposed SIS implies hardness of ℓ -succinct SIS, defined analogously by replacing the LWE distinguishing task with that of finding a short nonzero vector in the kernel of the public matrix. Indeed, the public matrix and trapdoor

parts of the reduction apply unchanged, while the challenge rerandomization step is unnecessary. Thus, $(\mathbf{A}^\star, \mathbf{W}^\star, \mathbf{T}^\star)$ is statistically close to a valid ℓ -succinct SIS instance.

Suppose that a solver returns a short vector $\mathbf{z} \neq \mathbf{0}$ such that $\mathbf{A}^\star \mathbf{z} = \mathbf{0} \bmod q$. The reduction outputs $\mathbf{z}' := \mathbf{S}\mathbf{z}$. Since $\mathbf{A}^\star = \mathbf{A}_{\text{dec}}\mathbf{S}$, we have

$$\mathbf{A}_{\text{dec}}\mathbf{z}' = \mathbf{A}_{\text{dec}}\mathbf{S}\mathbf{z} = \mathbf{A}^\star \mathbf{z} = \mathbf{0} \bmod q.$$

Moreover, $\mathbf{S}$ is invertible over $\mathbb{Z}$, and hence $\mathbf{z}' \neq \mathbf{0}$. Finally,

$$\|\mathbf{z}'\|_2 \leq \|\mathbf{S}\|_{2 \rightarrow 2} \|\mathbf{z}\|_2 \leq m \|\mathbf{z}\|_2.$$

Here $\|\mathbf{S}\|_{2 \rightarrow 2} \leq \|\mathbf{S}\|_F \leq m$ because all entries of $\mathbf{S}$ have absolute value at most one. Thus, the reduction incurs only a factor- m loss in the SIS norm bound.

The reverse reduction, from succinct SIS to decomposed SIS, follows by adapting the succinct LWE-to-decomposed LWE transformation of [AMR25, Theorem 4.2] to short kernel vectors. Our reduction, therefore, establishes their equivalence.

Lemma 5.6 (From a short SIS vector to an LWE distinguisher). *Let q be an odd prime and let $(\mathbf{A}, H)$ be any distribution of a public matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and an auxiliary hint. Suppose that, given $(\mathbf{A}, H)$, a solver finds a nonzero $\mathbf{x} \in \mathbb{Z}_q^m$ with $\mathbf{A}\mathbf{x} = \mathbf{0} \bmod q$ and $\|\mathbf{x}\|_2 \leq \gamma < q$ with probability δ . For $\lambda \geq 2$ and $\chi \geq 1$, put $B = \chi\gamma\sqrt{m(\lambda + \log m)}$ and assume $B \leq q/8$. Then the solver yields a distinguisher between $(\mathbf{A}, H, \mathbf{A}^\top \mathbf{s} + \mathbf{e})$ and $(\mathbf{A}, H, \mathbf{u})$ with advantage at least*

$$\delta \left(1 - 2^{-\lambda} - \frac{2\lfloor B \rfloor + 1}{q} \right),$$

where $\mathbf{s} \sim \mathbb{Z}_q^n$, $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}_q^m, \chi}$, and $\mathbf{u} \sim \mathbb{Z}_q^m$ are independent of $(\mathbf{A}, H)$.

Proof. Run the SIS solver using only $(\mathbf{A}, H)$ and verify its answer. On failure, output zero. On success, output one exactly when the centered representative of $\mathbf{x}^\top \mathbf{b} \bmod q$ has absolute value at most B , where $\mathbf{b}$ is the challenge. For an LWE challenge, $\mathbf{x}^\top \mathbf{b} = \mathbf{x}^\top \mathbf{e} \bmod q$. By Lemma 3.3 and Cauchy–Schwarz, $|\mathbf{x}^\top \mathbf{e}| \leq B$ except with probability $2^{-\lambda}$. The error is independent of the solver’s output. For a uniform challenge, $\mathbf{x} \not\equiv \mathbf{0} \bmod q$ because $0 < \|\mathbf{x}\|_2 < q$. Since q is prime, $\mathbf{x}^\top \mathbf{u}$ is uniform over $\mathbb{Z}_q$, so the acceptance probability is exactly $(2\lfloor B \rfloor + 1)/q$ conditional on solver success. Subtracting the two acceptance probabilities proves the claim. The reduction preserves the public hint distribution. $\square$

Remark 5.7 (The two directions between decomposed LWE and SIS). Applying Lemma 5.6 with $H = (\overline{\mathbf{W}}, \mathbf{R})$ gives a classical reduction from decisional decomposed LWE to decomposed SIS in the stated noise and solution norm regime. Composing it with Theorem 6.1 gives a reduction from search decomposed LWE to decomposed SIS, under that theorem’s standard LWE assumption and parameter conditions. In the other direction, Remark 5.5, Theorem 7.1, and the search version of Lemma 8.8 compose to give a quantum reduction from decomposed SIS to search decomposed LWE. For the last step, take $r = 0$, $t = \ell n \lceil \log q \rceil$, and $\text{Matrix}(\overline{\mathbf{W}}, \mathbf{M}) = \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G}$. This specializes generalized decomposed LWE to ordinary decomposed LWE and preserves the original secret. Each composition uses the parameters required at its successive steps; these are not claims of equivalence at identical noise widths, block counts, or solution norm bounds.

5.1 Rerandomization of the Public Matrix

First, we provide a lemma specifying the rerandomization of the decomposed LWE public matrix.

Let $\mathcal{X}_{\text{col}}$ denote the distribution over $\mathbb{Z}_q^{n \lceil \log q \rceil}$ obtained by sampling uniformly from $\{-1, 0, 1\}^{n \lceil \log q \rceil}$ and reducing modulo q . We will assume that q is an odd prime, implying $q \geq 3$, so that reduction modulo q on $\{-1, 0, 1\}$ is injective. It follows that $H_\infty(\mathcal{X}_{\text{col}}) = n \lceil \log q \rceil \log 3$.

Lemma 5.8 (Public matrix rerandomization). *Let $\ell = \ell(\lambda)$, $n = n(\lambda)$, $k = k(\lambda)$, $q = q(\lambda) \in \mathbb{N}$, and $\sigma = \sigma(\lambda) \in \mathbb{R}$ such that ℓ , $\log q$, n , and k are polynomially bounded in λ . Suppose*

$$\begin{aligned} q \text{ odd prime}, \quad k &\geq 2n \log q, \quad \sigma \geq 2, \\ n \lceil \log q \rceil &= \omega(\log \lambda), \quad \frac{q}{\sigma} = \omega(\sqrt{\log \lambda}). \end{aligned}$$

Let $\overline{\mathbf{W}}$ and $\mathbf{A}$ be sampled according to the public-matrix marginal of the ℓ -decomposed LWE distribution with parameters (n, k, σ, q) , i.e.,

$$\begin{aligned} \mathbf{A} &= \overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G}, \\ \overline{\mathbf{W}} &\sim \mathbb{Z}_q^{n \times \ell k}, \quad \mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times \ell n \lceil \log q \rceil}, \sigma}. \end{aligned}$$

Independently sample $\mathbf{X} \sim \{-1, 0, 1\}^{n \lceil \log q \rceil \times (\ell^2 - 1)n \lceil \log q \rceil}$ and parse $\mathbf{A} = [\mathbf{A}_0 \mid \mathbf{A}_{\text{rest}}]$, with $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times n \lceil \log q \rceil}$ and $\mathbf{A}_{\text{rest}} \in \mathbb{Z}_q^{n \times (\ell^2 - 1)n \lceil \log q \rceil}$. Then

$$(\overline{\mathbf{W}}, \mathbf{A}_0, \mathbf{A}_0 \mathbf{X} + \mathbf{A}_{\text{rest}}) \approx_s (\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{U}_{\text{rest}}),$$

where $\mathbf{U}_0 \sim \mathbb{Z}_q^{n \times n \lceil \log q \rceil}$ and $\mathbf{U}_{\text{rest}} \sim \mathbb{Z}_q^{n \times (\ell^2 - 1)n \lceil \log q \rceil}$.

Proof. The argument proceeds through four experiments, each defining a random output tuple. The first experiment is exactly the distribution obtained by applying the public matrix transformation to a decomposed LWE instance, whereas the last experiment is the required uniform succinct LWE public matrix distribution. The changes from each experiment to the preceding one are highlighted in blue.

All four experiments use the following common convention: if $\mathbf{W}_1$ does not have full row rank, the experiment outputs a fixed symbol $\perp$ and terminates. This ensures that every conditional Gaussian distribution appearing below is well defined. The common abort event has negligible probability. Indeed, since $\mathbf{W}_1$ is uniform over $\mathbb{Z}_q^{n \times k}$,

$$\begin{aligned} \Pr[\text{rank}(\mathbf{W}_1) < n] &\leq \frac{q^n - 1}{q - 1} q^{-k} < \frac{q^{n-k}}{q - 1} \\ &\leq q^{-n} = 2^{-n \log q} \\ &= 2^{-\Omega(n \lceil \log q \rceil)} = \text{negl}(\lambda). \end{aligned}$$

The third inequality follows from $k \geq 2n \log q \geq 2n$, and the last equality follows from $n \lceil \log q \rceil = \omega(\log \lambda)$. At the end of the proof, we remove this abort from the two endpoint experiments, losing at most the negligible probability that $\mathbf{W}_1$ does not have full row rank. We only require the matrix $\mathbf{W}_1$ to be full row rank and need no such condition for the matrices $\mathbf{W}_i$ for $i \geq 2$.

Expt_{real}:

1. Sample

$$\begin{aligned} \mathbf{W}_i &\sim \mathbb{Z}_q^{n \times k}, & \mathbf{R}_i &\sim \mathcal{D}_{\mathbb{Z}^{k \times n \lceil \log q \rceil}, \sigma} & \text{for all } i \in [\ell], \\ \mathbf{X} &\sim \{-1, 0, 1\}^{n \lceil \log q \rceil \times (\ell^2 - 1)n \lceil \log q \rceil}. \end{aligned}$$

2. Let

$$\begin{aligned} \overline{\mathbf{W}} &= [\mathbf{W}_1 \ \cdots \ \mathbf{W}_\ell] \in \mathbb{Z}_q^{n \times \ell k}, \\ \mathbf{A}_0 &= \mathbf{W}_1 \cdot \mathbf{R}_1 + \mathbf{G} \in \mathbb{Z}_q^{n \times n \lceil \log q \rceil}, \\ \mathbf{A}_{\text{rest}} &= [\mathbf{W}_i \cdot \mathbf{R}_j + \delta_{i,j} \mathbf{G}]_{(i,j) \neq (1,1)} \in \mathbb{Z}_q^{n \times (\ell^2 - 1)n \lceil \log q \rceil}. \end{aligned}$$

3. Output

$$(\overline{\mathbf{W}}, \mathbf{A}_0, \mathbf{A}_0 \cdot \mathbf{X} + \mathbf{A}_{\text{rest}}).$$

Expt₁:

1. Sample

$$\begin{aligned} \mathbf{W}_i &\sim \mathbb{Z}_q^{n \times k}, & \mathbf{R}_i &\sim \mathcal{D}_{\mathbb{Z}^{k \times n \lceil \log q \rceil}, \sigma} & \text{for all } i \in [\ell], \\ \widehat{\mathbf{R}}_1 &\leftarrow \mathbf{W}_1^{-1}(\mathbf{W}_1 \cdot \mathbf{R}_1, \sigma), \\ \mathbf{X} &\sim \{-1, 0, 1\}^{n \lceil \log q \rceil \times (\ell^2 - 1)n \lceil \log q \rceil}. \end{aligned}$$

2. Let

$$\begin{aligned} \overline{\mathbf{W}} &= [\mathbf{W}_1 \ \cdots \ \mathbf{W}_\ell] \in \mathbb{Z}_q^{n \times \ell k}, \\ \widehat{\mathbf{R}}_j &= \mathbf{R}_j & \text{for all } j \in [2, \ell], \\ \mathbf{A}_0 &= \mathbf{W}_1 \cdot \mathbf{R}_1 + \mathbf{G} \in \mathbb{Z}_q^{n \times n \lceil \log q \rceil}, \\ \widehat{\mathbf{A}}_{\text{rest}} &= [\mathbf{W}_i \cdot \widehat{\mathbf{R}}_j + \delta_{i,j} \mathbf{G}]_{(i,j) \neq (1,1)} \in \mathbb{Z}_q^{n \times (\ell^2 - 1)n \lceil \log q \rceil}. \end{aligned}$$

3. Output

$$(\overline{\mathbf{W}}, \mathbf{A}_0, \mathbf{A}_0 \cdot \mathbf{X} + \widehat{\mathbf{A}}_{\text{rest}}).$$

Expt₂:

1. Sample

$$\begin{aligned} \mathbf{W}_i &\sim \mathbb{Z}_q^{n \times k}, & \mathbf{R}_i &\sim \mathcal{D}_{\mathbb{Z}^{k \times n \lceil \log q \rceil}, \sigma} & \text{for all } i \in [\ell], \\ \mathbf{U}_0 &\sim \mathbb{Z}_q^{n \times n \lceil \log q \rceil}, \\ \widehat{\mathbf{R}}_1 &\leftarrow \mathbf{W}_1^{-1}(\mathbf{U}_0 - \mathbf{G}, \sigma), \\ \mathbf{X} &\sim \{-1, 0, 1\}^{n \lceil \log q \rceil \times (\ell^2 - 1)n \lceil \log q \rceil}. \end{aligned}$$

2. Let

$$\begin{aligned}\overline{\mathbf{W}} &= [\mathbf{W}_1 \ \cdots \ \mathbf{W}_\ell] \in \mathbb{Z}_q^{n \times \ell k}, \\ \widehat{\mathbf{R}}_j &= \mathbf{R}_j \quad \text{for all } j \in [2, \ell], \\ \widehat{\mathbf{A}}_{\text{rest}} &= [\mathbf{W}_i \cdot \widehat{\mathbf{R}}_j + \delta_{i,j} \mathbf{G}]_{(i,j) \neq (1,1)} \in \mathbb{Z}_q^{n \times (\ell^2 - 1)n \lceil \log q \rceil}.\end{aligned}$$

3. Output

$$(\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{U}_0 \cdot \mathbf{X} + \widehat{\mathbf{A}}_{\text{rest}}).$$

$\text{Expt}_{\text{rand}}$:

1. Sample

$$\begin{aligned}\mathbf{W}_i &\sim \mathbb{Z}_q^{n \times k} \quad \text{for all } i \in [\ell], \\ \mathbf{U}_0 &\sim \mathbb{Z}_q^{n \times n \lceil \log q \rceil}, \\ \mathbf{U}_{\text{rest}} &\sim \mathbb{Z}_q^{n \times (\ell^2 - 1)n \lceil \log q \rceil}.\end{aligned}$$

2. Let

$$\overline{\mathbf{W}} = [\mathbf{W}_1 \ \cdots \ \mathbf{W}_\ell] \in \mathbb{Z}_q^{n \times \ell k}$$

3. Output

$$(\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{U}_{\text{rest}}).$$

We now bound the statistical distance between $\text{Expt}_{\text{real}}$ and $\text{Expt}_{\text{rand}}$. We apply Lemma 3.2 in two columnwise hybrids. The first hybrid ranges over the $n \lceil \log q \rceil$ columns of $\mathbf{R}_1$, whereas the second ranges over the $(\ell^2 - 1)n \lceil \log q \rceil$ columns of $\mathbf{X}$.

