

A Subexponential Algorithm for Binary LWE with Sublinear Samples

Zhili Chen¹, Kaijie Jiang², Anyu Wang², and Chuanqi Zhang³

¹ Centre for Quantum Technologies, National University of Singapore, Singapore

² Institute for Advanced Study, Tsinghua University, Beijing, China

³ Department of Software Systems & Cybersecurity, Faculty of Information
Technology, Monash University, Melbourne, Australia

Abstract. We study Binary LWE with polynomial modulus in the regime of a linear number of samples. Prior work of Kirchner and Fouque and of Herold, Kirshanova, and May (CRYPTO '15, DCC '18) gives an algorithm running in time $\exp(O(n/\log \log n))$ using $O(n \log n)$ samples, leaving open whether the same complexity can be achieved with only $m = \Theta(n)$ samples.

We resolve this problem by showing that the same complexity is achievable with even a sublinear number of samples. Fix any constant $a > 1/2$, let $q = n^{a+o(1)}$ be an odd prime, and let $m = \lceil n/\log \log n \rceil$. We give a randomized algorithm that distinguishes $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{e}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ from a uniformly random pair, for every fixed secret-error pair satisfying $\|\mathbf{s}\|_2^2 + \|\mathbf{e}\|_2^2 = O(m + n)$. Its expected running time and space are $\exp(O(n/\log \log n))$.

A crucial component is a new instantiation of the auxiliary superlattices in the Gaussian–Wagner framework of Ducas, Engelberts, and Loyer (CRYPTO '25). These superlattices have sufficiently small smoothing parameters even when the approximation parameter for the entire output list is subexponentially small, while still admitting exact shifted discrete Gaussian sampling within the claimed running time. A Euclidean dual-minimum argument simultaneously controls the smoothing parameters of all projected relevant lattices. Combining the resulting short vectors with an Aharonov–Regev Fourier distinguisher yields the claimed sublinear sample LWE algorithm.

1 Introduction

The Learning with Errors problem (LWE) is one of the fundamental average-case problems underlying lattice based cryptography [32, 9]. In its decisional form, one is given a pair $(\mathbf{A}, \mathbf{b})$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ is uniform, and the goal is to distinguish $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e} \bmod q \in \mathbb{Z}_q^m$ from a uniformly random vector independent of $\mathbf{A}$. Here, $\mathbf{s} \in \mathbb{Z}^n$ is the secret, $\mathbf{e} \in \mathbb{Z}^m$ is the error, and m is the number of available samples. LWE has led to constructions of public key encryption [32], identity based encryption [17], oblivious transfer [30], attribute based encryption for circuits [19], predicate encryption for circuits [20], noninteractive zero knowledge for NP [29], succinct noninteractive arguments for P [13], and fully

homomorphic encryption (FHE) [10]. We refer to [33,28] for surveys of LWE and its cryptographic applications.

In concrete constructions, one is often interested not only in standard LWE, but also in variants in which the secret or the error has particularly small coordinates. Small secrets and errors appear in FrodoKEM, whose underlying public key encryption scheme uses unstructured LWE with secret and error coefficients sampled from small, bounded distributions [18]. Binary secret LWE, which restricts the secret but not necessarily the error to a binary alphabet, also underlies the torus LWE components of TFHE [12].

These parameter choices have motivated both hardness reductions and specialized cryptanalysis. The hardness of binary secret LWE has been established through reductions from standard LWE [9,24]. For small errors, Micciancio and Peikert [25] proved hardness even for binary errors when the number of samples is suitably restricted, including linear sample regimes $m = \Theta(n)$. Subsequent work extends this study to biased binary errors [34]. On the algorithmic side, small secrets and errors have been exploited by algebraic methods [5,34], lattice decoding and hybrid attacks [6,11], dual lattice attacks [4], and BKW-type algorithms [22]. In particular, the complexity of binary-error LWE depends strongly on the number of available samples [25,34], making sample complexity an essential part of the problem.

This work focuses on the regime in which both the secret and the error are narrow. We use *Binary LWE* to denote the case $\mathbf{s} \in \{0,1\}^n$ and $\mathbf{e} \in \{0,1\}^m$, distinguishing it from variants that restrict only one of the two vectors. More generally, we consider *narrow LWE*, where the secret and error satisfy $\|\mathbf{s}\|_2^2 + \|\mathbf{e}\|_2^2 = O(n + m)$. At CRYPTO 2015, Kirchner and Fouque [22] claimed an $\exp(O(n/\ln \ln n))$ time algorithm for narrow LWE with polynomial modulus q , including Binary LWE with only $m = \Theta(n)$ samples. Their approach refines the Blum–Kalai–Wasserman (BKW) algorithm [7], which solves Learning Parity with Noise (LPN) in subexponential time but also requires subexponentially many samples. To reduce this sample requirement, they employed a sample-amplification technique originating with Lyubashevsky [23], generating many new samples by taking short combinations of the original equations. However, Herold, Kirshanova, and May [21, Section 7.3] subsequently identified a gap in the statistical analysis of this amplification step: the argument did not establish that the generated samples were sufficiently close to the distribution required by the subsequent BKW analysis. Their corrected analysis recovers the slightly subexponential bound when $m = O(n \ln n)$, but leaves the linear sample claim unresolved [21,15]. Reducing the sample requirement strengthens the result, since an algorithm using fewer samples can always discard any additional ones.

