Decomposed LWE is Equivalent to Succinct LWE

Damiano Abram [*]
University of Edinburgh

Gal Arnon [†]
Bocconi University

Valerio Cini [‡]
Bocconi University

Paul Lou [§]
Bocconi University

Giulio Malavolta [¶]
Bocconi University

Lawrence Roy [‖]
IBM Research Zürich

Abstract

We prove that the Succinct Learning with Errors assumption, introduced by Wee (CRYPTO '24), and the Decomposed Learning with Errors assumption, introduced by Abram, Malavolta, and Roy (CRYPTO '25), are equivalent under appropriate parameter settings. Abram, Malavolta, and Roy proved that Succinct LWE implies Decomposed LWE. We establish the converse implication, showing that Decomposed LWE implies Succinct LWE.

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 $\mathbf{s}^\mathsf{T}\mathbf{A} + \mathbf{e}^\mathsf{T}$ 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, WW23, 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

[*] abram.damiano@protonmail.com
[†] galarnon42@gmail.com
[‡] cini.valerio@gmail.com
[§] paul96lou@gmail.com
[¶] giulio.malavolta@hotmail.it
[‖] ldr709@gmail.com

them. Examples include the ongoing Lattice Assumption Zoo project [LatZoo26] and the 2026 ICMS workshop *Novel Lattice Assumptions* [ICMS26].

Among other assumptions, the succinct LWE problem [Wee24] and the decomposed LWE problem [AMR25] have gained particular momentum. Leveraging such assumptions, one can prove the security of many new advanced cryptographic primitives, e.g., more efficient attribute-based encryption [Wee24, Wee25, AMR25], broadcast encryption [CHW25, WW25, GY26b, BBD+26], laconic function evaluation [AMR25], threshold encryption [AGY26, GY26a, CWY26], succinct randomized encodings [AMR25, LLL26], registration-based encryption [CHW25, WW25], correlation-intractable hash functions [AMR26], and more. In light of these developments, there is an urge to better understand the hardness of these problems, with the goal of strengthening and unifying the foundations of new advanced cryptographic primitives.

One reason why some believed that the succinct LWE and decomposed LWE assumptions might be separated from a computational complexity standpoint is that the setup in succinct LWE generates a secret trapdoor, whereas the setup in decomposed LWE is entirely public. However, as we will shortly explain through a "hinted-LWE" framework, this distinction turns out to be superficial. Indeed, we show that these two assumptions are in fact equivalent (for an appropriate choice of parameters) with only a negligible loss in distinguishing advantage. That is, we prove the following theorem.

Theorem 1.1 (Informal)**.** *Succinct LWE ≡ Decomposed LWE.*

One direction of the equivalence (if succinct LWE is hard then so is decomposed LWE) was shown in [AMR25]. As such, our main technical contribution is the proof of the converse. We refer the reader to Theorem 1.2 for a formal theorem statement.

Before stating our results more precisely, we recall the definitions of succinct and decomposed LWE and review structural similarities between the two.

1.1 The "Hinted-LWE" Framework

The following "hinted-LWE" framework was introduced by David 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 LWE with respect to $\mathbf{A}$ remains hard even when this additional information is revealed. Succinct LWE [Wee24] and BASIS [WW23] are naturally formulated in this way.

Decomposed LWE [AMR25], however, was originally introduced in a different form. Its definition does not contain an explicit trapdoor; instead, it requires LWE to remain hard for a structured public matrix. Wu observed that this structure implicitly determines a short gadget trapdoor for a related matrix, and hence that decomposed LWE can also be expressed within the hinted-LWE framework. We make this reformulation explicit below.

The general template. Consider a wide matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, and matrices $\mathbf{W}_1, \ldots, \mathbf{W}_\ell \in \mathbb{Z}_q^{n\times k}$, and define

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

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]. Typically, $\mathbf{T}$ is sampled from a discrete Gaussian conditioned on Equation (1).

The general security question is whether LWE with respect to $\mathbf{A}$ remains hard when both $\mathbf{D}_\ell$ and the short trapdoor $\mathbf{T}$ are made public. For the LWE variants considered here, the corresponding security statement has the form

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

where $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$, and $\mathbf{s}$, $\mathbf{e}$, which are referred to as the secret and error term of the LWE sample are sampled from appropriate distributions over $\mathbb{Z}_q^n$ and $\mathbb{Z}^m$ respectively.

Succinct LWE. The ℓ-succinct LWE assumption, introduced in [Wee24], considers the matrices $\mathbf{A}$ and $\mathbf{W}_i$ uniformly sampled, and a short preimage $\mathbf{T} \leftarrow \mathbf{D}_\ell^{-1}\left(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma\right)$.[1] In other words, $\mathbf{T}$ is sampled from a discrete Gaussian of parameter σ, conditioned on Equation (1) holding. The secret vector $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ is uniform random, the error term $\mathbf{e}$ is sampled from a discrete Gaussian error distribution, i.e., $\mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^m,\chi}$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^m$ is uniform. Although $\mathbf{T}$ provides short preimages under $\mathbf{D}_\ell$ of all the columns of $\mathbf{I}_\ell \otimes \mathbf{G}$, it should not help distinguish LWE samples associated with the uniformly sampled matrix $\mathbf{A}$ from uniform.

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 [Wee24, Wee25, 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 $o(\ell) \cdot \mathrm{poly}(n, \log q)$; see, for example, [Wu25].

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

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

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

$$\mathbf{A} \coloneqq \overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \mathrm{vec}(\mathbf{I}_\ell)^\intercal \otimes \mathbf{G} \in \mathbb{Z}_q^{n\times \ell^2 n\lceil \log q\rceil}. \tag{2}$$

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

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

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

where $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^{\ell^2 n\lceil \log q\rceil}$ 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.