Let $\mathcal{R}$ denote the distribution over $\mathbb{Z}_q^k$ obtained by sampling a column from $\mathcal{D}_{\mathbb{Z}^k, \sigma}$ and reducing it modulo q . By the assumptions $k \geq 2n \log q$ and $\sigma \geq 2$, together with Fact 3.5, we have

$$H_\infty(\mathcal{R}) - n \log q \geq k \log \sigma - n \log q - \text{negl}(\lambda) \geq n \log q - \text{negl}(\lambda).$$

Consequently, $H_\infty(\mathcal{R}) - n \log q = \Omega(n \lceil \log q \rceil)$. Since $H_\infty(\mathcal{X}_{\text{col}}) = n \lceil \log q \rceil \log 3$, we have

$$H_\infty(\mathcal{X}_{\text{col}}) - n \log q = n(\lceil \log q \rceil \log 3 - \log q),$$

which gives $H_\infty(\mathcal{X}_{\text{col}}) - n \log q = \Omega(n \lceil \log q \rceil)$.

Thus, in both applications, the source min-entropy exceeds the output length $n \log q$ by $\Omega(n \lceil \log q \rceil)$. By Lemma 3.2, each column replacement incurs statistical distance at most

$$2^{-\Omega(n \lceil \log q \rceil)}.$$

The total number of column replacements is $\ell^2 n \lceil \log q \rceil$, which is polynomial in λ . Since $n \lceil \log q \rceil = \omega(\log \lambda)$, the total statistical distance accumulated by both columnwise hybrids is negligible in λ .

1. $\text{Expt}_{\text{real}} \equiv \text{Expt}_1$: The first two output coordinates, $\overline{\mathbf{W}}$ and $\mathbf{W}_1 \mathbf{R}_1 + \mathbf{G}$, have the same marginal distribution in both experiments. Fix any values $\overline{\mathbf{W}}^* = [\mathbf{W}_1^* \ \cdots \ \mathbf{W}_\ell^*] \in \mathbb{Z}_q^{n \times \ell k}$ and $\mathbf{A}_0^* \in \mathbb{Z}_q^{n \times n \lceil \log q \rceil}$. Conditioned on the event that in $\text{Expt}_{\text{real}}$ the first two coordinates satisfy $\overline{\mathbf{W}} = \overline{\mathbf{W}}^*$ and $\mathbf{W}_1 \mathbf{R}_1 + \mathbf{G} = \mathbf{A}_0^*$, the conditional marginal distribution of $\mathbf{R}_1$ in $\text{Expt}_{\text{real}}$ is the Gaussian distribution $\mathcal{D}_{\mathbb{Z}^{k \times n \lceil \log q \rceil}, \sigma}$ conditioned on $\mathbf{A}_0^* = \mathbf{W}_1 \mathbf{R}_1 + \mathbf{G} \bmod q$, i.e.,

$$\mathbf{R}_1 \sim (\mathbf{W}_1^*)^{-1} (\mathbf{A}_0^* - \mathbf{G}, \sigma).$$

In Expt_1 , the matrix $\widehat{\mathbf{R}}_1$ is a fresh sample from exactly the same conditional distribution, because $\mathbf{W}_1 \mathbf{R}_1 = \mathbf{A}_0^* - \mathbf{G} \bmod q$.

The remaining matrices $\mathbf{R}_2, \dots, \mathbf{R}_\ell$ and $\mathbf{X}$ are sampled identically in both experiments. The third output coordinate is then obtained by the same deterministic computation, using respectively $\mathbf{R}_1$ or $\widehat{\mathbf{R}}_1$. Hence its conditional distribution given the first two coordinates is identical in the two experiments. Therefore,

$$\text{Expt}_{\text{real}} \equiv \text{Expt}_1.$$

2. $\Delta(\text{Expt}_1, \text{Expt}_2) = \text{negl}(\lambda)$. Recall that both experiments output $\perp$ when $\mathbf{W}_1$ does not have full row rank. Write the columns of $\mathbf{R}_1$ as $\mathbf{r}_{1,1}, \dots, \mathbf{r}_{1,n \lceil \log q \rceil}$. Each column, reduced modulo q , is an independent sample from $\mathcal{R}$. For each column, Lemma 3.2, applied with $\mathbf{A} = \mathbf{W}_1$ and $\mathbf{x} = \mathbf{r}_{1,j} \bmod q$, gives

$$(\mathbf{W}_1, \mathbf{W}_1 \mathbf{r}_{1,j}) \approx_s (\mathbf{W}_1, \mathbf{u}_j),$$

with statistical distance $2^{-\Omega(n \lceil \log q \rceil)}$, where $\mathbf{u}_j$ is independently uniform over $\mathbb{Z}_q^n$. A standard hybrid argument over the $n \lceil \log q \rceil$ columns, plus adding the fixed matrix $\mathbf{G}$ at the end, gives

$$\Delta((\mathbf{W}_1, \mathbf{W}_1 \mathbf{R}_1 + \mathbf{G}), (\mathbf{W}_1, \mathbf{U}_0)) \leq n \lceil \log q \rceil 2^{-\Omega(n \lceil \log q \rceil)},$$

where $\mathbf{U}_0$ is uniform over $\mathbb{Z}_q^{n \times n \lceil \log q \rceil}$.

Define the following randomized postprocessing K on input $(\mathbf{W}_1, \mathbf{A}_0)$ and apply it to both distributions:

- If $\mathbf{W}_1$ does not have full row rank, output $\perp$.
- Otherwise, independently sample

$$\mathbf{W}_2, \dots, \mathbf{W}_\ell, \quad \mathbf{R}_2, \dots, \mathbf{R}_\ell, \quad \mathbf{X}$$

according to their respective distributions, and sample

$$\widehat{\mathbf{R}}_1 \sim \mathbf{W}_1^{-1} (\mathbf{A}_0 - \mathbf{G}, \sigma).$$

Set $\widehat{\mathbf{R}}_j = \mathbf{R}_j$ for $j \in \{2, \dots, \ell\}$ and define

$$\widehat{\mathbf{A}}_{\text{rest}} = \left[\mathbf{W}_i \widehat{\mathbf{R}}_j + \delta_{i,j} \mathbf{G} \right]_{(i,j) \neq (1,1)}.$$

Finally, output

$$(\overline{\mathbf{W}}, \mathbf{A}_0, \mathbf{A}_0 \mathbf{X} + \widehat{\mathbf{A}}_{\text{rest}}).$$

This matrix $\widehat{\mathbf{R}}_1$ is well defined because the matrix $\mathbf{W}_1$ has full row rank. That is, $\mathbf{W}_1$ defines a surjective map so that for every value of $\mathbf{A}_0$ there exists at least one inverse $\widehat{\mathbf{R}}_1$ such that $\mathbf{W}_1 \widehat{\mathbf{R}}_1 = \mathbf{A}_0 - \mathbf{G}$. For the first input distribution, this postprocessing gives Expt_1 ; for the second, it gives Expt_2 . Therefore, by Fact 3.1, and the assumption that $n \lceil \log q \rceil = \omega(\log \lambda)$, one obtains

$$\begin{aligned} \Delta(\text{Expt}_1, \text{Expt}_2) &\leq \Delta((\mathbf{W}_1, \mathbf{W}_1 \mathbf{R}_1 + \mathbf{G}), (\mathbf{W}_1, \mathbf{U}_0)) \\ &\leq n \lceil \log q \rceil 2^{-\Omega(n \lceil \log q \rceil)} \\ &= \text{negl}(\lambda). \end{aligned}$$

3. $\Delta(\text{Expt}_2, \text{Expt}_{\text{rand}}) = \text{negl}(\lambda)$: A columnwise application of Lemma 3.2 to $\mathbf{X}$ gives

$$\Delta((\mathbf{U}_0, \mathbf{U}_0 \mathbf{X}), (\mathbf{U}_0, \mathbf{U}_{\text{rest}})) \leq (\ell^2 - 1) n \lceil \log q \rceil 2^{-\Omega(n \lceil \log q \rceil)},$$

where $\mathbf{U}_{\text{rest}}$ is independently uniform over $\mathbb{Z}_q^{n \times (\ell^2 - 1) n \lceil \log q \rceil}$.

Apply the following randomized postprocessing to a pair $(\mathbf{U}_0, \mathbf{Y})$.

- Sample $\mathbf{W}_1, \dots, \mathbf{W}_\ell$ and $\mathbf{R}_2, \dots, \mathbf{R}_\ell$ according to their prescribed distributions.
- If $\mathbf{W}_1$ does not have full row rank, output $\perp$.
- Otherwise, sample

$$\widehat{\mathbf{R}}_1 \sim \mathbf{W}_1^{-1}(\mathbf{U}_0 - \mathbf{G}, \sigma),$$

set $\widehat{\mathbf{R}}_j = \mathbf{R}_j$ for $j \in \{2, \dots, \ell\}$, construct $\widehat{\mathbf{A}}_{\text{rest}}$ as in Expt_2 , and output

$$(\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{Y} + \widehat{\mathbf{A}}_{\text{rest}}).$$

Applied to $(\mathbf{U}_0, \mathbf{U}_0 \mathbf{X})$, this postprocessing gives Expt_2 . Applied to $(\mathbf{U}_0, \mathbf{U}_{\text{rest}})$, it gives

$$(\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{U}_{\text{rest}} + \widehat{\mathbf{A}}_{\text{rest}}).$$

Since $\mathbf{U}_{\text{rest}}$ is independently uniform, its translation by $\widehat{\mathbf{A}}_{\text{rest}}$ is also uniform and independent of the first two coordinates. Thus, the latter tuple is distributed exactly as $\text{Expt}_{\text{rand}}$. Therefore, as before, by Fact 3.1 and the assumption that $n \lceil \log q \rceil = \omega(\log \lambda)$, one gets

$$\begin{aligned} \Delta(\text{Expt}_2, \text{Expt}_{\text{rand}}) &\leq \Delta((\mathbf{U}_0, \mathbf{U}_0 \mathbf{X}), (\mathbf{U}_0, \mathbf{U}_{\text{rest}})) \\ &\leq (\ell^2 - 1) n \lceil \log q \rceil 2^{-\Omega(n \lceil \log q \rceil)} \\ &= \text{negl}(\lambda). \end{aligned}$$

By the triangle inequality, we get

$$\Delta(\text{Expt}_{\text{real}}, \text{Expt}_{\text{rand}}) \leq \text{negl}(\lambda).$$

Finally, removing the common abort from both experiments changes their statistical distance by at most $2 \Pr[\text{rank}(\mathbf{W}_1) < n] = \text{negl}(\lambda)$. Thus, the same conclusion holds for the unconditioned distributions in the statement. $\square$

6 A Search-to-Decision Reduction for Decomposed LWE

We use the search and decisional variants of ℓ -decomposed LWE from Definition 3.8. In particular, the parameter tuple is always (n, k, χ, σ, q) ; the block parameter ℓ is recorded in the name of the problem.

Theorem 6.1 (Search-to-decision reduction for decomposed LWE). *Let $\lambda \in \mathbb{N}$ be the security parameter and let*

$$\begin{aligned} n &= n(\lambda), & k &= k(\lambda), & \ell &= \ell(\lambda) \in \mathbb{N}, \\ \chi &= \chi(\lambda), & \sigma &= \sigma(\lambda), & \varepsilon &= \varepsilon(\lambda) \in \mathbb{R}^+, & \text{and} \\ q &= q(\lambda) \text{ odd prime.} \end{aligned}$$

Define $L := \ell \lceil 32(\lambda + 2) \ln 2 / \varepsilon^2 \rceil$. Suppose that

$$k \geq 6n \lceil \log q \rceil^2, \quad \sigma = \Omega(\sqrt{n} \lceil \log q \rceil), \quad \varepsilon \geq \frac{1}{\text{poly}(\lambda)}.$$

Assume moreover that $n, k, \ell, q, \chi, \sigma$ are $\text{poly}(\lambda)$ -bounded, and that decisional LWE with parameters $(n \lceil \log q \rceil, nL \lceil \log q \rceil, \chi, q)$ is hard for PPT algorithms.

Then every PPT algorithm that distinguishes decisional ℓ -decomposed LWE with parameters (n, k, χ, σ, q) with advantage at least ε yields a PPT algorithm that solves search L -decomposed LWE with parameters (n, k, χ, σ, q) with probability $1 - \text{negl}(\lambda)$.

Proof. Let $\mathcal{A}$ be the PPT distinguisher from the theorem statement, and set

$$N := \left\lceil \frac{32(\lambda + 2) \ln 2}{\varepsilon^2} \right\rceil, \quad \text{so that} \quad L = N\ell.$$

Since $\varepsilon \geq 1/\text{poly}(\lambda)$, the amplification factor N is polynomially bounded. We first amplify $\mathcal{A}$ so that the error of each of the nq invocations made below is exponentially small, allowing a union bound over all of them. Apply Lemma 6.2 to $\mathcal{A}$, and let $\mathcal{B}$ denote the resulting algorithm. Thus $\mathcal{B}$ distinguishes uniform from planted L -decomposed LWE samples with overwhelming probability, including for every fixed secret. We construct a PPT search algorithm $\mathcal{C}$ using $\mathcal{B}$.

The Reduction. On input a search L -decomposed LWE sample

$$(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top) \in \mathbb{Z}_q^{n \times Lk} \times \mathbb{Z}^{k \times Ln \lceil \log q \rceil} \times \mathbb{Z}_q^{1 \times L^2 n \lceil \log q \rceil},$$

the algorithm $\mathcal{C}$ proceeds as follows.