We now spell out Wu's observation [Wu25]. The short matrix $\mathbf{R}$ implicitly determines a short trapdoor of the form appearing in the hinted-LWE framework. Let

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

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

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

Consider the matrix

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

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

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

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

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

Indeed, $(\mathbf{A}, \mathbf{W}, \mathbf{T})$ can be computed from $(\overline{\mathbf{W}}, \mathbf{R})$, while $\overline{\mathbf{W}}$ is obtained by reordering the blocks of $\mathbf{W}$ and $\mathbf{R}$ is recovered from the bottom k rows of $-\mathbf{T}$. Thus, the two formulations of the problem are equivalent.

1.2 Our Work

Previously, the hardness of decomposed LWE was known to follow from that of an appropriate instantiation of succinct LWE [AMR25, Theorem 4.2]. Our main result establishes the converse reduction, namely that the hardness of succinct LWE follows from an appropriate instantiation of decomposed LWE. Our work therefore shows that the two assumptions are equivalent. More formally, we prove the following theorem:

Theorem 1.2 (Decomposed LWE implies Succinct LWE)**.** *Let $\lambda \in \mathbb{N}$ be the security parameter and let*

$$n = n(\lambda), \quad k = k(\lambda), \quad \ell = \ell(\lambda) \in \mathbb{N},$$
$$\sigma_{\mathsf{dec}} = \sigma_{\mathsf{dec}}(\lambda), \quad \sigma_{\mathsf{suc}} = \sigma_{\mathsf{suc}}(\lambda), \quad \chi_{\mathsf{dec}} = \chi_{\mathsf{dec}}(\lambda), \quad \chi_{\mathsf{suc}} = \chi_{\mathsf{suc}}(\lambda) \in \mathbb{R}, \quad \textit{and}$$
$$q = q(\lambda) \textit{ odd prime.}$$

[2] Under the dimension and Gaussian-width conditions of Lemma 2.6, $\mathbf{T}$ can be used as auxiliary input to SamplePre to sample a fresh trapdoor whose distribution is statistically close to $\mathbf{D}_\ell^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$ if necessary.

Define $m \coloneqq \ell^2 n\lceil \log q \rceil$ *and suppose that*

$$n = \omega(\log \lambda), \qquad k \geq 2n\lceil \log q \rceil, \qquad \frac{q}{\sigma_{\mathsf{dec}}} = \omega(\sqrt{\log \lambda}),$$
$$\chi_{\mathsf{dec}} = \omega(\sqrt{\log \lambda}), \qquad \sigma_{\mathsf{dec}} \geq 2,$$
$$\chi_{\mathsf{suc}} \geq 3m\chi_{\mathsf{dec}}, \qquad \sigma_{\mathsf{suc}} \geq 2\left(\ell^3 n\lceil \log q \rceil + k\right) \log(\ell n) \cdot \sigma_{\mathsf{dec}} \sqrt{k\ell n\lceil \log q \rceil}.$$

Then every probabilistic polynomial-time distinguisher for ℓ*-succinct LWE with parameters*

$$(n, m, k, \chi_{\mathsf{suc}}, \sigma_{\mathsf{suc}}, q)$$

yields a probabilistic polynomial-time distinguisher for ℓ*-decomposed LWE with parameters*

$$(n, k, \chi_{\mathsf{dec}}, \sigma_{\mathsf{dec}}, q),$$

with only a negligible loss in distinguishing advantage.

Discussion of the parameters. It is important to observe that the 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

$$nk\lceil \log q \rceil = o(\ell) \cdot \mathrm{poly}(n, \log q),$$

then the succinct LWE instance produced by the reduction lies in *the same regime*. Thus, the reduction maps an instantiation of decomposed LWE that is useful for these applications to an equally application-relevant instantiation of succinct LWE.

Indeed, the only parameters that change are the Gaussian parameters:

$$\chi_{\mathsf{suc}} \geq 3m\chi_{\mathsf{dec}},$$
$$\sigma_{\mathsf{suc}} \geq 2\left(\ell^3 n\lceil \log q \rceil + k\right) \log(\ell n) \cdot \sigma_{\mathsf{dec}} \sqrt{k\ell n\lceil \log q \rceil}.$$

Consequently, as the increase in the above parameters is only polynomial, the reduction is compatible with a polynomial modulus q.

2 Preliminaries

Notations. We use boldface lower case for column vectors (e.g. $\mathbf{r}$) and boldface upper case for matrices (e.g., $\mathbf{R}$). For integral vectors and matrices (i.e., those over $\mathbb{Z}$), we use the notation $\|\mathbf{r}\|$, $\|\mathbf{R}\|$ to denote the maximum absolute value over all the entries. We denote the security parameter by λ. We say that a function $f : \mathbb{N} \to \mathbb{R}$ is negligible if $|f(\lambda)| = 2^{-\omega(\log \lambda)}$. We say that a function $f : \mathbb{N} \to \mathbb{R}$ is polynomially bounded if $|f(\lambda)| = 2^{O(\log \lambda)} = \lambda^{O(1)}$. For two distributions $\mathcal{X}$ and $\mathcal{Y}$ over the same countable set Ω, their statistical distance is

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

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

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

For a distribution $\mathcal{X}$ over a finite set Ω, we denote by $\mathsf{Unif}(\Omega)$ the uniform distribution over Ω and define min-entropy as

$$\mathrm{H}_\infty(\mathcal{X}) \coloneqq -\log\left(\max_{x \in \Omega} \Pr_{x' \xleftarrow{\$} \mathcal{X}}[x' = x]\right).$$

Unless stated otherwise, all logarithms are taken to base 2.

2.1 Statistical Distance Facts

Fact 2.1 (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*

$$\Delta\big(K(A), K(B)\big) \le \Delta(A, B).$$

Lemma 2.2 (Leftover Hash Lemma [HILL99])**.** *Let q be prime, $\mathcal{X}$ be a distribution over $\mathbb{Z}^m$ and $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{x} \xleftarrow{\$} \mathcal{X}, \mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^n$. Then*

$$\Delta\big((\mathbf{A}, \mathbf{A}\mathbf{x}), (\mathbf{A}, \mathbf{u})\big) \le \frac{1}{2}\sqrt{q^n \cdot 2^{-\mathrm{H}_\infty(\mathcal{X})}}.$$

Consequently, if $\mathrm{H}_\infty(\mathcal{X}) \ge n \log q + 2\lambda - 2$, then $(\mathbf{A}, \mathbf{A}\mathbf{x}) \approx_s (\mathbf{A}, \mathbf{u})$, with statistical distance at most $2^{-\lambda}$.

2.2 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 $\mathbf{x}$, $\Pr[\mathbf{x}] \propto e^{-\pi \cdot (x_1^2 + \cdots + x_m^2)/\chi^2}$.

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

Lemma 2.3 (Norm of Gaussian, adapted from [GPV08, Lemma 4.2, ePrint version])**.** *For all $\chi \ge 1$ and $k \ge 2$,*

$$\Pr\left[|x| \ge \chi\sqrt{k} \mid x \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}, \chi}\right] \le 2^{-k}.$$

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

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

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

Lemma 2.4 (Noise rerandomization, adapted from [KY16, Lemma 1]). *Let $m, q \in \mathbb{N}$, let $\chi = \omega(\sqrt{\log m})$, and let $\mathbf{V} \in \mathbb{Z}^{m\times m}$. For every $\chi' > 2\chi m\|\mathbf{V}\|$, there exists a probabilistic polynomial-time algorithm* ReRand *with the following property:*

- *For every $\mathbf{a} \in \mathbb{Z}_q^m$ and $\mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^m,\chi}$, given only $\mathbf{a}^\mathsf{T} + \mathbf{e}^\mathsf{T} \bmod q$, the algorithm outputs*

$$\mathsf{ReRand}\left(\mathbf{V}, \mathbf{a}^\mathsf{T} + \mathbf{e}^\mathsf{T}, \chi, \frac{\chi'}{2\chi}\right) = \mathbf{a}^\mathsf{T}\mathbf{V} + (\mathbf{e}')^\mathsf{T} \bmod q,$$

  where the distribution of $\mathbf{e}'$ is statistically close to $\mathcal{D}_{\mathbb{Z}^m,\chi'}$. In particular, if $\mathbf{V}$ is invertible modulo q and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ is uniform, then

$$\mathsf{ReRand}\left(\mathbf{V}, \mathbf{u}^\mathsf{T}, \chi, \frac{\chi'}{2\chi}\right) \approx_s \mathsf{Unif}(\mathbb{Z}_q^m).$$

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 2.5 (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 $k = \mathsf{poly}(\lambda)$, then $H_\infty(\mathcal{R}) \geq k \log_2 \sigma - \mathsf{negl}(\lambda)$.*

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

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

By definition of min-entropy,

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

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

Trapdoor and Preimage Sampling

Our reduction will use gadget trapdoors [MP12], which builds on a long sequence of works on constructing lattice trapdoors [Ajt96, GPV08, AP09, ABB10a, ABB10b, CHKP10]. We use the adaptation of such results from [WWW25].

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

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

For any $t \in \mathbb{N}$, the bit-decomposition map $\mathbf{G}^{-1} : \mathbb{Z}_q^{n \times t} \to \{0,1\}^{(n \cdot \lceil \log q \rceil) \times t}$ 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,

$$\mathbf{x} \xleftarrow{\$} \mathbf{A}^{-1}(\mathbf{v}, \tau) \quad \Longleftrightarrow \quad \mathbf{x} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^m, \tau} \;\Big|\; \mathbf{A} \cdot \mathbf{x} = \mathbf{v} \bmod q. \tag{3}$$

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

Lemma 2.6 (Trapdoor generation and sampling, adapted from [MP12])**.** *Let n, m, q be lattice parameters with $m \geq 3n\lceil \log q \rceil$. Then there exist efficient algorithms* (TrapGen, SamplePre) *with the following syntax:*

- $\mathsf{TrapGen}(1^n, q, m) \to (\mathbf{A}, \mathbf{T})$*: On input the lattice dimension n, the modulus q, and the number of samples m, the trapdoor generation algorithm outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor $\mathbf{T}$.*
- $\mathsf{SamplePre}(\mathbf{A}, \mathbf{T}, \mathbf{v}, s) \to \mathbf{x}$*: On input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a trapdoor $\mathbf{T}$, a target vector $\mathbf{v} \in \mathbb{Z}_q^n$, and a Gaussian width parameter s, the preimage-sampling algorithm outputs a vector $\mathbf{x} \in \mathbb{Z}^m$.*

Moreover, the above algorithms satisfy the following properties:

- ***Trapdoor distribution:*** *If $(\mathbf{A}, \mathbf{T}) \leftarrow \mathsf{TrapGen}(1^n, q, m)$, then the distribution of $\mathbf{A}$ is 2^{-n} close to the uniform distribution over $\mathbb{Z}_q^{n \times m}$. Moreover $\mathbf{A} \cdot \mathbf{T} = \mathbf{G} \bmod q$ and $\|\mathbf{T}\| = 1$.*
- ***Preimage sampling:*** *For all matrices $\mathbf{T}$, width parameters s, and all target vector $\mathbf{v} \in \mathbb{Z}_q^n$ in the column span of $\mathbf{A}$, if we sample $\mathbf{x} \leftarrow \mathsf{SamplePre}(\mathbf{A}, \mathbf{T}, \mathbf{v}, s)$, then $\mathbf{A}\mathbf{x} = \mathbf{v} \bmod q$.*
- ***Preimage distribution:*** *Suppose $\mathbf{T}$ is a gadget trapdoor for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ (i.e., $\mathbf{A}\mathbf{T} = \mathbf{G} \bmod q$). Then, for all $s \geq m\|\mathbf{T}\| \log n$, and all target vectors $\mathbf{v} \in \mathbb{Z}_q^n$, the statistical distance between the following distribution is at most 2^{-n}:*

$$\{\mathbf{x} \leftarrow \mathsf{SamplePre}(\mathbf{A}, \mathbf{T}, \mathbf{v}, s)\} \qquad \textit{and} \qquad \{\mathbf{x} \leftarrow \mathbf{A}^{-1}(\mathbf{v}, s)\}.$$

We extend SamplePre to matrix targets by applying it independently columnwise.