1. For every $i \in [n]$ and $t \in \mathbb{Z}_q$, independently sample $\mathbf{z}_{i,t} \sim \mathbb{Z}_q^n$, $\mathbf{p}_{i,t} \sim \mathbb{Z}_q^{Lk}$. Let $\boldsymbol{\varepsilon}_i \in \mathbb{Z}_q^n$ be the i -th standard basis vector, and define

$$\overline{\mathbf{V}}_{i,t} := \overline{\mathbf{W}} + \boldsymbol{\varepsilon}_i \mathbf{p}_{i,t}^\top$$

and

$$\mathbf{c}_{i,t}^\top := \mathbf{b}^\top + t \mathbf{p}_{i,t}^\top (\mathbf{I}_L \otimes \mathbf{R}) + \mathbf{z}_{i,t}^\top \left(\overline{\mathbf{V}}_{i,t} (\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G} \right).$$

2. For every $i \in [n]$ and $t \in \mathbb{Z}_q$, using fresh internal randomness in each invocation, set

$$\text{acc}_{i,t} \leftarrow \mathcal{B}(\bar{\mathbf{V}}_{i,t}, \mathbf{R}, \mathbf{c}_{i,t}^\top).$$

3. For every $i \in [n]$, if there is exactly one $t \in \mathbb{Z}_q$ for which $\text{acc}_{i,t} = 1$, set $\tilde{s}_i := t$; otherwise, return $\perp$ and terminate. Return $\tilde{\mathbf{s}}^\top := (\tilde{s}_1, \dots, \tilde{s}_n)$.

Correctness. We show that C recovers the secret except with negligible probability. Let $\mathbf{s} \in \mathbb{Z}_q^n$ be the secret of the input. There is an error vector $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{L^2 n \lceil \log q \rceil}, \chi}$ such that

$$\mathbf{b}^\top = \mathbf{s}^\top (\bar{\mathbf{W}}(\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}) + \mathbf{e}^\top.$$

Consequently,

$$\mathbf{c}_{i,t}^\top = (\mathbf{s} + \mathbf{z}_{i,t})^\top (\bar{\mathbf{V}}_{i,t}(\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}) + (t - s_i)\mathbf{p}_{i,t}^\top(\mathbf{I}_L \otimes \mathbf{R}) + \mathbf{e}^\top. \quad (36)$$

To interpret Equation (36), fix (i, t) and temporarily omit the term proportional to $t - s_i$. We claim that the resulting tuple

$$(\bar{\mathbf{V}}_{i,t}, \mathbf{R}, (\mathbf{s} + \mathbf{z}_{i,t})^\top (\bar{\mathbf{V}}_{i,t}(\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}) + \mathbf{e}^\top)$$

is distributed exactly as a planted L -decomposed LWE sample. Indeed, since $(\bar{\mathbf{W}}, \mathbf{p}_{i,t})$ is uniform and the map

$$(\bar{\mathbf{W}}, \mathbf{p}_{i,t}) \mapsto (\bar{\mathbf{W}} + \varepsilon_i \mathbf{p}_{i,t}^\top, \mathbf{p}_{i,t})$$

is a bijection, the pair $(\bar{\mathbf{V}}_{i,t}, \mathbf{p}_{i,t})$ is uniform over $\mathbb{Z}_q^{n \times Lk} \times \mathbb{Z}_q^{Lk}$. Moreover, $\mathbf{R}$ has its prescribed Gaussian distribution, $\mathbf{s} + \mathbf{z}_{i,t}$ is a fresh uniform secret, and $\mathbf{e}$ has its prescribed error distribution, with all these variables mutually independent for this fixed query. Different queries may remain correlated through $(\bar{\mathbf{W}}, \mathbf{R}, \mathbf{e})$.

Consequently, the only term in Equation (36) that does not already belong to a planted decomposed LWE challenge is

$$(t - s_i)\mathbf{p}_{i,t}^\top(\mathbf{I}_L \otimes \mathbf{R}).$$

We now analyze this term separately for correct and wrong guesses.

- *Correct guesses.* Fix $(i, t) \in [n] \times \mathbb{Z}_q$ with $t = s_i$. The additional term above vanishes, so

$$(\bar{\mathbf{V}}_{i,t}, \mathbf{R}, \mathbf{c}_{i,t}^\top)$$

is a planted L -decomposed LWE sample with parameters (n, k, χ, σ, q) . By Lemma 6.2, $\mathcal{B}$ outputs 1 except with probability at most $2^{-\lambda}$.

- *Wrong guesses.* Fix $(i, t) \in [n] \times \mathbb{Z}_q$ with $t \neq s_i$, and set

$$\mathbf{u} := \mathbf{s} + \mathbf{z}_{i,t}, \quad \mathbf{v} := (t - s_i)\mathbf{p}_{i,t}.$$

Since q is prime, $t - s_i$ is invertible in $\mathbb{Z}_q$. Thus $\mathbf{v}$ is uniform in $\mathbb{Z}_q^{Lk}$, and the independence established above shows that $\mathbf{u}, \bar{\mathbf{V}}_{i,t}, \mathbf{R}, \mathbf{v}$ and $\mathbf{e}$ are mutually independent. Apply Lemma 6.4 to Equation (36). Its LWE assumption is precisely the one in the theorem. It follows that $(\bar{\mathbf{V}}_{i,t}, \mathbf{R}, \mathbf{c}_{i,t}^\top)$ is computationally indistinguishable from a uniform L -decomposed LWE sample with parameters (n, k, χ, σ, q) .

Recovery. The correct-guess analysis shows that $\mathcal{B}$ accepts the guess $t = s_i$ except with probability at most $2^{-\lambda}$. For every $t \neq s_i$, Lemma 6.4 shows that the tuple given to $\mathcal{B}$ is computationally indistinguishable from a uniform L -decomposed LWE sample, so $\mathcal{B}$ accepts with probability at most $2^{-\lambda} + \text{negl}(\lambda)$. A union bound over the $nq = \text{poly}(\lambda)$ pairs (i, t) shows that, except with negligible probability, the unique accepted guess for every i is $t = s_i$. This union bound uses only the marginal guarantee for each query; the queries may be correlated through the reused input instance. $\square$

The amplification, splitting, and wrong-guess results used above are stated and proved in Lemmas 6.2 to 6.4, respectively.

Lemma 6.2 (Secret random self-reduction). *Let $\lambda, n, k, \ell, L, q, \chi, \sigma$ and ε be as in Theorem 6.1, and set*

$$N := \left\lceil \frac{32(\lambda + 2) \ln 2}{\varepsilon^2} \right\rceil, \quad \text{so that} \quad L = N\ell.$$

If a PPT distinguisher $\mathcal{A}$ distinguishes decisional ℓ -decomposed LWE with parameters (n, k, χ, σ, q) with advantage at least ε , then there is a PPT algorithm $\mathcal{B}$ with the following properties. All probabilities below also include the internal randomness of the algorithms.

- Uniform.

$$\Pr_{\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}} \left[\mathcal{B}(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top) = 1 \right] \leq 2^{-\lambda},$$

where $\overline{\mathbf{W}} \sim \mathbb{Z}_q^{n \times Lk}$, $\mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times Ln \lceil \log q \rceil}, \sigma}$, and $\mathbf{b} \sim \mathbb{Z}_q^{L^2 n \lceil \log q \rceil}$.

- Planted. For every fixed $\mathbf{s} \in \mathbb{Z}_q^n$,

$$\Pr_{\overline{\mathbf{W}}, \mathbf{R}, \mathbf{e}} \left[\mathcal{B}(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{s}^\top (\overline{\mathbf{W}}(\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}) + \mathbf{e}^\top) = 1 \right] \geq 1 - 2^{-\lambda},$$

where $\overline{\mathbf{W}} \sim \mathbb{Z}_q^{n \times Lk}$, $\mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times Ln \lceil \log q \rceil}, \sigma}$, and $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{L^2 n \lceil \log q \rceil}, \chi}$.

Proof. Define the acceptance probabilities of $\mathcal{A}$ on the uniform and planted ℓ -decomposed LWE distributions, respectively, by

$$p_{\text{unif}} := \Pr_{\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}} \left[\mathcal{A}(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top) = 1 \right],$$

where $\overline{\mathbf{W}} \sim \mathbb{Z}_q^{n \times \ell k}$, $\mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times \ell n \lceil \log q \rceil}, \sigma}$, and $\mathbf{b} \sim \mathbb{Z}_q^{\ell^2 n \lceil \log q \rceil}$, and

$$p_{\text{pl}} := \Pr_{\overline{\mathbf{W}}, \mathbf{R}, \mathbf{s}, \mathbf{e}} \left[\mathcal{A}(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{s}^\top (\overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G}) + \mathbf{e}^\top) = 1 \right],$$

where $\overline{\mathbf{W}} \sim \mathbb{Z}_q^{n \times \ell k}$, $\mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times \ell n \lceil \log q \rceil}, \sigma}$, $\mathbf{s} \sim \mathbb{Z}_q^n$, and $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{\ell^2 n \lceil \log q \rceil}, \chi}$. By hypothesis, $|p_{\text{pl}} - p_{\text{unif}}| \geq \varepsilon$.

We define $\mathcal{B}$, using independent internal randomness in every call to $\mathcal{A}$. On input

$$(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top) \in \mathbb{Z}_q^{n \times Lk} \times \mathbb{Z}^{k \times Ln \lceil \log q \rceil} \times \mathbb{Z}_q^{1 \times L^2 n \lceil \log q \rceil},$$

apply Lemma 6.3 with $\ell_0 = \ell$ and $\ell_1 = L = N\ell$ to obtain N tuples

$$\{(\overline{\mathbf{W}}_j, \mathbf{R}_j, \mathbf{b}_j^\top)\}_{j \in [N]},$$

each having the dimensions of an ℓ -decomposed LWE sample. Independently sample N uniform ℓ -decomposed LWE tuples $\mathbf{U}_1, \dots, \mathbf{U}_N$ and compute

$$\hat{p}_{\text{unif}} := \frac{1}{N} \sum_{j=1}^N \mathcal{A}(\mathbf{U}_j).$$

For every $j \in [N]$, sample $\mathbf{z}_j \sim \mathbb{Z}_q^n$ independently and set

$$\mathbf{Z}_j := \left(\overline{\mathbf{W}}_j, \mathbf{R}_j, \mathbf{b}_j^\top + \mathbf{z}_j^\top \left(\overline{\mathbf{W}}_j (\mathbf{I}_\ell \otimes \mathbf{R}_j) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G} \right) \right).$$

Compute

$$\hat{p}_{\text{pl}} := \frac{1}{N} \sum_{j=1}^N \mathcal{A}(\mathbf{Z}_j).$$

Output 1 if $|\hat{p}_{\text{pl}} - \hat{p}_{\text{unif}}| > \varepsilon/2$, and output 0 otherwise.

The variables contributing to $\hat{p}_{\text{unif}}$ are independent samples from the uniform ℓ -decomposed LWE distribution, so $\hat{p}_{\text{unif}}$ estimates p_{unif} . If the input to $\mathcal{B}$ is uniform, then Lemma 6.3 and the invariance of a uniform target under translation show that the $\mathbf{Z}_j$ are independent uniform ℓ -decomposed LWE samples; hence $\hat{p}_{\text{pl}}$ also estimates p_{unif} . If the input is planted with any fixed secret $\mathbf{s}$, the same lemma yields independent planted ℓ -decomposed LWE samples sharing $\mathbf{s}$. Adding the independent mask $\mathbf{z}_j$ makes the secret $\mathbf{s} + \mathbf{z}_j$ independently uniform in every $\mathbf{Z}_j$, so $\hat{p}_{\text{pl}}$ estimates p_{pl} .

Hoeffding's inequality and the definition of N give

$$\Pr [|\hat{p}_{\text{unif}} - p_{\text{unif}}| > \varepsilon/8] \leq 2e^{-N\varepsilon^2/32} \leq 2^{-(\lambda+1)},$$

and the same bound holds for $\hat{p}_{\text{pl}}$ relative to whichever of p_{unif} and p_{pl} it estimates. Thus, on a uniform input, with probability at least $1 - 2^{-\lambda}$,

$$|\hat{p}_{\text{pl}} - \hat{p}_{\text{unif}}| \leq |\hat{p}_{\text{pl}} - p_{\text{unif}}| + |\hat{p}_{\text{unif}} - p_{\text{unif}}| \leq \varepsilon/4 < \varepsilon/2.$$

On a planted input, with the same probability,

$$|\hat{p}_{\text{pl}} - \hat{p}_{\text{unif}}| \geq |p_{\text{pl}} - p_{\text{unif}}| - |\hat{p}_{\text{pl}} - p_{\text{pl}}| - |\hat{p}_{\text{unif}} - p_{\text{unif}}| \geq 3\varepsilon/4 > \varepsilon/2.$$

This proves both guarantees. $\square$

Lemma 6.3 (Splitting decomposed LWE instances). *Fix the public decomposed LWE parameters (n, k, χ, σ, q) . Let $\ell_0, N \in \mathbb{N}$ and define $\ell_1 := N\ell_0$. There is an efficient algorithm that, on input*

$$(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\top) \in \mathbb{Z}_q^{n \times \ell_1 k} \times \mathbb{Z}^{k \times \ell_1 n \lceil \log q \rceil} \times \mathbb{Z}_q^{1 \times \ell_1^2 n \lceil \log q \rceil},$$

outputs N tuples

$$(\overline{\mathbf{W}}_j, \mathbf{R}_j, \mathbf{b}_j^\top) \in \mathbb{Z}_q^{n \times \ell_0 k} \times \mathbb{Z}^{k \times \ell_0 n \lceil \log q \rceil} \times \mathbb{Z}_q^{1 \times \ell_0^2 n \lceil \log q \rceil} \quad (j \in [N])$$

such that:

- If the input is a planted ℓ_1 -decomposed LWE sample with parameters (n, k, χ, σ, q) and secret $\mathbf{s}$, then the outputs are, conditioned on $\mathbf{s}$, independent planted ℓ_0 -decomposed LWE samples with the same parameters and the shared secret $\mathbf{s}$.
- If the input is a uniform ℓ_1 -decomposed LWE sample with parameters (n, k, χ, σ, q) , then the outputs are independent uniform ℓ_0 -decomposed LWE samples with those parameters.