2.3 Computational Assumptions

Definition 2.7 (Decisional ℓ-succinct LWE, [Wee25])**.** *Fix a security parameter λ and parameters n, m, k, χ, σ,and q. The (decision) ℓ-succinct LWE assumption stipulates that*

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

where

$$\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi}, \mathbf{c} \leftarrow \mathbb{Z}_q^m,$$
$$\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n \times k}, \mathbf{T} \leftarrow \left[\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}\right]^{-1}\left(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma\right).$$

Definition 2.8 (Decisional ℓ-Decomposed LWE, [AMR25]). *Fix a security parameter λ and parameters n, k, χ, σ, and q. The (decision) ℓ-decomposed LWE assumption stipulates that*

$$\left(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{s}^\intercal(\overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \mathrm{vec}(\mathbf{I}_\ell)^\intercal \otimes \mathbf{G}) + \mathbf{e}^\intercal\right) \approx_c \left(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{u}^\intercal\right)$$

where $\overline{\mathbf{W}} \xleftarrow{\$} \mathbb{Z}_q^{n\times \ell k}$, $\mathbf{R} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^{k\times \ell n\lceil \log q\rceil},\sigma}$, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^{\ell^2 n\lceil \log q\rceil},\chi}$, *and* $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^{\ell^2 n\lceil \log q\rceil}$.

3 Proof of Equivalence

Before delving into the technical description of the proof of Theorem 1.2, we provide an exposition of it and in particular we discuss how we sample a succinct LWE instance using a given decomposed LWE instance.

3.1 Reduction Overview

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. Throughout this section, set $m = \ell^2 n\lceil \log q\rceil$. In both assumptions, the LWE challenge is associated with a matrix $\mathbf{A}_{\mathsf{x}} \in \mathbb{Z}_q^{n\times m}$, $\mathsf{x} \in \{\mathsf{suc}, \mathsf{dec}\}$, and the auxiliary information contains or implicitly determines a short gadget trapdoor $\mathbf{T}$ for the corresponding matrix

$$\mathbf{D}_\ell = [\mathbf{I}_\ell \otimes \mathbf{A}_{\mathsf{x}} \mid \mathbf{W}],$$

satisfying

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

For decomposed LWE, the matrices $\mathbf{A}_{\mathsf{dec}}$ 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 $\mathbf{A}_{\mathsf{x}}$. In succinct LWE, $\mathbf{A}_{\mathsf{suc}}$ is uniformly random, whereas in decomposed LWE the matrix $\mathbf{A}_{\mathsf{dec}}$ has the structured form from Equation (2).

The reduction must therefore transform $\mathbf{A}_{\mathsf{dec}}$ 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 $\widetilde{\mathbf{A}} = \mathbf{A}_{\mathsf{dec}} \cdot \mathbf{S}$. The structured leftover-hash lemma of Afshar, Goyal, and Yadugiri [AGY26, 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

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

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

$$\widetilde{\mathbf{D}}_\ell \coloneqq \mathbf{D}_\ell \overline{\mathbf{S}} = [\mathbf{I}_\ell \otimes \widetilde{\mathbf{A}} \mid \mathbf{W}],$$

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

The problem is that a random low-norm matrix, even if invertible modulo q, with overwhelming probability will not have a low-norm inverse. Consequently, $\overline{\mathbf{S}}^{-1}\mathbf{T}$ will not be a (low-norm) gadget trapdoor of $\widetilde{\mathbf{D}}_\ell$.

Rerandomization with a short inverse. The failure of the preceding attempt comes from sampling a short matrix $\mathbf{S}$ first and only afterwards asking whether $\mathbf{S}^{-1}$ is also short. We instead restrict $\mathbf{S}$ to a structured family for which invertibility and the shortness of both $\mathbf{S}$ and $\mathbf{S}^{-1}$ 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

$$\mathbf{A} = [\mathbf{A}_0 \mid \mathbf{A}_{\mathrm{rest}}], \qquad \mathbf{A}_0 \coloneqq \mathbf{W}_1 \cdot \mathbf{R}_1 + \mathbf{G} \in \mathbb{Z}_q^{n \times n\lceil \log q \rceil}, \qquad \mathbf{A}_{\mathrm{rest}} \in \mathbb{Z}_q^{n \times ((\ell^2-1)n \cdot \lceil \log q \rceil)}.$$

Under the parameter conditions of the theorem, 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

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

and define

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

Observe that both $\mathbf{S}$ and $\mathbf{S}^{-1}$ are short. Then consider the rerandomized matrix

$$\widetilde{\mathbf{A}} \coloneqq \mathbf{A} \cdot \mathbf{S} = [\, \mathbf{A}_0 \mid \mathbf{A}_0 \cdot \mathbf{X} + \mathbf{A}_{\mathrm{rest}} \,].$$

Although $\mathbf{A}_0$ is marginally close to uniform, it is correlated with $\mathbf{A}_{\mathrm{rest}}$ 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 3.1 proves that

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

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

We next transport the decomposed LWE trapdoor through this transformation. Define

$$\overline{\mathbf{S}} \coloneqq \operatorname{diag}(\mathbf{I}_\ell \otimes \mathbf{S}, \mathbf{I}_k).$$

Then

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

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

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

Moreover, using the explicit form of $\mathbf{T}$, we obtain

$$\left\|\widehat{\mathbf{T}}\right\| \le \max\{1, \|\mathbf{R}\|\}.$$

Thus, $\widehat{\mathbf{T}}$ is a valid short gadget trapdoor for $\widetilde{\mathbf{D}}_\ell$. It is not, however, distributed according to the conditional Gaussian distribution required by succinct LWE. Under the parameter conditions of the theorem, this trapdoor can be used as auxiliary input to the preimage sampler. The reduction computes

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

By the columnwise form of Lemma 2.6,

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

The freshly sampled matrix $\widetilde{\mathbf{T}}$ is the trapdoor given to the succinct-LWE distinguisher.

Finally, the reduction rerandomizes the noisy decomposed LWE challenge using Lemma 2.4 and sets

$$\widetilde{\mathbf{b}}^\mathsf{T} \leftarrow \mathsf{ReRand}\left(\mathbf{S}, \mathbf{b}^\mathsf{T}, \chi_{\text{dec}}, \frac{\chi_{\text{suc}}}{2\chi_{\text{dec}}}\right).$$

Every entry of $\mathbf{S}$ belongs to $\{-1, 0, 1\}$, and hence $\|\mathbf{S}\| = 1$. Therefore,

$$\chi_{\text{suc}} \ge 3m\chi_{\text{dec}} > 2\chi_{\text{dec}} m \|\mathbf{S}\|,$$

so Lemma 2.4 applies. In the planted case, its output is statistically close to $\mathbf{s}^\mathsf{T}\widetilde{\mathbf{A}} + \widetilde{\mathbf{e}}^\mathsf{T}$, where $\widetilde{\mathbf{e}} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^m, \chi_{\text{suc}}}$. 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, the last part of Lemma 2.4 gives $\widetilde{\mathbf{b}}^\mathsf{T} \approx_s \mathsf{Unif}(\mathbb{Z}_q^m)$.

3.2 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 $q \ge 3$, so that reduction modulo q on $\{-1, 0, 1\}$ is injective. It follows that $\mathrm{H}_\infty(\mathcal{X}_{\text{col}}) = n\lceil \log q \rceil \log 3$.

Lemma 3.1 (Public-matrix rerandomization). *Let $\ell = \ell(\lambda), n = n(\lambda), k = k(\lambda), q = q(\lambda) \in \mathbb{N}$, and $\chi = \chi(\lambda), \sigma = \sigma(\lambda) \in \mathbb{R}$. Suppose*

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

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

$$\mathbf{A} = \overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \mathrm{vec}(\mathbf{I}_\ell)^\intercal \otimes \mathbf{G}, \qquad \text{with} \quad \overline{\mathbf{W}} \xleftarrow{\$} \mathbb{Z}_q^{n \times \ell k}, \mathbf{R} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^{k \times \ell n\lceil \log q \rceil}, \sigma}.$$

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

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

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

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 $\mathbb{Z}_q^{n \times k}$,

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

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 $i \geq 2$.

$\mathsf{Expt}_{\mathsf{real}}$:

1. Sample

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

2. Let

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

3. Output

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

Expt_1:

1. Sample

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

2. Let

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

3. Output

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

Expt_2:

1. Sample

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

2. Let

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

3. Output

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

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

1. Sample

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

2. Let

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

3. Output

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

We now bound the statistical distance between $\mathsf{Expt}_{\text{real}}$ and $\mathsf{Expt}_{\text{rand}}$. We apply Lemma 2.2 in two column-wise 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 $\mathbf{X}$.

By the assumptions $k \geq 2n \log q$ and $\sigma \geq 2$, together with Fact 2.5, we have

$$\mathsf{H}_\infty(\mathcal{R}) - n\log q \geq k\log\sigma - n\log q - \mathrm{negl}(\lambda) \geq n\log q - \mathrm{negl}(\lambda).$$

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

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

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

Thus, in both applications, the source min-entropy exceeds the output length $n\log q$ by $\Omega(n\lceil \log q\rceil)$. By Lemma 2.2, 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 column-wise hybrids is negligible in λ.

1. $\mathsf{Expt}_{\text{real}} \equiv \mathsf{Expt}_1$: 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 $\overline{\mathbf{W}}^* = \begin{bmatrix}\mathbf{W}_1^* & \cdots & \mathbf{W}_\ell^*\end{bmatrix} \in \mathbb{Z}_q^{n\times \ell k}$ and $\mathbf{A}_0^* \in \mathbb{Z}_q^{n\times n\lceil \log q\rceil}$. Conditioned on the event that in $\mathsf{Expt}_{\text{real}}$ 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 $\mathsf{Expt}_{\text{real}}$ 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.,

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

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 $\mathbf{R}_2, \ldots, \mathbf{R}_\ell$ 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,

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

2. $\Delta(\mathsf{Expt}_1, \mathsf{Expt}_2) = \mathsf{negl}(\lambda)$. Recall that both experiments output $\perp$ when $\mathbf{W}_1$ does not have full row rank. Write the columns of $\mathbf{R}_1$ as $\mathbf{r}_{1,1}, \ldots, \mathbf{r}_{1,n\lceil \log q \rceil}$. Each column, reduced modulo q, is an independent sample from $\mathcal{R}$. For each column, Lemma 2.2, 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

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

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

Define the following randomized post-processing 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

$$\mathbf{W}_2, \ldots, \mathbf{W}_\ell, \qquad \mathbf{R}_2, \ldots, \mathbf{R}_\ell, \qquad \mathbf{X}$$

according to their respective distributions, and sample

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

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

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

Finally, output

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

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 post-processing gives Expt_1; for the second, it gives Expt_2. Therefore, by Fact 2.1, and the assumption that $n\lceil \log q \rceil = \omega(\log \lambda)$, one obtains

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

3. $\Delta(\mathsf{Expt}_2, \mathsf{Expt}_{\mathsf{rand}}) = \mathsf{negl}(\lambda)$: A column-wise application of Lemma 2.2 to $\mathbf{X}$ gives

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

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

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

- Sample $\mathbf{W}_1, \ldots, \mathbf{W}_\ell$ and $\mathbf{R}_2, \ldots, \mathbf{R}_\ell$ according to their prescribed distributions.