Proof. Write

$$\overline{\mathbf{W}} = (\mathbf{W}^{[1]} \mid \dots \mid \mathbf{W}^{[\ell_1]}), \quad \mathbf{R} = (\mathbf{R}^{[1]} \mid \dots \mid \mathbf{R}^{[\ell_1]}),$$

where $\mathbf{W}^{[a]} \in \mathbb{Z}_q^{n \times k}$ and $\mathbf{R}^{[b]} \in \mathbb{Z}^{k \times n \lceil \log q \rceil}$. View $\mathbf{b}^\top$ as an $\ell_1 \times \ell_1$ array of blocks $\mathbf{b}_{a,b}^\top \in \mathbb{Z}_q^{1 \times n \lceil \log q \rceil}$. If the input is planted with secret $\mathbf{s}$, then for every $a, b \in [\ell_1]$,

$$\mathbf{b}_{a,b}^\top = \mathbf{s}^\top (\mathbf{W}^{[a]} \mathbf{R}^{[b]} + \delta_{a,b} \mathbf{G}) + \mathbf{e}_{a,b}^\top,$$

where $\mathbf{e}_{a,b}^\top$ is the corresponding length- $n \lceil \log q \rceil$ block of the error vector.

For every $j \in [N]$, let

$$I_j := \{(j-1)\ell_0 + 1, \dots, j\ell_0\}.$$

Define $\overline{\mathbf{W}}_j$ and $\mathbf{R}_j$ by concatenating the blocks whose indices lie in I_j , and let $\mathbf{b}_j^\top$ be the concatenation, in the public decomposed LWE order, of the blocks $\mathbf{b}_{a,b}^\top$ with $a, b \in I_j$. Define $\mathbf{e}_j^\top$ analogously. Then

$$\mathbf{b}_j^\top = \mathbf{s}^\top (\overline{\mathbf{W}}_j (\mathbf{I}_{\ell_0} \otimes \mathbf{R}_j) + \text{vec}(\mathbf{I}_{\ell_0})^\top \otimes \mathbf{G}) + \mathbf{e}_j^\top.$$

Hence every output is exactly a planted ℓ_0 -decomposed LWE sample with the public parameter tuple (n, k, χ, σ, q) and secret $\mathbf{s}$. Conditioned on $\mathbf{s}$, outputs for distinct j are independent because they use disjoint blocks of $\overline{\mathbf{W}}$ and $\mathbf{R}$ and disjoint coordinates of $\mathbf{e}$.

If the original input is uniform, the same construction uses disjoint blocks of $\overline{\mathbf{W}}$ and $\mathbf{R}$ and disjoint coordinates of the uniform target $\mathbf{b}$. Therefore, the outputs are independent uniform ℓ_0 -decomposed LWE samples with parameters (n, k, χ, σ, q) . $\square$

Lemma 6.4 (Uniformity of masked decomposed LWE samples). *Let $\lambda, n, k, \ell, L, q, \chi, \sigma$ and ε be as in Theorem 6.1. Sample*

$$\overline{\mathbf{V}} \sim \mathbb{Z}_q^{n \times Lk}, \quad \mathbf{R} \sim \mathcal{D}_{\mathbb{Z}^{k \times Ln \lceil \log q \rceil}, \sigma}, \quad \mathbf{u} \sim \mathbb{Z}_q^n, \quad \mathbf{v} \sim \mathbb{Z}_q^{Lk}, \quad \mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{L^2 n \lceil \log q \rceil}, \chi},$$

and define

$$\mathbf{c}^\top := \mathbf{u}^\top (\overline{\mathbf{V}} (\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}) + \mathbf{v}^\top (\mathbf{I}_L \otimes \mathbf{R}) + \mathbf{e}^\top.$$

Then

$$(\overline{\mathbf{V}}, \mathbf{R}, \mathbf{c}^\top) \approx_c (\overline{\mathbf{V}}, \mathbf{R}, \mathbf{w}^\top),$$

where $\mathbf{w} \sim \mathbb{Z}_q^{L^2 n \lceil \log q \rceil}$ is independent and uniform. The distribution on the right is the uniform L -decomposed LWE distribution with parameters (n, k, χ, σ, q) .

Proof. It suffices to prove

$$(\mathbf{R}, \mathbf{v}^\top (\mathbf{I}_L \otimes \mathbf{R}) + \mathbf{e}^\top) \approx_c (\mathbf{R}, \mathbf{w}^\top), \tag{37}$$

where $\mathbf{w} \sim \mathbb{Z}_q^{L^2 n \lceil \log q \rceil}$ is uniform and independent of $\mathbf{R}$. Indeed, append independent $\bar{\mathbf{V}} \sim \mathbb{Z}_q^{n \times Lk}$ and $\mathbf{u} \sim \mathbb{Z}_q^n$, and apply the randomized postprocessing

$$(\mathbf{R}, \mathbf{x}^\top) \mapsto (\bar{\mathbf{V}}, \mathbf{R}, \mathbf{u}^\top (\bar{\mathbf{V}}(\mathbf{I}_L \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_L)^\top \otimes \mathbf{G}) + \mathbf{x}^\top).$$

Applied to the left-hand distribution, this produces the tuple in the lemma. Applied to $(\mathbf{R}, \mathbf{w}^\top)$, its third component is uniform and independent of $(\bar{\mathbf{V}}, \mathbf{R})$. Thus the right-hand distribution becomes exactly the uniform L -decomposed LWE distribution appearing in the lemma.

It remains to prove Equation (37). Write

$$\mathbf{v}^\top = (\mathbf{v}_1^\top, \dots, \mathbf{v}_L^\top) \quad \text{and} \quad \mathbf{e}^\top = (\mathbf{e}_1^\top, \dots, \mathbf{e}_L^\top),$$

where every $\mathbf{v}_a \in \mathbb{Z}_q^k$ and every $\mathbf{e}_a \in \mathbb{Z}^{Ln \lceil \log q \rceil}$. Then

$$\mathbf{v}^\top(\mathbf{I}_L \otimes \mathbf{R}) + \mathbf{e}^\top = (\mathbf{v}_1^\top \mathbf{R} + \mathbf{e}_1^\top, \dots, \mathbf{v}_L^\top \mathbf{R} + \mathbf{e}_L^\top).$$

For $a \in \{0, \dots, L\}$, let Hyb_a be the distribution of

$$(\mathbf{R}, \mathbf{y}_1^\top, \dots, \mathbf{y}_L^\top),$$

where the first a vectors $\mathbf{y}_b$ are independent and uniform in $\mathbb{Z}_q^{Ln \lceil \log q \rceil}$, whereas, for every $b > a$,

$$\mathbf{y}_b^\top = \mathbf{v}_b^\top \mathbf{R} + \mathbf{e}_b^\top, \quad \mathbf{v}_b \sim \mathbb{Z}_q^k, \quad \mathbf{e}_b \sim \mathcal{D}_{\mathbb{Z}^{Ln \lceil \log q \rceil}, \chi}.$$

Thus Hyb_0 is the distribution on the left-hand side of Equation (37), whereas Hyb_L is the distribution on its right-hand side.

Fix $a \in [L]$. A distinguisher between Hyb_{a-1} and Hyb_a yields a distinguisher for decisional NLWE with parameters $(k, Ln \lceil \log q \rceil, \sigma, \chi, q)$. Given an NLWE challenge $(\mathbf{R}^\star, \mathbf{y}^\top)$, set $\mathbf{R} := \mathbf{R}^\star$, sample $\mathbf{y}_b \sim \mathbb{Z}_q^{Ln \lceil \log q \rceil}$ uniformly for every $b < a$, and set $\mathbf{y}_a := \mathbf{y}$. For every $b > a$, sample

$$\mathbf{v}_b \sim \mathbb{Z}_q^k, \quad \mathbf{e}_b \sim \mathcal{D}_{\mathbb{Z}^{Ln \lceil \log q \rceil}, \chi},$$

and set $\mathbf{y}_b^\top := \mathbf{v}_b^\top \mathbf{R}^\star + \mathbf{e}_b^\top$. The resulting tuple follows Hyb_{a-1} when the challenge is an NLWE sample and Hyb_a when it is uniform.

By the polynomial bounds in Theorem 6.1, the sample count $Ln \lceil \log q \rceil$ is polynomially bounded in λ . By Theorem 3.13, the hypotheses

$$k \geq 6n \lceil \log q \rceil^2, \quad \sigma = \Omega(\sqrt{n} \lceil \log q \rceil)$$

and the assumed hardness of decisional LWE with parameters $(n \lceil \log q \rceil, Ln \lceil \log q \rceil, \chi, q)$ imply that consecutive hybrids are computationally indistinguishable. Since L is polynomially bounded, $\text{Hyb}_0 \approx_c \text{Hyb}_L$. The randomized postprocessing above therefore proves the lemma. $\square$

7 A Quantum Reduction from Succinct SIS to Succinct LWE

We give a quantum reduction from succinct SIS to succinct LWE, following the template of Regev [Reg09]; see also [CLZ22, Section 3.1]. We identify each element of $\mathbb{Z}_q$ with its unique integer representative in $[-(q-1)/2, (q-1)/2]$. All register arithmetic is modulo q , and vector norms are Euclidean norms of these representatives.

Theorem 7.1. *Let $\lambda \geq 4$, let q be an odd prime, and suppose*

$$m \geq \max\{n \log_2 q, 4\lambda + 4\},$$

$$4\sqrt{\lambda + \lceil \log_2(m+1) \rceil} \leq \chi \leq \frac{q}{8\sqrt{\lambda + \lceil \log_2(m+1) \rceil}}. \quad (38)$$

For any parameters (ℓ, k, s) for which the problem is defined, a polynomial-time algorithm solving $\text{sLWE}_{n,m,q,\ell,k,s,\chi}$ with probability $\delta \geq 1/\text{poly}(\lambda)$ gives a quantum polynomial-time algorithm solving $\text{sSIS}_{n,m,q,\ell,k,s,\gamma}$ with probability at least $\delta^2 - \text{negl}(\lambda)$, provided

$$\gamma \geq \frac{q\sqrt{m}}{2\chi}. \quad (39)$$

The reduction uses the solver's circuit once coherently and once in reverse.

For a state $|v\rangle$, write $|\hat{v}\rangle$ for its quantum Fourier transform over $\mathbb{Z}_q^m$, defined on basis states by

$$|\mathbf{e}\rangle \mapsto q^{-m/2} \sum_{\mathbf{x} \in \mathbb{Z}_q^m} e^{2\pi i \langle \mathbf{e}, \mathbf{x} \rangle / q} |\mathbf{x}\rangle.$$

This transformation is unitary and can be implemented in quantum polynomial time.

Lemma 7.2. *Under Eq. (38), one can efficiently prepare a state negligibly close in trace distance to the Gaussian state*

$$|f\rangle = \frac{1}{\sqrt{\sum_{\mathbf{e} \in \mathbb{Z}_q^m} e^{-\pi \|\mathbf{e}\|_2^2 / \chi^2}}} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} e^{-\pi \|\mathbf{e}\|_2^2 / (2\chi^2)} |\mathbf{e}\rangle = \sum_{\mathbf{e}} f(\mathbf{e}) |\mathbf{e}\rangle. \quad (40)$$

Measuring $|f\rangle$ gives a distribution within negligible statistical distance of $\mathcal{D}_{\mathbb{Z}^m, \chi} \bmod q$. Writing $|\hat{f}\rangle = \sum_{\mathbf{x}} \hat{f}(\mathbf{x}) |\mathbf{x}\rangle$, we have

$$\sum_{\|\mathbf{x}\|_2 > q\sqrt{m}/(2\chi)} |\hat{f}(\mathbf{x})|^2 = \text{negl}(\lambda), \quad q^n |\hat{f}(\mathbf{0})|^2 \leq 2^{-m}. \quad (41)$$

Proof. The Grover–Rudolph algorithm, as used in [Reg09, proof of Lemma 3.12], gives each one-dimensional Gaussian and we can prepare the above state by taking the tensor product.

Write $\rho_u(S) = \sum_{a \in S} e^{-\pi a^2 / u^2}$. For any u satisfying the bounds on χ in Eq. (38), Poisson summation and the Gaussian integral give

$$u \leq \rho_u(\mathbb{Z}) \leq 1 + u, \quad \rho_u(a + q\mathbb{Z}) \leq \rho_u(q\mathbb{Z}) \leq 1 + u/q.$$

For $|a| > q/2$,

$$e^{-\pi a^2/u^2} = e^{-\pi a^2/(2u^2)} e^{-\pi a^2/(2u^2)} \leq e^{-\pi q^2/(8u^2)} e^{-\pi a^2/(2u^2)}.$$

Summing this inequality and dividing by $\rho_u(\mathbb{Z})$ gives

$$\Pr_{a \sim \mathcal{D}_{\mathbb{Z},u}} [|a| > q/2] \leq e^{-\pi q^2/(8u^2)} \frac{\rho_{\sqrt{2}u}(\mathbb{Z})}{\rho_u(\mathbb{Z})} \leq 2e^{-\pi q^2/(8u^2)}, \quad \rho_u(\mathbb{Z}_q) \geq 0.9u. \quad (42)$$

Since $q^2/u^2 \geq 64\lambda$, a union bound proves the assertion about measuring $|f\rangle$.

Consider the two one-dimensional, unnormalized states

$$|f_u\rangle = \sum_{a \in \mathbb{Z}_q} e^{-\pi a^2/(2u^2)} |a\rangle, \quad |g_u\rangle = \sum_{a \in \mathbb{Z}_q} \rho_{\sqrt{2}u}(a + q\mathbb{Z}) |a\rangle.$$

Their amplitude differences are nonnegative and at most $1 + \sqrt{2}u/q$. Hence

$$\begin{aligned} \| |g_u\rangle - |f_u\rangle \|_2^2 &\leq (1 + \sqrt{2}u/q) \sum_{|a| > q/2} e^{-\pi a^2/(2u^2)} \\ &\leq (1 + \sqrt{2}u/q)(1 + 2u)e^{-\pi q^2/(16u^2)} \\ &\leq 3ue^{-\pi q^2/(16u^2)}. \end{aligned}$$

Consequently,

$$\left\| \frac{|g_u\rangle}{\| |g_u\rangle \|_2} - \frac{|f_u\rangle}{\| |f_u\rangle \|_2} \right\|_2 \leq \frac{2\| |g_u\rangle - |f_u\rangle \|_2}{\sqrt{\rho_u(\mathbb{Z}_q)}} \leq 4e^{-\pi q^2/(32u^2)} \leq 4e^{-2\pi\lambda}. \quad (43)$$

The Poisson summation formula gives the exact one-dimensional Fourier identity

$$|\hat{g}_u\rangle = \frac{\sqrt{2}u}{\sqrt{q}} \sum_{x \in \mathbb{Z}_q} \rho_{q/(\sqrt{2}u)}(x + q\mathbb{Z}) |x\rangle = \frac{\sqrt{2}u}{\sqrt{q}} |g_{q/(2u)}\rangle.$$

Taking reciprocals of the bounds on χ in Eq. (38) gives

$$4\sqrt{\lambda + \lceil \log_2(m+1) \rceil} \leq \frac{q}{2\chi} \leq \frac{q}{8\sqrt{\lambda + \lceil \log_2(m+1) \rceil}}.$$

Thus Eq. (43) applies at both widths. Unitarity and the positive Fourier scaling factor imply

$$\frac{|\hat{g}_\chi\rangle}{\| |g_\chi\rangle \|_2} = \frac{|g_{q/(2\chi)}\rangle}{\| |g_{q/(2\chi)}\rangle \|_2}.$$

Consequently, the triangle inequality and unitarity give

$$\begin{aligned} \left\| \frac{|\hat{f}_\chi\rangle}{\| |f_\chi\rangle \|_2} - \frac{|f_{q/(2\chi)}\rangle}{\| |f_{q/(2\chi)}\rangle \|_2} \right\|_2 &\leq \left\| \frac{|\hat{f}_\chi\rangle}{\| |f_\chi\rangle \|_2} - \frac{|\hat{g}_\chi\rangle}{\| |g_\chi\rangle \|_2} \right\|_2 + \left\| \frac{|g_{q/(2\chi)}\rangle}{\| |g_{q/(2\chi)}\rangle \|_2} - \frac{|f_{q/(2\chi)}\rangle}{\| |f_{q/(2\chi)}\rangle \|_2} \right\|_2 \\ &= \left\| \frac{|f_\chi\rangle}{\| |f_\chi\rangle \|_2} - \frac{|g_\chi\rangle}{\| |g_\chi\rangle \|_2} \right\|_2 + \left\| \frac{|g_{q/(2\chi)}\rangle}{\| |g_{q/(2\chi)}\rangle \|_2} - \frac{|f_{q/(2\chi)}\rangle}{\| |f_{q/(2\chi)}\rangle \|_2} \right\|_2 \leq 8e^{-2\pi\lambda}. \end{aligned}$$

For unit states $|h\rangle, |v\rangle$, the tensor product identity

$$|h\rangle^{\otimes m} - |v\rangle^{\otimes m} = \sum_{j=0}^{m-1} |h\rangle^{\otimes j} \otimes (|h\rangle - |v\rangle) \otimes |v\rangle^{\otimes (m-j-1)}$$

gives $\| |h\rangle^{\otimes m} - |v\rangle^{\otimes m} \|_2 \leq m \| |h\rangle - |v\rangle \|_2$. Since $|f\rangle = (|f_\chi\rangle / \| |f_\chi\rangle \|_2)^{\otimes m}$,

$$\left\| |\hat{f}\rangle - \left(\frac{|f_{q/(2\chi)}\rangle}{\| |f_{q/(2\chi)}\rangle \|_2} \right)^{\otimes m} \right\|_2 \leq 8me^{-2\pi\lambda} = \text{negl}(\lambda). \quad (44)$$

For any such width u , the full discrete Gaussian satisfies

$$\Pr_{\mathbf{x} \sim \mathcal{D}_{\mathbb{Z}^m, u}} [\|\mathbf{x}\|_2 > u\sqrt{m}] \leq e^{-\pi m/2} \left(\frac{\rho_{\sqrt{2}u}(\mathbb{Z})}{\rho_u(\mathbb{Z})} \right)^m \leq [e^{-\pi/2}(\sqrt{2} + 1/u)]^m \leq 2^{-m}.$$

Conditioning every coordinate to lie in $\mathbb{Z}_q$ multiplies this bound by at most two, since Eq. (42) gives

$$\Pr_{\mathbf{x} \sim \mathcal{D}_{\mathbb{Z}^m, u}} [\mathbf{x} \in \mathbb{Z}_q^m] \geq 1 - 2me^{-\pi q^2/(8u^2)} \geq 1/2.$$

Taking $u = q/(2\chi)$ and using Eq. (44) proves the tail bound. At $\mathbf{0}$ every Fourier phase equals one, and the Gaussian amplitudes factor across coordinates, so

$$q^n |\hat{f}(\mathbf{0})|^2 = q^n \left(\frac{\rho_{\sqrt{2}\chi}(\mathbb{Z}_q)^2}{q\rho_\chi(\mathbb{Z}_q)} \right)^m.$$

The mass bounds above and $\chi \geq 8$ give

$$\frac{\rho_{\sqrt{2}\chi}(\mathbb{Z}_q)^2}{\rho_\chi(\mathbb{Z}_q)} \leq \frac{(1 + \sqrt{2}\chi)^2}{0.9\chi} \leq \frac{(\sqrt{2} + 1/8)^2}{0.9} \chi < 3\chi.$$

Finally, $m \geq n \log_2 q$ and the upper bound on χ imply $q^n \leq 2^m$ and $\chi/q \leq 1/16$. Therefore

$$q^n |\hat{f}(\mathbf{0})|^2 \leq q^n (3\chi/q)^m \leq (6\chi/q)^m \leq (3/8)^m \leq 2^{-m}.$$

□

Proof of Theorem 7.1. Fix $(\mathbf{A}, \mathbf{W}, \mathbf{T})$. Let $p(\mathbf{s} \mid \mathbf{b})$ be the solver's probability of outputting $\mathbf{s}$ on this tuple and $\mathbf{b}$, with the convention that the solver maps invalid answers to $\mathbf{0}$. Implement the solver by a unitary U that preserves $\mathbf{b}$ and retains all random coins, measurement records, and work registers. Expectations below are over the input tuple $(\mathbf{A}, \mathbf{W}, \mathbf{T})$. The reduction prepares

$$q^{-n/2} \sum_{\mathbf{s} \in \mathbb{Z}_q^n, \mathbf{e} \in \mathbb{Z}_q^m} f(\mathbf{e}) |\mathbf{A}^\top \mathbf{s} + \mathbf{e}, \mathbf{e}\rangle. \quad (45)$$

This can be done efficiently by preparing the uniform secret and the Gaussian state of Lemma 7.2, computing $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e}$, and erasing the secret register by reversibly solving $\mathbf{A}^\top \mathbf{s} = \mathbf{b} - \mathbf{e}$ using

Gaussian elimination; full row rank guarantees a unique solution. Apply U , subtract $\mathbf{b} - \mathbf{A}^\top \mathbf{s}'$ from the error register using the solver's answer $\mathbf{s}'$, and apply $U^\dagger$. Measure the error and ancilla registers and abort unless both are zero. Fourier transform the remaining register, measure $\mathbf{x}$, and return it if

$$\mathbf{Ax} = \mathbf{0} \pmod{q}, \quad \mathbf{x} \neq \mathbf{0}, \quad \|\mathbf{x}\|_2 \leq \gamma. \quad (46)$$

Otherwise return failure.

The error register becomes zero exactly for the correct answer, since

$$\mathbf{e} - (\mathbf{b} - \mathbf{A}^\top \mathbf{s}') = \mathbf{A}^\top (\mathbf{s}' - \mathbf{s}) = \mathbf{0} \iff \mathbf{s}' = \mathbf{s}.$$

Write Π_s for projection onto solver answer $\mathbf{s}$. After $U^\dagger$, the additional projection of all solver work registers onto zero contributes the amplitude

$$\langle \mathbf{b}, 0 | U^\dagger \Pi_s U | \mathbf{b}, 0 \rangle = p(\mathbf{s} | \mathbf{b}).$$

Consequently, the *unnormalized* state after both zero tests is

$$\begin{aligned} |v\rangle &:= q^{-n/2} \sum_{\mathbf{b}, \mathbf{s}} f(\mathbf{b} - \mathbf{A}^\top \mathbf{s}) p(\mathbf{s} | \mathbf{b}) |\mathbf{b}\rangle, \\ \| |v\rangle \|_2^2 &= \Pr[\text{both zero tests pass}] \leq 1. \end{aligned} \quad (47)$$

Define the unnormalized state

$$|h\rangle := q^{-n/2} \sum_{\mathbf{s}, \mathbf{e}} f(\mathbf{e}) |\mathbf{A}^\top \mathbf{s} + \mathbf{e}\rangle.$$

Since all amplitudes in these two states are nonnegative, we obtain

$$\mathbb{E} \langle h | v \rangle \geq \mathbb{E} q^{-n} \sum_{\mathbf{b}, \mathbf{s}} f(\mathbf{b} - \mathbf{A}^\top \mathbf{s})^2 p(\mathbf{s} | \mathbf{b}) \geq \delta - \text{negl}(\lambda). \quad (48)$$

The sum is exactly the solver's success probability for a uniform secret and error density $f(\mathbf{e})^2$, which is statistically close to the required Gaussian by Lemma 7.2. The Fourier transform of $|h\rangle$ is

$$\begin{aligned} |\hat{h}\rangle &= q^{-(m+n)/2} \sum_{\mathbf{x}, \mathbf{s}, \mathbf{e}} f(\mathbf{e}) e^{2\pi i \langle \mathbf{A}^\top \mathbf{s} + \mathbf{e}, \mathbf{x} \rangle / q} |\mathbf{x}\rangle \\ &= q^{-n/2} \sum_{\mathbf{x}} \hat{f}(\mathbf{x}) \prod_{j=1}^n \left(\sum_{s_j \in \mathbb{Z}_q} e^{2\pi i s_j (\mathbf{Ax})_j / q} \right) |\mathbf{x}\rangle = q^{n/2} \sum_{\mathbf{Ax}=\mathbf{0}} \hat{f}(\mathbf{x}) |\mathbf{x}\rangle. \end{aligned} \quad (49)$$

Let Q_A project onto the basis states satisfying Eq. (46); thus $\mathbb{E} \|Q_A |\hat{v}\rangle\|_2^2$ is exactly the algorithm's success probability, including both zero tests. For every nonzero $\mathbf{x}$, uniform full row rank $\mathbf{A}$ satisfies

$$\Pr[\mathbf{Ax} = \mathbf{0}] = \frac{q^{m-n} - 1}{q^m - 1} \leq q^{-n}.$$

Hence Eqs. (41) and (49) give

$$\begin{aligned}\mathbb{E}\| |h\rangle\|_2^2 &\leq q^n |\hat{f}(\mathbf{0})|^2 + \sum_{\mathbf{x} \neq \mathbf{0}} |\hat{f}(\mathbf{x})|^2 \leq 1 + 2^{-m}, \\ \mathbb{E}\|(I - Q_A)|\hat{h}\rangle\|_2^2 &\leq q^n |\hat{f}(\mathbf{0})|^2 + \sum_{\|\mathbf{x}\|_2 > \gamma} |\hat{f}(\mathbf{x})|^2 = \text{negl}(\lambda).\end{aligned}\tag{50}$$

Using unitarity and Cauchy–Schwarz, first for each tuple and then for the expectation, we obtain

$$\begin{aligned}\delta - \text{negl}(\lambda) &\leq \mathbb{E}\langle \hat{h} | \hat{v} \rangle \\ &= \mathbb{E}[\langle \hat{h} | Q_A | \hat{v} \rangle + \langle \hat{h} | (I - Q_A) | \hat{v} \rangle] \\ &\leq \mathbb{E}[|\langle \hat{h} | Q_A | \hat{v} \rangle| + |\langle \hat{h} | (I - Q_A) | \hat{v} \rangle|] \\ &\leq \mathbb{E}[\|Q_A |\hat{h}\rangle\|_2 \|Q_A |\hat{v}\rangle\|_2 + \|(I - Q_A) |\hat{h}\rangle\|_2 \|(I - Q_A) |\hat{v}\rangle\|_2] \\ &\leq \sqrt{\mathbb{E}\| |h\rangle\|_2^2 \mathbb{E}\|Q_A |\hat{v}\rangle\|_2^2} + \sqrt{\mathbb{E}\|(I - Q_A) |\hat{h}\rangle\|_2^2 \mathbb{E}\| |\hat{v}\rangle\|_2^2} \\ &\leq \sqrt{(1 + 2^{-m}) \mathbb{E}\|Q_A |\hat{v}\rangle\|_2^2} + \text{negl}(\lambda).\end{aligned}$$

Therefore

$$\Pr[\text{SIS success}] = \mathbb{E}\|Q_A |\hat{v}\rangle\|_2^2 \geq \frac{(\delta - \text{negl}(\lambda))^2}{1 + 2^{-m}} \geq \delta^2 - \text{negl}(\lambda).$$

All steps take polynomial time. Implementing the state preparation and Fourier transform to negligible total error changes this unconditional success probability only negligibly. $\square$

8 Generalizations of Decomposed LWE

In this section, we prove a series of simple but useful reductions between variants of succinct LWE and decomposed LWE introduced below. We summarize the relations between them in Figure 2.

Before presenting our results, let us generalize the definitions of succinct LWE and decomposed LWE by allowing the secret $\mathbf{s}$ to be sampled from an arbitrary distribution Sample .

Definition 8.1 ((ℓ , Sample)-Succinct LWE). *Fix a security parameter λ and parameters n, m, k, χ, σ , and q . Let $\text{Sample}(1^n)$ be any efficiently sampleable distribution over $\mathbb{Z}_q^n$. The (decision) (ℓ , Sample)-succinct LWE assumption stipulates that*

$$(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{W}, \mathbf{T}) \approx_c (\mathbf{A}, \mathbf{c}^\top, \mathbf{W}, \mathbf{T}),$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \text{Sample}(1^n)$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi}$, $\mathbf{c} \leftarrow \mathbb{Z}_q^m$, and $\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n \times k}$ are independent, and $\mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$. When Sample is not specified, we refer to the version where $\mathbf{s}$ is sampled uniformly over $\mathbb{Z}_q^n$.