- If $\mathbf{W}_1$ does not have full row rank, output $\perp$.
- Otherwise, sample
$$\widehat{\mathbf{R}}_1 \leftarrow \mathbf{W}_1^{-1}(\mathbf{U}_0 - \mathbf{G}, \sigma),$$
set $\widehat{\mathbf{R}}_j = \mathbf{R}_j$ for $j \in \{2, \ldots, \ell\}$, construct $\widehat{\mathbf{A}}_{\mathsf{rest}}$ as in Expt_2, and output
$$\left(\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{Y} + \widehat{\mathbf{A}}_{\mathsf{rest}}\right).$$

Applied to $(\mathbf{U}_0, \mathbf{U}_0\mathbf{X})$, this post-processing gives Expt_2. Applied to $(\mathbf{U}_0, \mathbf{U}_{\mathsf{rest}})$, it gives
$$\left(\overline{\mathbf{W}}, \mathbf{U}_0, \mathbf{U}_{\mathsf{rest}} + \widehat{\mathbf{A}}_{\mathsf{rest}}\right).$$
Since $\mathbf{U}_{\mathsf{rest}}$ is independently uniform, its translation by $\widehat{\mathbf{A}}_{\mathsf{rest}}$ is also uniform and independent of the first two coordinates. Thus, the latter tuple is distributed exactly as $\mathsf{Expt}_{\mathsf{rand}}$. Therefore, as before, by Fact 2.1 and the assumption that $n\lceil \log q \rceil = \omega(\log \lambda)$, one gets
$$\Delta\left(\mathsf{Expt}_2, \mathsf{Expt}_{\mathsf{rand}}\right) \le \Delta\left((\mathbf{U}_0, \mathbf{U}_0\mathbf{X}), (\mathbf{U}_0, \mathbf{U}_{\mathsf{rest}})\right) \le (\ell^2 - 1)n\lceil \log q \rceil 2^{-\Omega(n\lceil \log q \rceil)} = \mathsf{negl}(\lambda).$$

By the triangle inequality, we get
$$\Delta\left(\mathsf{Expt}_{\mathsf{real}}, \mathsf{Expt}_{\mathsf{rand}}\right) \le \mathsf{negl}(\lambda).$$

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

3.3 Proof of Main Theorem

We now give a proof of Theorem 1.2 by combining Lemmas 2.4, 2.6 and 3.1. We first describe the reduction and establish three statistical properties of its output.

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

Write the decomposed LWE input as $(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{b}^\intercal)$, where
$$\overline{\mathbf{W}} = [\mathbf{W}_1 \mid \cdots \mid \mathbf{W}_\ell] \in \mathbb{Z}_q^{n \times \ell k}, \qquad \mathbf{R} = [\mathbf{R}_1 \mid \cdots \mid \mathbf{R}_\ell] \in \mathbb{Z}^{k \times \ell n \lceil \log q \rceil}.$$

The vector $\mathbf{b}^\intercal$ is either uniform over $\mathbb{Z}_q^m$ or is of the form $\mathbf{s}^\intercal\mathbf{A}_{\mathsf{dec}} + \mathbf{e}^\intercal \bmod q$, where $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ and $\mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^m, \chi_{\mathsf{dec}}}$.

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

1. If $\|\mathbf{R}\| > 2\sigma_{\mathsf{dec}}\sqrt{k\ell n\lceil \log q \rceil}$, return an independent uniform bit and terminate.
2. Compute
$$\mathbf{A}_{\mathsf{dec}} \coloneqq \overline{\mathbf{W}}(\mathbf{I}_\ell \otimes \mathbf{R}) + \mathrm{vec}(\mathbf{I}_\ell)^\intercal \otimes \mathbf{G} \in \mathbb{Z}_q^{n \times m}$$
and write $\mathbf{A}_{\mathsf{dec}} = [\mathbf{A}_0 \mid \mathbf{A}_{\mathsf{rest}}]$, where $\mathbf{A}_0 = \mathbf{W}_1\mathbf{R}_1 + \mathbf{G}$.

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

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

Define

$$\mathbf{A}^\star \coloneqq \mathbf{A}_{\mathrm{dec}}\mathbf{S}, \qquad \mathbf{W}^\star \coloneqq \begin{bmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_\ell \end{bmatrix}.$$

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

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

Let

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

Sample

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

5. Compute

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

6. Give $(\mathbf{A}^\star, \mathbf{c}^{\star\mathsf{T}}, \mathbf{W}^\star, \mathbf{T}^\star)$ 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 2.3 to every entry of $\mathbf{R}$ with tail parameter $4k\ell n\lceil \log q\rceil$ and taking a union bound gives

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

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 $\mathbf{A}^\star\mathbf{Q} = \mathbf{A}_{\mathrm{dec}}$, 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 $\widetilde{\mathbf{D}}_\ell\widehat{\mathbf{T}}$ 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

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

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

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

$$\left\|\widehat{\mathbf{T}}\right\| \le \max\{1, \|\mathbf{R}\|\} \le 2\sigma_{\mathsf{dec}}\sqrt{k\ell n\lceil \log q\rceil}.$$

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

$$(\ell m + k)\left\|\widehat{\mathbf{T}}\right\| \log(\ell n) \le 2\left(\ell^3 n\lceil \log q\rceil + k\right)\sigma_{\mathsf{dec}}\sqrt{k\ell n\lceil \log q\rceil}\,\log(\ell n) \le \sigma_{\mathsf{suc}}. \tag{4}$$

Moreover,

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

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

Thus, the call to SamplePre in Step 4 satisfies all the hypotheses of Lemma 2.6, 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 $\mathsf{x} \in \{\mathsf{pl}, \mathsf{unif}\}$, let $\mathsf{Hyb}_{0,\mathsf{x}}$ denote the distribution of the view

   $$\left(\mathbf{A}^\star, \mathbf{c}^{\star\mathsf{T}}, \mathbf{W}^\star, \mathbf{T}^\star\right)$$

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

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

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

   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

   $$\mathsf{Hyb}_{1,\mathsf{x}} \coloneqq \left(\mathbf{A}^\star, \mathbf{c}^{\star\mathsf{T}}, \mathbf{W}^\star, \mathbf{T}\right).$$

3. For $\mathsf{x} \in \{\mathsf{pl}, \mathsf{unif}\}$, we define $\mathsf{Hyb}_{2,\mathsf{x}}$ by modifying the distribution of the challenge vector $\mathbf{c}$ that lies in the second coordinate of $\mathsf{Hyb}_{1,\mathsf{x}}$. 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 $\mathsf{Hyb}_{2,\mathsf{x}}$, we still have