Definition 8.2 ((ℓ , Sample)-Decomposed LWE). *Fix a security parameter λ and parameters n, k, χ, σ , and q . Let $\text{Sample}(1^n)$ be any efficiently sampleable distribution over $\mathbb{Z}_q^n$. The (decision) (ℓ , Sample)-decomposed LWE assumption stipulates that*

$$(\mathbf{W}, \mathbf{R}, \mathbf{s}^\top (\mathbf{W}(\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G}) + \mathbf{e}^\top) \approx_c (\mathbf{W}, \mathbf{R}, \mathbf{c}^\top),$$




Figure 2: Relations between the variants of succinct LWE and decomposed LWE studied in this section. Arrows denote the existence of a reduction that solves instances of the problem at the tail using a solver for the problem at its head. Blue arrows indicate our contributions; black arrows indicate prior or elementary reductions. The dashed arrow denotes the quantum reduction to the separately shown search problem. The classical LWE-to-SIS arrow uses the noise and solution norm conditions of Lemma 5.6. All arrows require the parameter conditions of their cited results. The gray box at the top represents SIS variants of the assumptions; the four gray boxes below represent decisional LWE variants that differ in the distribution of the secret: box 2 uses a uniformly random secret, box 3 uses a discrete Gaussian secret, box 4 uses a uniformly random secret with a fixed public prefix, box 5 uses a discrete Gaussian secret with a fixed public prefix. Extended succinct LWE is introduced in Definition 8.3, with its SIS analog in Remark 8.5, while generalized decomposed LWE is introduced in Definition 8.7, with its SIS analog in Remark 8.9.

where $\mathbf{W} \leftarrow \mathbb{Z}_q^{n \times \ell k}$, $\mathbf{R} \leftarrow \mathcal{D}_{\mathbb{Z}^{k \times \ell n \lceil \log q \rceil}, \sigma}$, $\mathbf{s} \leftarrow \text{Sample}(1^n)$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^{\ell^2 n \lceil \log q \rceil}, \chi}$, and $\mathbf{c} \leftarrow \mathbb{Z}_q^{\ell^2 n \lceil \log q \rceil}$ are independent. When Sample is not specified, we refer to the version where $\mathbf{s}$ is sampled uniformly over $\mathbb{Z}_q^n$.

8.1 Extended Succinct LWE

We first introduce *extended succinct LWE*. This states that succinct LWE holds even if we provide the adversary with additional LWE samples under the succinct LWE secret.

Definition 8.3 (Extended ℓ -Succinct LWE). *Fix a security parameter λ and parameters n, m, k, χ, σ , and q . Let r be a nonnegative integer and let $\text{Sample}(1^n)$ be an efficiently sampleable distribution over $\mathbb{Z}_q^n$. The (decision) r -extended (ℓ, Sample) -succinct LWE assumption stipulates that*

$$(\mathbf{A}, \mathbf{M}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{M} + \tilde{\mathbf{e}}^\top, \mathbf{W}, \mathbf{T}) \approx_c (\mathbf{A}, \mathbf{M}, \mathbf{c}^\top, \mathbf{u}^\top, \mathbf{W}, \mathbf{T}),$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{M} \leftarrow \mathbb{Z}_q^{n \times r}$, $\mathbf{s} \leftarrow \text{Sample}(1^n)$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi}$, $\tilde{\mathbf{e}} \leftarrow \mathcal{D}_{\mathbb{Z}^r, \chi}$, $\mathbf{c} \leftarrow \mathbb{Z}_q^m$, $\mathbf{u} \leftarrow \mathbb{Z}_q^r$, and $\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n \times k}$ are independent, and $\mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$.

Lemma 8.4 (Extended Succinct LWE). *Let q be an odd prime, let $n, m, k, \ell \geq 1$ and $r \geq 0$ be $\text{poly}(\lambda)$ -bounded integers where $n = \Omega(\lambda)$, and $\chi, \chi', \sigma > 0$ reals such that*

- (i) $m \geq 3n \lceil \log q \rceil + 3\lambda$;
- (ii) $\chi \geq \sqrt{\lambda}$ and $\chi' \geq 3\chi \sqrt{m(r+1)}$;
- (iii) $\sigma \geq (\ell m + k) \log(\ell n)$.

Let $\text{Sample}(1^n)$ be an efficiently sampleable distribution over $\mathbb{Z}_q^n$. Then, every probabilistic polynomial-time distinguisher for r -extended (ℓ, Sample) -succinct LWE with parameters $(n, m, k, \chi', \sigma, q)$ yields a probabilistic polynomial-time distinguisher for (ℓ, Sample) -succinct LWE with parameters $(n, m, k, \chi, \sigma, q)$ with only a $2^{-\Omega(\lambda)}$ loss in distinguishing advantage.

Proof. We show a reduction from (ℓ, Sample) -succinct LWE to extended (ℓ, Sample) -succinct LWE. Suppose that the challenger provided us with $(\mathbf{A}, \mathbf{c}^\top, \mathbf{W}, \mathbf{T})$ where $\mathbf{D} := [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$ satisfies $\mathbf{D}\mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G}$, and $\mathbf{c}$ is either sampled uniformly at random or $\mathbf{c}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$ where $\mathbf{s} \leftarrow \text{Sample}(1^n)$, and $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi}$ is a noise term.

In the reduction, we sample $\mathbf{S} \leftarrow \{0, 1\}^{m \times r}$ and compute $\mathbf{v}^\top = \text{ReRand}([\mathbf{I}_m \mid \mathbf{S}], \mathbf{c}^\top, \chi, \chi'/(2\chi))$. Then, we output $(\mathbf{A}, \mathbf{AS}, \mathbf{v}_0^\top, \mathbf{v}_1^\top, \mathbf{W}, \mathbf{T})$ where $\mathbf{v}_0$ consists of the first m entries of $\mathbf{v}$ and $\mathbf{v}_1$ consists of the remaining r . Notice that by Lemma 3.4, when $\mathbf{c}^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$,

$$\mathbf{v}^\top = \mathbf{s}^\top [\mathbf{A} \mid \mathbf{AS}] + (\mathbf{e}')^\top,$$

where $\mathbf{e}'$ is statistically close to a sample from $\mathcal{D}_{\mathbb{Z}^{m+r}, \chi'}$. Moreover, by the leftover hash lemma (Lemma 3.2), we have that $\mathbf{AS}$ is statistically close to a random matrix in $\mathbb{Z}_q^{n \times r}$.

If instead $\mathbf{c}$ is random, by Lemma 3.4, we have that

$$\mathbf{v}^\top = (\mathbf{c} - \mathbf{e})^\top [\mathbf{I}_m \mid \mathbf{S}] + (\mathbf{e}')^\top.$$

By the leftover hash lemma (Lemma 3.2), we conclude that $(\mathbf{v}, \mathbf{AS})$ are statistically close to a random sample from $\mathbb{Z}_q^{m+r} \times \mathbb{Z}_q^{n \times r}$. $\square$

Remark 8.5. The same construction reduces succinct SIS to its extended version, whose goal is to find a nonzero short integer vector $\mathbf{x} = (\mathbf{x}_0, \mathbf{x}_1) \in \mathbb{Z}^m \times \mathbb{Z}^r$ satisfying $[\mathbf{A} \mid \mathbf{M}]\mathbf{x} = \mathbf{0} \bmod q$. Give the solver $(\mathbf{A}, \mathbf{AS}, \mathbf{W}, \mathbf{T})$ and return $\mathbf{x}_0 + \mathbf{S}\mathbf{x}_1$. This vector satisfies the required kernel relation and $\|\mathbf{x}_0 + \mathbf{S}\mathbf{x}_1\|_2 \leq \sqrt{mr+1}\|\mathbf{x}\|_2$.

If $\mathbf{x}_1 = \mathbf{0}$, the output is nonzero whenever the solver succeeds. Otherwise, fix an index j with $(\mathbf{x}_1)_j \neq 0$. Conditional on the public tuple, the solver's coins, and all other columns of $\mathbf{S}$, a zero output determines its j th column uniquely over the integers. For fixed $\mathbf{A}$, a uniform binary column is uniform on the fiber determined by its image under $\mathbf{A}$. The average reciprocal fiber size is at most $q^n 2^{-m}$, since there are at most q^n nonempty fibers. A union bound over the r possible indices therefore bounds the probability of a zero output by $r q^n 2^{-m}$. Thus a successful extended SIS solution yields a valid succinct SIS solution except with negligible probability.

8.2 Succinct LWE with a Small Secret

Next, we show that the hardness of succinct LWE also implies hardness when the secret is sampled from a discrete Gaussian distribution.

Lemma 8.6 (Succinct LWE with a Small Secret). *Let q be an odd prime, let $n, m, k, \ell \geq 1$ be $\text{poly}(\lambda)$ -bounded integers where $n = \Omega(\lambda)$, and $\chi, \chi', \sigma, \sigma' > 0$ reals such that*

- (i) $m \geq 3n\lceil \log q \rceil + 3\lambda$;
- (ii) $\chi \geq \sqrt{\lambda}$ and $\chi' \geq 4\chi\sqrt{m(n+\lambda)}$;
- (iii) $\sigma' \geq (\ell m + k)\sigma\sqrt{\lambda}\log(\ell n)$ and $\sigma \geq (\ell m + k)\log(\ell n)$.

Then, every probabilistic polynomial-time distinguisher for $(\ell, \mathcal{D}_{\mathbb{Z}^n, \chi'})$ -succinct LWE with parameters $(n, m, k, \chi', \sigma', q)$ yields a probabilistic polynomial-time distinguisher for ℓ -succinct LWE with parameters $(n, m, k, \chi, \sigma, q)$ with only a $2^{-\Omega(\lambda)}$ loss in distinguishing advantage.

Proof. Set $r = n + \lambda$. Apply Lemma 8.4 with output noise width χ' ; its noise requirement follows from $3\sqrt{r+1} \leq 4\sqrt{r}$. It therefore suffices to reduce r -extended succinct LWE with noise width χ' to $(\ell, \mathcal{D}_{\mathbb{Z}^n, \chi'})$ -succinct LWE. Suppose that the challenger provides $(\mathbf{A}, \mathbf{M}, \mathbf{u}_0^\top, \mathbf{u}_1^\top, \mathbf{W}, \mathbf{T})$, where $\mathbf{D} := [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$ satisfies $\mathbf{D}\mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G}$. In the planted case, $\mathbf{u}_0^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}_0^\top$ and $\mathbf{u}_1^\top = \mathbf{s}^\top \mathbf{M} + \mathbf{e}_1^\top$, where $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi'}$ and $\mathbf{e}_1 \leftarrow \mathcal{D}_{\mathbb{Z}^r, \chi'}$ are independent. In the uniform case, $\mathbf{u}_0$ and $\mathbf{u}_1$ are independent and uniform.

Choose an invertible $n \times n$ submatrix $\mathbf{C}$ of $\mathbf{M}$. Such a choice exists except with probability at most $\frac{q}{q-1}q^{n-r-1}$; otherwise, output a random bit. Let $\mathbf{v}$ be the subvector of $\mathbf{u}_1$ corresponding to these columns. By Lemma 4.9 and the lower bound on σ , the Gaussian coordinate tail bound gives $\|\mathbf{T}\| \leq \sigma\sqrt{\lambda}$ except with probability $2^{-\Omega(\lambda)}$. Thus the stated lower bound on σ' satisfies the sampling condition in Lemma 3.6. Sample $\mathbf{T}' \leftarrow \text{SamplePre}(\mathbf{D}, \mathbf{T}, \mathbf{I}_\ell \otimes \mathbf{C}\mathbf{G}, \sigma')$ and output

$$(\mathbf{C}^{-1}\mathbf{A}, -\mathbf{u}_0^\top + \mathbf{v}^\top \mathbf{C}^{-1}\mathbf{A}, (\mathbf{I}_\ell \otimes \mathbf{C}^{-1})\mathbf{W}, \mathbf{T}').$$

The matrices $\mathbf{C}^{-1}\mathbf{A}$ and $(\mathbf{I}_\ell \otimes \mathbf{C}^{-1})\mathbf{W}$ are independent and uniform. By Lemma 3.6, $\mathbf{T}'$ is statistically close to a Gaussian sample of width σ' conditioned on $\mathbf{D}\mathbf{T}' = \mathbf{I}_\ell \otimes \mathbf{C}\mathbf{G}$. Multiplying this identity on the left by $\mathbf{I}_\ell \otimes \mathbf{C}^{-1}$ gives

$$[\mathbf{I}_\ell \otimes (\mathbf{C}^{-1}\mathbf{A}) \mid (\mathbf{I}_\ell \otimes \mathbf{C}^{-1})\mathbf{W}]\mathbf{T}' = \mathbf{I}_\ell \otimes \mathbf{G},$$

so the output has the required public distribution.

In the planted case, write $\mathbf{v}^\top = \mathbf{s}^\top \mathbf{C} + \mathbf{e}^\top$. Since the columns are chosen using only $\mathbf{M}$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \chi'}$ is independent of the output public matrices and of $\mathbf{e}_0$. The identity

$$-\mathbf{u}_0^\top + \mathbf{v}^\top \mathbf{C}^{-1} \mathbf{A} = \mathbf{e}^\top \mathbf{C}^{-1} \mathbf{A} - \mathbf{e}_0^\top$$

therefore gives the desired secret and error distributions, using symmetry of the discrete Gaussian. In the uniform case, the same expression is uniform because $\mathbf{u}_0$ is uniform and independent. $\square$

8.3 Generalized Decomposed LWE

Next, we consider a variant of decomposed LWE, where we generate the LWE matrix by shifting $\mathbf{W}(\mathbf{I}_\ell \otimes \mathbf{R})$ by an arbitrary term $\mathbf{F}$ (which may depend on $\mathbf{W}$, but not on $\mathbf{R}$), instead of $\text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{G}$ as in the standard version of the assumption. We show that even this variant of the assumption is implied by the hardness of decomposed LWE. Variants of this assumption have been used in [CHW25, AMR26b].

Definition 8.7 (Generalized ℓ -Decomposed LWE). *Fix a security parameter λ and parameters n, k, t, χ, σ , and q . Let r be a nonnegative integer and let $\text{Matrix} : \mathbb{Z}_q^{n \times \ell k} \times \mathbb{Z}_q^{n \times r} \rightarrow \mathbb{Z}_q^{n \times \ell t}$ be an efficiently computable deterministic function, and let $\text{Sample}(1^n)$ be an efficiently sampleable distribution over $\mathbb{Z}_q^n$. The (decision) $(r, \text{Matrix}, \text{Sample})$ -generalized ℓ -decomposed LWE assumption stipulates that*

$$\left(\mathbf{W}, \mathbf{R}, \mathbf{M}, \underbrace{\mathbf{s}^\top (\mathbf{W}(\mathbf{I}_\ell \otimes \mathbf{R}) + \mathbf{F})}_{=: \mathbf{D}_F} + \mathbf{e}^\top, \mathbf{s}^\top \mathbf{M} + \tilde{\mathbf{e}}^\top \right) \approx_c \left(\mathbf{W}, \mathbf{R}, \mathbf{M}, \mathbf{u}^\top, \mathbf{v}^\top \right),$$

where $\mathbf{W} \leftarrow \mathbb{Z}_q^{n \times \ell k}$, $\mathbf{R} \leftarrow \mathcal{D}_{\mathbb{Z}^{k \times t}, \sigma}$, $\mathbf{M} \leftarrow \mathbb{Z}_q^{n \times r}$, $\mathbf{F} \leftarrow \text{Matrix}(\mathbf{W}, \mathbf{M})$, $\mathbf{s} \leftarrow \text{Sample}(1^n)$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^{\ell t}, \chi}$, $\tilde{\mathbf{e}} \leftarrow \mathcal{D}_{\mathbb{Z}^r, \chi'}$, $\mathbf{u} \leftarrow \mathbb{Z}_q^{\ell t}$ and $\mathbf{v} \leftarrow \mathbb{Z}_q^r$. The search version uses the planted tuple and asks to recover $\mathbf{s}$.