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

but now we freshly sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\widetilde{\mathbf{e}} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^m,\chi_{\mathsf{suc}}}$, and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ independently of the $\mathbf{A}^\star$ and $\mathbf{W}^\star$and modify the challenge vector as follows

$$\begin{aligned}\mathsf{Hyb}_{2,\mathsf{pl}} &\coloneqq \left(\mathbf{A}^\star, \mathbf{s}^\intercal\mathbf{A}^\star + \widetilde{\mathbf{e}}^\intercal, \mathbf{W}^\star, \mathbf{T}\right),\\ \mathsf{Hyb}_{2,\mathsf{unif}} &\coloneqq \left(\mathbf{A}^\star, \mathbf{u}^\intercal, \mathbf{W}^\star, \mathbf{T}\right).\end{aligned}$$

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

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

Define

$$\begin{aligned}\mathsf{Hyb}_{3,\mathsf{pl}} &\coloneqq (\mathbf{U}_\mathsf{A}, \mathbf{s}^\intercal\mathbf{U}_\mathsf{A} + \widetilde{\mathbf{e}}^\intercal, \mathbf{U}_\mathsf{W}, \mathbf{T}),\\ \mathsf{Hyb}_{3,\mathsf{unif}} &\coloneqq (\mathbf{U}_\mathsf{A}, \mathbf{u}^\intercal, \mathbf{U}_\mathsf{W}, \mathbf{T}).\end{aligned}$$

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

Lemma 3.2 (Trapdoor hybrid)**.** *For every* $\mathsf{x} \in \{\mathsf{pl}, \mathsf{unif}\}$*, we have* $\mathsf{Hyb}_{0,\mathsf{x}} \approx_s \mathsf{Hyb}_{1,\mathsf{x}}$.

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

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

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

Since Equation (4) holds, the preimage-distribution guarantee of Lemma 2.6, followed by a hybrid over the $\ell n\lceil\log q\rceil$ target columns, gives

$$\Delta\left((\mathsf{Hyb}_{0,\mathsf{x}} \mid \Omega_\mathsf{x} = \omega, \mathcal{E}), (\mathsf{Hyb}_{1,\mathsf{x}} \mid \Omega_\mathsf{x} = \omega, \mathcal{E})\right) \le \ell n\lceil\log q\rceil\, 2^{-\ell n}.$$

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

$$\Delta(\mathsf{Hyb}_{0,\mathsf{x}}, \mathsf{Hyb}_{1,\mathsf{x}}) \le \mathbb{E}_{\omega \xleftarrow{\$} \Omega_\mathsf{x}\mid\mathcal{E}}\left[\Delta\left((\mathsf{Hyb}_{0,\mathsf{x}} \mid \omega, \mathcal{E}), (\mathsf{Hyb}_{1,\mathsf{x}} \mid \omega, \mathcal{E})\right)\right] \le \ell n\lceil\log q\rceil\, 2^{-\ell n} = \mathsf{negl}(\lambda).$$

□

Lemma 3.3 (Challenge hybrid)**.** *For every* $\mathsf{x} \in \{\mathsf{pl}, \mathsf{unif}\}$*, we have* $\mathsf{Hyb}_{1,\mathsf{x}} \approx_s \mathsf{Hyb}_{2,\mathsf{x}}$.

Proof. Since $\|\mathbf{S}\| = 1$ and $\chi_{\mathsf{suc}} \geq 3m\chi_{\mathsf{dec}}$, we have $\chi_{\mathsf{suc}} > 2\chi_{\mathsf{dec}} m\|\mathbf{S}\|$. Moreover, $m = \mathrm{poly}(\lambda)$ and $\chi_{\mathsf{dec}} = \omega(\sqrt{\log \lambda})$ imply $\chi_{\mathsf{dec}} = \omega(\sqrt{\log m})$. Therefore, Lemma 2.4 applies with $\mathbf{V} = \mathbf{S}$.

Let

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

In the planted branch, we may use the same $\mathbf{s}$ in both hybrids, since it is uniform and independent of $(\overline{\mathbf{W}}, \mathbf{R}, \mathbf{X}, \mathbf{T})$. Conditioned on $Z_{\mathsf{pl}} = z$ and $\mathcal{E}$, the matrices $\mathbf{S}$, $\mathbf{A}_{\mathsf{dec}}$, and $\mathbf{A}^\star$ are fixed, whereas the original error and the coins of ReRand remain independent. Hence Lemma 2.4 gives

$$\mathbf{c}^{\star\mathsf{T}} \approx_s \mathbf{s}^\mathsf{T}\mathbf{A}_{\mathsf{dec}}\mathbf{S} + \widetilde{\mathbf{e}}^\mathsf{T} = \mathbf{s}^\mathsf{T}\mathbf{A}^\star + \widetilde{\mathbf{e}}^\mathsf{T} \bmod q.$$

In the uniform branch, conditioned on $Z_{\mathsf{unif}} = z$ 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 2.4 therefore gives $\mathbf{c}^{\star\mathsf{T}} \approx_s \mathbf{u}^\mathsf{T}$.

Let $\varepsilon_{\mathsf{RR}}(\lambda)$ denote the negligible statistical error in Lemma 2.4. In either branch, for every z in the support of $Z_{\mathsf{x}} \mid \mathcal{E}$, the preceding pointwise comparison gives

$$\Delta\left(\mathsf{Hyb}_{1,\mathsf{x}} \mid Z_{\mathsf{x}} = z, \mathcal{E}, \mathsf{Hyb}_{2,\mathsf{x}} \mid Z_{\mathsf{x}} = z, \mathcal{E}\right) \leq \varepsilon_{\mathsf{RR}}(\lambda).$$

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

Lemma 3.4 (Public-matrix hybrid)**.** *For every* $\mathsf{x} \in \{\mathsf{pl}, \mathsf{unif}\}$*, we have* $\mathsf{Hyb}_{2,\mathsf{x}} \approx_s \mathsf{Hyb}_{3,\mathsf{x}}$.