Lemma 8.8 (Generalized Decomposed LWE). *Let q be an odd prime, let $n, m, k, t, \ell \geq 1$ and $r \geq 0$ be poly(λ)-bounded integers where $n = \Omega(\lambda)$, and $\chi, \chi', \sigma, \sigma' > 0$ reals such that*

- (i) $m \geq 3n \lceil \log q \rceil + 3\lambda$;
- (ii) $\chi \geq \sqrt{\lambda}$ and $\chi' \geq 16\chi\sigma' \sqrt{\lambda m \ell t (mr + 1)}$;
- (iii) $\sigma' \geq (\ell m + k)\sigma \sqrt{\lambda} \log(\ell n)$ and $\sigma \geq (\ell m + k) \log(\ell n)$.

Let $\text{Matrix} : \mathbb{Z}_q^{n \times \ell k} \times \mathbb{Z}_q^{n \times r} \rightarrow \mathbb{Z}_q^{n \times \ell t}$ be a deterministic efficiently computable function and let $\text{Sample}(1^n)$ be an efficiently sampleable distribution over $\mathbb{Z}_q^n$. Then, every probabilistic polynomial-time distinguisher for $(r, \text{Matrix}, \text{Sample})$ -generalized ℓ -decomposed LWE with parameters $(n, k, t, \chi', \sigma', q)$ yields a probabilistic polynomial-time distinguisher for (ℓ, Sample) -succinct LWE with parameters $(n, m, k, \chi, \sigma, q)$ with only a $2^{-\Omega(\lambda)}$ loss in distinguishing advantage.

The same reduction applies to the search versions: a solver that recovers the generalized decomposed LWE secret with probability δ yields a solver for search succinct LWE with success at least $\delta - 2^{-\Omega(\lambda)}$.

Proof. We first apply the construction in the proof of Lemma 8.4 with intermediate noise width $\tau := 3\chi\sqrt{mr+1}$. This width suffices because $\|[\mathbf{I}_m \mid \mathbf{S}]\|_{2 \rightarrow 2} \leq \sqrt{mr+1}$ for $\mathbf{S} \in \{0, 1\}^{m \times r}$. It remains to reduce the resulting r -extended (ℓ, Sample) -succinct LWE instance to generalized decomposed LWE. Suppose that the challenger provided us with $(\mathbf{A}, \mathbf{M}, \mathbf{u}_0^\top, \mathbf{u}_1^\top, \mathbf{W}, \mathbf{T})$ where $\mathbf{D} := [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$ satisfies $\mathbf{D}\mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G}$, and $\mathbf{u}_0, \mathbf{u}_1$ are either sampled uniformly at random or $\mathbf{u}_0^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}_0^\top$, $\mathbf{u}_1^\top = \mathbf{s}^\top \mathbf{M} + \mathbf{e}_1^\top$ where $\mathbf{s} \leftarrow \text{Sample}(1^n)$ and $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^m, \tau}$, $\mathbf{e}_1 \leftarrow \mathcal{D}_{\mathbb{Z}^r, \tau}$ are noise terms.

Split $\mathbf{W}$ into blocks $\mathbf{W}_1, \dots, \mathbf{W}_\ell \in \mathbb{Z}_q^{n \times k}$ and define $\bar{\mathbf{W}} := [\mathbf{W}_1 \mid \dots \mid \mathbf{W}_\ell]$. Compute $\mathbf{F} \leftarrow \text{Matrix}(\bar{\mathbf{W}}, \mathbf{M})$ and split it into ℓ blocks $\mathbf{F}_1, \dots, \mathbf{F}_\ell \in \mathbb{Z}_q^{n \times t}$. Let $\bar{\mathbf{F}} \in \mathbb{Z}_q^{\ell n \times t}$ be their vertical concatenation. We compute $\mathbf{T}' \in \mathbb{Z}^{(\ell m + k) \times t}$ as $\mathbf{T}' \leftarrow \text{SamplePre}(\mathbf{D}, \mathbf{T}, \bar{\mathbf{F}}, \sigma')$. Then, we split $\mathbf{T}'$ into blocks $\mathbf{V}_i \in \mathbb{Z}^{m \times t}$ and $\mathbf{R} \in \mathbb{Z}^{k \times t}$ such that

$$\mathbf{A}\bar{\mathbf{V}} = \underbrace{-\bar{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \mathbf{F}}_{=: \mathbf{D}_F}, \quad \bar{\mathbf{V}} = [\mathbf{V}_1 \mid \dots \mid \mathbf{V}_\ell].$$

The output of the reduction is the tuple $(\bar{\mathbf{W}}, -\mathbf{R}, \mathbf{M}, \mathbf{v}_0^\top, \mathbf{v}_1^\top)$ where

$$\mathbf{v}_0^\top = \text{ReRand}(\bar{\mathbf{V}}, \mathbf{u}_0^\top, \tau, \chi'/(2\tau)), \quad \mathbf{v}_1^\top = \text{ReRand}(\mathbf{I}_r, \mathbf{u}_1^\top, \tau, \chi'/(2\tau)).$$

We justify the public distribution jointly. Fix $\varepsilon = 2^{-\lambda}$. Except with probability $2^{-\Omega(\lambda)}$ over $\mathbf{A}$, the smoothing argument in Lemma 4.9 applies both to $\Lambda^\perp(\mathbf{A})$ and to $\Lambda^\perp(\mathbf{D})$. The stated lower bounds on σ and σ' exceed the corresponding smoothing bounds. Completing the square and comparing shifted coset masses gives the Gaussian coordinate tail bounds, which imply $\|\mathbf{T}\| \leq \sigma\sqrt{\lambda}$ except with probability $2^{-\Omega(\lambda)}$, so Lemma 3.6 allows us to replace $\mathbf{T}'$ by independent canonical Gaussian preimages of width σ' .

For this canonical distribution, conditioning on $\mathbf{R}$ leaves independent Gaussian columns of the blocks $\mathbf{V}_i$, with syndromes $\mathbf{F}_i - \mathbf{W}_i \mathbf{R}$ under $\mathbf{A}$. Smoothing makes the Gaussian masses of all these cosets agree within a factor $1 + O(\varepsilon)$. Consequently, the joint distribution is within $2^{-\Omega(\lambda)}$ of the experiment that samples $\mathbf{R} \leftarrow \mathcal{D}_{\mathbb{Z}^{k \times t}, \sigma'}$ independently of $(\mathbf{A}, \mathbf{W}, \mathbf{M})$ and then samples each column of $\bar{\mathbf{V}}$ from its prescribed Gaussian coset under $\mathbf{A}$. This proves the required public distribution, including the sign change of $\mathbf{R}$ by Gaussian symmetry. Moreover, Lemma 3.20 and a union bound give $\|\bar{\mathbf{V}}\|_{2 \rightarrow 2} \leq \sigma'\sqrt{m\ell t}$ except with probability $2^{-\Omega(\lambda)}$. The noise hypothesis therefore implies $\chi' > 2\tau \max\{\|\bar{\mathbf{V}}\|_{2 \rightarrow 2}, 1\}$, as required for both rerandomization calls. When $r = 0$, the second call and its output are empty.

Finally, by Lemma 3.4, when $\mathbf{u}_0^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e}_0^\top$ and $\mathbf{u}_1^\top = \mathbf{s}^\top \mathbf{M} + \mathbf{e}_1^\top$, we have that

$$\mathbf{v}_0^\top = \mathbf{s}^\top \underbrace{\mathbf{A}\bar{\mathbf{V}}}_{\mathbf{D}_F} + (\mathbf{e}_0')^\top, \quad \mathbf{v}_1^\top = \mathbf{s}^\top \mathbf{M} + (\mathbf{e}_1')^\top,$$

where $\mathbf{e}_0'$ and $\mathbf{e}_1'$ are statistically close to samples from $\mathcal{D}_{\mathbb{Z}^{\ell t}, \chi'}$ and $\mathcal{D}_{\mathbb{Z}^r, \chi'}$. If instead $\mathbf{u}_0$ and $\mathbf{u}_1$ are uniform, write $\mathbf{h} := \mathbf{u}_0 - \mathbf{e}_0$, introducing an independent $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^m, \tau}$. Then $\mathbf{h}$ is uniform and independent of $\mathbf{e}_0$ and all public matrices. The rerandomization guarantee lets us analyze $\mathbf{v}_0^\top = \mathbf{h}^\top \bar{\mathbf{V}} + (\mathbf{e}_0')^\top$ with an independent error of width χ' , up to statistical error. Apply the smoothing argument from Lemma 4.9 to the random $(n+1) \times m$ matrix $[\mathbf{A}^\top \mid \mathbf{h}]^\top$. The bound $m \geq 3n[\log q] + 3\lambda$ implies $m \geq 2(n+1)\log q + 2\lambda$ for all sufficiently large λ , since $n = \Omega(\lambda)$. Thus, except with probability $2^{-\Omega(\lambda)}$, Gaussian coset masses for this augmented matrix are flat at width σ' . For every

fixed syndrome under $\mathbf{A}$, the additional image under $\mathbf{h}^\top$ is therefore statistically close to uniform. Applying this fact to the conditionally independent columns of $\bar{\mathbf{V}}$ proves joint uniformity of $\mathbf{h}^\top \bar{\mathbf{V}}$, retaining $(\bar{\mathbf{W}}, -\mathbf{R}, \mathbf{M})$. The identity transformation in the second rerandomization call preserves uniformity independently, so both output vectors have the required uniform distribution up to $2^{-\Omega(\lambda)}$ error. All comparisons in the planted experiment hold jointly with the secret $\mathbf{s}$, which is unchanged. Returning the oracle's recovered secret therefore proves the search claim. $\square$

Remark 8.9. The same construction reduces extended succinct SIS to generalized decomposed SIS. Given $(\bar{\mathbf{W}}, -\mathbf{R}, \mathbf{M})$, the latter problem asks for a nonzero short integer vector $\mathbf{x} = (\mathbf{x}_0, \mathbf{x}_1) \in \mathbb{Z}^{\ell t} \times \mathbb{Z}^r$ satisfying $\mathbf{D}_F \mathbf{x}_0 + \mathbf{M} \mathbf{x}_1 = \mathbf{0} \bmod q$. Return $(\bar{\mathbf{V}} \mathbf{x}_0, \mathbf{x}_1)$. The identity $\mathbf{A} \bar{\mathbf{V}} = \mathbf{D}_F$ proves the required extended SIS kernel relation. The preceding proof gives $\|\bar{\mathbf{V}}\|_{2 \rightarrow 2} \leq \sigma' \sqrt{m \ell t}$ except with negligible probability, so the output norm is at most $\max\{\sigma' \sqrt{m \ell t}, 1\} \|\mathbf{x}\|_2$.

If $\mathbf{x}_1 \neq \mathbf{0}$, the output is nonzero. Otherwise, success implies $\mathbf{x}_0 \neq \mathbf{0}$. Condition on $\mathbf{A}$, the public tuple, the solver's coins, and all but one column of $\bar{\mathbf{V}}$ whose coefficient in $\mathbf{x}_0$ is nonzero. A zero output determines that column uniquely over the integers. In the canonical experiment above, it is a Gaussian sample of width σ' from a prescribed coset of $\Lambda^\perp(\mathbf{A})$. Since $\sigma' \geq 2\eta_\varepsilon(\Lambda^\perp(\mathbf{A}))$, Lemma 3.20 bounds every point probability by $2^{-m}(1 + \varepsilon)/(1 - \varepsilon)$. A union bound over the ℓt columns, together with the negligible simulation and smoothing errors, proves that the output is nonzero except with negligible probability.

8.4 Decomposed LWE with a Structured Secret

Finally, we show that decomposed LWE is hard even if we structure the LWE secret by prepending a public prefix $\mathbf{r}$ (notice that this is not straightforward as for standard LWE due to the particular structure of the decomposed LWE matrix). This may be useful, for instance, in applications that mix decomposed LWE techniques with FHE schemes where the secret key is required to have the first entry set to 1.

Lemma 8.10 (Decomposed LWE with a Structured Secret). *Let q be an odd prime, let $n', n, m, a, k, \ell \geq 1$ be $\text{poly}(\lambda)$ -bounded integers where $n = \Omega(\lambda)$, and $\chi, \chi', \sigma, \sigma' > 0$ reals such that*

- (i) $n' = n + a$;
- (ii) $m \geq 3n \lceil \log q \rceil + 3\lambda$;
- (iii) $\chi \geq \sqrt{\lambda}$ and $\chi' \geq 16\chi\sigma'\ell\sqrt{\lambda m(a+n)\lceil \log q \rceil}$;
- (iv) $\sigma' \geq (\ell m + k)\sigma\sqrt{\lambda} \log(\ell n)$ and $\sigma \geq (\ell m + k) \log(\ell n)$.

Let $\mathbf{r} \in \mathbb{Z}_q^a$ be an arbitrary vector and let $\text{Sample}(1^n)$ be an efficiently sampleable distribution over $\mathbb{Z}_q^n$. Define $\text{Sample}'(1^{n'})$ to be the distribution that outputs the concatenation of $\mathbf{r}$ and $\mathbf{s}$, where $\mathbf{s} \leftarrow \text{Sample}(1^n)$. Then, every probabilistic polynomial-time distinguisher for $(\ell, \text{Sample}')\text{-decomposed LWE}$ with parameters $(n', k, \chi', \sigma', q)$ yields a probabilistic polynomial-time distinguisher for $(\ell, \text{Sample})\text{-succinct LWE}$ with parameters $(n, m, k, \chi, \sigma, q)$ with only a $2^{-\Omega(\lambda)}$ loss in distinguishing advantage.