Proof. For a pair of matrices $(\mathbf{A}, \mathbf{W})$, define $\mathbf{D}(\mathbf{A}, \mathbf{W}) \coloneqq [\mathbf{I}_\ell \otimes \mathbf{A} \mid \mathbf{W}]$. By Lemma 3.1 and since $\overline{\mathbf{W}}$ is the deterministic block-transpose of $\mathbf{W}$, we have $(\mathbf{A}^\star, \mathbf{W}^\star) \approx_s (\mathbf{U}_\mathsf{A}, \mathbf{U}_\mathsf{W})$. Moreover,

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

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

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 $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\widetilde{\mathbf{e}} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}^m, \chi_{\mathsf{suc}}}$, and

$$\mathbf{T} \xleftarrow{\$} \mathbf{D}(\mathbf{A}, \mathbf{W})^{-1}\left(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma_{\mathsf{suc}}\right),$$

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

Define K_{unif} analogously, except that it samples $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and outputs $K_{\mathsf{unif}}(\mathbf{A}, \mathbf{W}) = (\mathbf{A}, \mathbf{u}^\mathsf{T}, \mathbf{W}, \mathbf{T})$. Let $K_{\mathsf{x}}((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E})$ 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 $\mathbf{D}(\mathbf{A}^\star, \mathbf{W}^\star)\widehat{\mathbf{T}} = \mathbf{I}_\ell \otimes \mathbf{G} \bmod q$. Since $\mathbf{I}_\ell \otimes \mathbf{G}$ has full row rank, $\mathbf{D}(\mathbf{A}^\star, \mathbf{W}^\star)$ has full row rank. Hence $K_{\mathsf{x}}((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}) \equiv \mathsf{Hyb}_{2,\mathsf{x}}$.

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

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

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

By Fact 2.1,

$$K_{\mathsf{x}}\left((\mathbf{A}^\star, \mathbf{W}^\star) \mid \mathcal{E}\right) \approx_s K_{\mathsf{x}}(\mathbf{U}_{\mathsf{A}}, \mathbf{U}_{\mathsf{W}}).$$

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

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

□

Proof of Theorem 1.2. Let $\mathcal{A}$ and $\mathcal{B}^{\mathcal{A}}$ be as above. Combining Lemmas 3.2 to 3.4 and using the triangle inequality gives

$$\mathsf{Hyb}_{0,\mathsf{pl}} \approx_s \mathsf{Hyb}_{3,\mathsf{pl}}, \qquad \mathsf{Hyb}_{0,\mathsf{unif}} \approx_s \mathsf{Hyb}_{3,\mathsf{unif}}.$$

The final experiments are, up to the negligible failure probability above, the distributions of planted and uniform 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 3.5 (SIS analogue)**.** Decomposed SIS was introduced in [AMR26] as the natural search analogue of decomposed LWE.

Our proof of Theorem 1.2 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 $\mathbf{z}' \coloneqq \mathbf{S}\mathbf{z}$. Since $\mathbf{A}^\star = \mathbf{A}_{\mathsf{dec}}\mathbf{S}$, we have

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

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

$$\|\mathbf{z}'\| \leq m\|\mathbf{S}\|\|\mathbf{z}\| = m\|\mathbf{z}\|.$$

Thus, the reduction incurs only a factor-m loss in the SIS norm bound.

Although it is not formally shown in [AMR26] that succinct SIS is reducible to decomposed SIS, this fact can be proven in an identical manner to that of their reduction from succinct LWE to decomposed LWE. Our reduction, therefore, establishes their equivalence.

Acknowledgements

The proof was found by humans. ChatGPT 5.6 Sol was used solely for writing assistance. The authors take full responsibility for all claims and proofs made in this paper.

4 References

[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. [7]

[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. [7]

[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. [1]

[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. [1]

[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. [10]

[AGY26] Abtin Afshar, Rishab Goyal, and Saikumar Yadugiri. Silent distributed cryptography for dnfs and threshold policies from lattices. *IACR Cryptol. ePrint Arch.*, 2026:1547, 2026. [2, 9]

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. [7]

[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. [1]

[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. [1, 2, 3, 4, 9]

[AMR26] Damiano Abram, Giulio Malavolta, and Lawrence Roy. Correlation intractability for all batched relations. Cryptology ePrint Archive, Paper 2026/1140, 2026. [2, 21]

[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. [7]

[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. [1]

[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. [2]

[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. [1]

[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. [1]

[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. [1]

[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. [1]

[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. [7]

[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. [2]

[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. [1]

[CWY26] Jeffrey Champion, David J. Wu, and Shota Yamada. Silent threshold encryption from lattices. Cryptology ePrint Archive, Paper 2026/1716, 2026. [2]

[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. [1, 10]

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009. [1]

[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. [6, 7]

[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. [1]

[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. [1]

[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. [1]

[GY26a] 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. [2]

[GY26b] 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. [2]

[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. [6]

[ICMS26] International Centre for Mathematical Sciences. Novel lattice assumptions. Research workshop at the Edinburgh Futures Institute, Edinburgh, July 2026. 13–17 July 2026. [2]

[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. [7]

[LatZoo26] Lattice Assumption Zoo contributors. Lattice assumption zoo. Available at https://latticeassumptionzoo.org/, 2026. Accessed 24 August 2026. [2]

[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. [2]

[Mah18] Urmila Mahadev. Classical verification of quantum computations. In Mikkel Thorup, editor, *59th FOCS*, pages 259–267. IEEE Computer Society Press, October 2018. [1]

[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. [3, 7, 8, 10]

[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. [7]

[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. [1]

[Wee24] Hoeteck Wee. Circuit ABE with poly(depth, λ)-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. [1, 2, 3]

[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. [2, 3, 8]

[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. [2, 3, 4]

[WW23] 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. [1, 2]

[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. [2, 3]

[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. [7]