Proof. Let $\overline{\mathbf{G}} := [0|\mathbf{G}] \in \mathbb{Z}_q^{n \times (a+n) \lceil \log q \rceil}$ where a is the dimension of $\mathbf{r}$ and let *Matrix* be the function that constantly outputs $\text{vec}(\mathbf{I}_\ell)^\top \otimes \overline{\mathbf{G}}$. We show a reduction from $(0, \text{Matrix}, \text{Sample})$ -generalized ℓ -decomposed LWE to (ℓ, Sample') -decomposed LWE. Together with Lemma 8.8, this gives the claimed statement. Suppose that the challenger provided us with $(\mathbf{W}, \mathbf{R}, \mathbf{u}^\top)$, where $\mathbf{u}$ is either sampled uniformly at random or

$$\mathbf{u}^\top = \mathbf{s}^\top (\mathbf{W} \cdot (\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \overline{\mathbf{G}}) + \mathbf{e}^\top,$$

where $\mathbf{s} \leftarrow \text{Sample}(1^n)$ and $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}_q^{\ell^2(n+a) \lceil \log q \rceil}, \chi'}$ is a noise term.

The reduction samples $\overline{\mathbf{W}} \leftarrow \mathbb{Z}_q^{a \times \ell k}$ and sets

$$\widehat{\mathbf{W}} := \begin{bmatrix} \overline{\mathbf{W}} \\ \mathbf{W} \end{bmatrix}, \quad \widehat{\mathbf{u}}^\top := \mathbf{u}^\top + \mathbf{r}^\top (\overline{\mathbf{W}} (\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes [\mathbf{I}_a \otimes \mathbf{g}^\top | \mathbf{0}]).$$

Finally, it outputs $(\widehat{\mathbf{W}}, \mathbf{R}, \widehat{\mathbf{u}}^\top)$. Observe that the distribution of $\widehat{\mathbf{W}}$ is uniform over $\mathbb{Z}_q^{(n+a) \times \ell k}$, whereas $\mathbf{R}$ is distributed as $\mathcal{D}_{\mathbb{Z}_q^{k \times \ell(n+a) \lceil \log q \rceil}, \sigma'}$. If $\mathbf{u}$ is random, $\widehat{\mathbf{u}}$ is random too. In the planted case, we have

$$\widehat{\mathbf{u}}^\top = [\mathbf{r}^\top | \mathbf{s}^\top] (\widehat{\mathbf{W}} \cdot (\mathbf{I}_\ell \otimes \mathbf{R}) + \text{vec}(\mathbf{I}_\ell)^\top \otimes \mathbf{I}_{a+n} \otimes \mathbf{g}^\top) + \mathbf{e}^\top.$$

□

Acknowledgements. D.A. is supported by Input Output (iog.io) through their funding of the Edinburgh Blockchain Technology Lab. V.C. is supported by the European Union’s Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101202597 (ACryL). G.A and P.L. are supported by the European Research Council (ERC) under the EU’s Horizon 2020 research and innovation programme (Grant agreement No. 101019547) and by a Cariplo CRYPTONOMEX grant. The majority of this work was done while G.A. was a postdoctoral fellow at Bocconi University. G.M. is supported by the European Research Council through an ERC Starting Grant (Grant agreement No. 101077455, ObfusQation) and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA – 390781972. L. Roy was supported by the Swiss National Science Foundation (SNSF) under the Starting grant “Low-Entropy Cryptography for the Quantum Era”. Part of this work was conducted while D. Abram, G. Arnon, G. Malavolta and L. Roy were visiting the Simons Institute for the Theory of Computing. This work was supported by the fast Internet connection at Rifugio Brentei.

Disclosure of AI Usage The proofs were found by humans. ChatGPT was used solely for writing assistance. The authors take full responsibility for all claims and proofs made in this paper.

9 References

- [AAC⁺26] Damiano Abram, Gal Arnon, Valerio Cini, Paul Lou, Giulio Malavolta, and Lawrence Roy. Decomposed LWE is equivalent to succinct LWE. Cryptology ePrint Archive, Paper 2026/1829, 2026. ⁴⁶

- [ABB10a] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Berlin, Heidelberg, May / June 2010. ²³
- [ABB10b] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 98–115. Springer, Berlin, Heidelberg, August 2010. ²³
- [ABCR26] Shweta Agrawal, Kaartik Bhushan, Geoffroy Couteau, and Mahshid Riahinia. Post-quantum public-key pseudorandom correlation functions for ot. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology – CRYPTO 2026*, pages 527–558, Cham, 2026. Springer Nature Switzerland. ³
- [ACL⁺22] Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri Aravinda Krishnan Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 102–132. Springer, Cham, August 2022. ³
- [AFLN24] Martin R. Albrecht, Giacomo Fenzi, Oleksandra Lapiha, and Ngoc Khanh Nguyen. SLAP: Succinct lattice-based polynomial commitments from standard assumptions. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VII*, volume 14657 of *LNCS*, pages 90–119. Springer, Cham, May 2024. ¹⁶
- [AGY26a] Abtin Afshar, Rishab Goyal, and Saikumar Yadugiri. Distributed monotone-policy encryption with silent setup from lattices. *Cryptology ePrint Archive*, Paper 2026/372, 2026. ³
- [AGY26b] Abtin Afshar, Rishab Goyal, and Saikumar Yadugiri. Silent distributed cryptography for dnfs and threshold policies from lattices. *IACR Cryptol. ePrint Arch.*, 2026:1547, 2026. ^{3, 15}
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. ²³
- [AKSY22] Shweta Agrawal, Elena Kirshanova, Damien Stehlé, and Anshu Yadav. Practical, round-optimal lattice-based blind signatures. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 39–53. ACM Press, November 2022. ³
- [AMR25] Damiano Abram, Giulio Malavolta, and Lawrence Roy. Key-homomorphic computations for RAM: Fully succinct randomised encodings and more. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 236–268. Springer, Cham, August 2025. ^{3, 5, 6, 7, 24, 52, 69}
- [AMR26a] Damiano Abram, Giulio Malavolta, and Lawrence Roy. Correlation intractability for all batched relations. *Cryptology ePrint Archive*, Paper 2026/1140, 2026. ^{3, 5, 25, 51}

- [AMR26b] Damiano Abram, Giulio Malavolta, and Lawrence Roy. How to authenticate a non-deterministic computation. *Cryptology ePrint Archive*, Paper 2026/741, 2026. ⁷²
- [AP09] Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. In Susanne Albers and Jean-Yves Marion, editors, *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009, Freiburg, Germany, February 26-28, 2009*, volume 3 of *LIPIcs*, pages 75–86. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2009. ²³
- [ARS24] Damiano Abram, Lawrence Roy, and Peter Scholl. Succinct homomorphic secret sharing. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 301–330. Springer, Cham, May 2024. ³
- [Bab85] László Babai. On lovász’ lattice reduction and the nearest lattice point problem (shortened version). In *Proceedings of the 2nd Symposium of Theoretical Aspects of Computer Science, STACS ’85*, pages 13–20, Berlin, Heidelberg, 1985. Springer-Verlag. ²⁸
- [BBD⁺26] James Bartusek, Nir Bitansky, Yevgeniy Dodis, Rachit Garg, and David J. Wu. Fair-weather no more: Guaranteed efficiency in secure group messaging. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology – CRYPTO 2026*, pages 200–232, Cham, 2026. Springer Nature Switzerland. ³
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Berlin, Heidelberg, May 2014. ³
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012. ³
- [BLMR13] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Berlin, Heidelberg, August 2013. ²⁶
- [BLNS23] Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Alessandro Sorniotti. A framework for practical anonymous credentials from lattices. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 384–417. Springer, Cham, August 2023. ³
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011. ³

- [CHKP10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 523–552. Springer, Berlin, Heidelberg, May / June 2010. ²³
- [CHW25] Jeffrey Champion, Yao-Ching Hsieh, and David J. Wu. Registered ABE and adaptively-secure broadcast encryption from succinct LWE. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 3–34. Springer, Cham, August 2025. ^{3, 4, 72}
- [CLM23] Valerio Cini, Russell W. F. Lai, and Giulio Malavolta. Lattice-based succinct arguments from vanishing polynomials - (extended abstract). In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 72–105. Springer, Cham, August 2023. ³
- [CLZ22] Yilei Chen, Qipeng Liu, and Mark Zhandry. Quantum algorithms for variants of average-case lattice problems via filtering. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 372–401. Springer, Cham, May / June 2022. ^{19, 64}
- [CW24] Jeffrey Champion and David J. Wu. Distributed broadcast encryption from lattices. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part III*, volume 15366 of *LNCS*, pages 156–189. Springer, Cham, December 2024. ^{3, 4}
- [CW26a] Jeffrey Champion and David J. Wu. Distributed monotone-policy encryption for DNFs from lattices. In Joan Daemen and Emmanuel Thomé, editors, *Advances in Cryptology – EUROCRYPT 2026, Part V*, volume 16545 of *Lecture Notes in Computer Science*, pages 334–365. Springer Nature Switzerland, 2026. ³
- [CW26b] Jeffrey Champion and David J. Wu. Optimal distributed monotone-policy encryption for DNFs and more from lattices. Cryptology ePrint Archive, Paper 2026/1464, 2026. ³
- [CWY26] Jeffrey Champion, David J. Wu, and Shota Yamada. Silent threshold encryption from lattices. Cryptology ePrint Archive, Paper 2026/1716, 2026. ³
- [FMN24] Giacomo Fenzi, Hossein Moghaddas, and Ngoc Khanh Nguyen. Lattice-based polynomial commitments: Towards asymptotic and concrete efficiency. *Journal of Cryptology*, 37(3):31, July 2024. ^{3, 16}
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009. ³
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. ^{21, 23, 27, 28}
- [GR02] Lov Grover and Terry Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions. arXiv:quant-ph/0208112, 2002. ¹⁹

- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Berlin, Heidelberg, August 2013. ³
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 545–554. ACM Press, June 2013. ³
- [GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Predicate encryption for circuits from LWE. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 503–523. Springer, Berlin, Heidelberg, August 2015. ³
- [GY26a] Rishab Goyal and Saikumar Yadugiri. Adaptively-secure flexible and identity-based broadcast encryption from decomposed LWE. *Cryptology ePrint Archive*, Paper 2026/862, 2026. ³
- [GY26b] Rishab Goyal and Saikumar Yadugiri. Distributed monotone policy encryption with stronger security for DNFs and threshold policies from lattices. *Cryptology ePrint Archive*, Paper 2026/1609, 2026. ³
- [GY26c] Rishab Goyal and Saikumar Yadugiri. Equivocal broadcast encryption: Adaptively-secure optimal distributed broadcast encryption from lattices. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology – CRYPTO 2026*, pages 414–445, Cham, 2026. Springer Nature Switzerland. ³
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. ²¹
- [ICMS26] International Centre for Mathematical Sciences. Novel lattice assumptions. Research workshop at the Edinburgh Futures Institute, Edinburgh, July 2026. 13–17 July 2026. ³
- [KB79] Ravindran Kannan and Achim Bachem. Polynomial algorithms for computing the smith and hermite normal forms of an integer matrix. *SIAM Journal on Computing*, 8(4):499–507, 1979. ^{29, 31}
- [KY16] Shuichi Katsumata and Shota Yamada. Partitioning via non-linear polynomial functions: More compact IBEs from ideal lattices and bilinear maps. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 682–712. Springer, Berlin, Heidelberg, December 2016. ^{17, 22}
- [LatZoo26] Lattice Assumption Zoo contributors. Lattice assumption zoo. Available at <https://latticeassumptionzoo.org/>, 2026. Accessed 24 August 2026. ³

- [LLL26] Hanjun Li, Huijia Lin, and George Lu. Succinct garbled circuits with low-depth garbling algorithms. In Joan Daemen and Emmanuel Thomé, editors, *Advances in Cryptology – EUROCRYPT 2026*, pages 246–275, Cham, 2026. Springer Nature Switzerland. ³
- [LZF26] Jiaqi Liu, Yuanyi Zhang, and Fang-Wei Fu. Ciphertext-policy ABE for NC^1 circuits with constant-size ciphertexts from succinct LWE. Cryptology ePrint Archive, Paper 2026/534, 2026. ³
- [Mah18] Urmila Mahadev. Classical verification of quantum computations. In Mikkel Thorup, editor, *59th FOCS*, pages 259–267. IEEE Computer Society Press, October 2018. ³
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. ^{4, 10, 16, 23}
- [MR04] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th FOCS*, pages 372–381. IEEE Computer Society Press, October 2004. ^{8, 10, 27}
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007. ²²
- [MSM25] Shayee Murshid, Ramprasad Sarkar, and Mriganka Mandal. Registered attribute-based encryption with publicly verifiable certified deletion, everlasting security, and more. Cryptology ePrint Archive, Paper 2025/2314, 2025. ³
- [Pei08] Chris Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, 2008. ²⁷
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, Berlin, Heidelberg, March 2006. ²⁷
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6), 2009. Lemma numbering follows the full version dated May 2, 2009. ^{7, 17, 19, 64}
- [vDGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 24–43. Springer, Berlin, Heidelberg, May / June 2010. ³
- [Wee24] Hoeteck Wee. Circuit ABE with $\text{poly}(\text{depth}, \lambda)$ -sized ciphertexts and keys from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 178–209. Springer, Cham, August 2024. ^{3, 4, 8, 25}

- [Wee25] Hoeteck Wee. Almost optimal KP and CP-ABE for circuits from succinct LWE. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 34–62. Springer, Cham, May 2025. [3](#), [4](#), [24](#)
- [Wu25] David J. Wu. Lattice assumptions with hints: Succinct LWE and its applications. Invited talk at the workshop *Obfuscation*, Simons Institute for the Theory of Computing, June 2025. Video recording available online. [4](#), [5](#)
- [WW23a] Hoeteck Wee and David J. Wu. Lattice-based functional commitments: Fast verification and cryptanalysis. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part V*, volume 14442 of *LNCS*, pages 201–235. Springer, Singapore, December 2023. [3](#), [4](#)
- [WW23b] Hoeteck Wee and David J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 385–416. Springer, Cham, April 2023. [3](#)
- [WW25] Hoeteck Wee and David J. Wu. Unbounded distributed broadcast encryption and registered ABE from succinct LWE. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 204–235. Springer, Cham, August 2025. [3](#), [4](#)
- [WWW25] Brent Waters, Hoeteck Wee, and David J. Wu. New techniques for preimage sampling: Improved NIZKs and more from LWE. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part IV*, volume 15604 of *LNCS*, pages 3–33. Springer, Cham, May 2025. [23](#)