At CRYPTO 2025, Ducas, Engelberts, and Loyer [15] revisited the Kirchner–Fouque approach from a modular perspective. Their starting point is the folklore interpretation of BKW as a combination of Wagner’s algorithm [35], run over the corresponding dual lattice⁴, and the Aharonov–Regev Fourier distinguisher [2].

⁴ The dual lattice of $\mathcal{L}$ is $\mathcal{L}^* := \{\mathbf{y} \in \text{span}_{\mathbb{R}}(\mathcal{L}) : \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}, \forall \mathbf{x} \in \mathcal{L}\}$.

This suggests studying the Wagner component separately: can it provably generate sufficiently short lattice vectors in subexponential time?

The connection comes from the lattice interpretation of LWE. An integer representative of $\mathbf{b}$ differs by $\mathbf{e}$ from a point in the q -ary lattice $\mathcal{L}_q(\mathbf{A}) = \mathbf{A}^\top \mathbb{Z}^n + q\mathbb{Z}^m$. When the error lies within the unique decoding radius, recovering this lattice point is an instance of bounded distance decoding (BDD) [33,28]. The dual satisfies $q\mathcal{L}_q(\mathbf{A})^* = \mathcal{L}_q^\perp(\mathbf{A})$, where $\mathcal{L}_q^\perp(\mathbf{A}) := \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = \mathbf{0} \bmod q\}$. Given a uniform matrix $\mathbf{A}$, finding a nonzero vector $\mathbf{x} \in \mathcal{L}_q^\perp(\mathbf{A})$ with $\|\mathbf{x}\|_2 \leq \beta$ is the Short Integer Solution (SIS) problem, another foundational problem in lattice based cryptography [3,26]. Its infinity norm variant, denoted SIS^∞ , instead requires $\|\mathbf{x}\|_\infty \leq \beta$. For an LWE input, such a relation satisfies $\langle \mathbf{b}, \mathbf{x} \rangle = \langle \mathbf{e}, \mathbf{x} \rangle \bmod q$, eliminating the secret and providing the basis for a dual distinguishing attack.

Ducas, Engelberts, and Loyer answered the above question by turning Wagner’s algorithm into an approximate discrete Gaussian sampler for random SIS lattices $\mathcal{L}_q^\perp(\mathbf{A})$. For suitable prime polynomial moduli q , their sampler reaches Gaussian width $q/\text{poly} \log(n)$ in time $\exp(O(n/\ln \ln n))$ with only $m = n + \omega(n/\ln \ln n)$ SIS variables. This yields a provable algorithm for SIS^∞ with norm bound $\beta = q/\text{poly} \log(n)$. In fact, their result is stronger than finding a single short vector: the sampler produces a list of lattice vectors whose distributions and dependencies are controlled through conditional similarity. Their Gaussian–Wagner framework achieves this through discrete Gaussian lifting and combining, without relying on sample amplification.

It is therefore natural to ask whether combining this sampler with the Aharonov–Regev distinguisher settles the linear sample LWE question. However, as explained in [15, Section 4.3], the precision available in their subexponential time sampler is insufficient for a direct application. Here, ε is the accuracy parameter in the smoothing conditions used for Gaussian lifting and combining. More precisely, [15, Theorem 3.3] guarantees that the output list is conditionally $q^4\varepsilon$ -similar⁵ to independent discrete Gaussian samples on the SIS lattice. Their subexponential time SIS applications use sufficiently small inverse-polynomial ε , yielding inverse-polynomial conditional error. For LWE, however, the distinguisher uses subexponentially many vectors simultaneously, and a straightforward statistical distance comparison of the entire list with independent Gaussian samples accumulates this error across the outputs. This requires ε to be subexponentially small, of the form $\varepsilon = 2^{-\tilde{\Omega}(n)}$ ⁶, at which point their sampler has exponential complexity. Although this obstacle resembles the sample-amplification issue, it is technically different and cannot be removed by increasing the number of input samples to $\Theta(n \ln n)$. Thus, obtaining the required distributional accuracy without losing the subexponential running time remained a barrier to resolving the linear sample question.

These developments lead us to ask:

⁵ see Definition 2.2 for conditional similarity.

⁶ Here $\tilde{\Omega}$ hides polylogarithmic factors: $\varepsilon \leq 2^{-c_0 n / (\ln n)^{c_1}}$ for some constants $c_0, c_1 > 0$ and all sufficiently large n .

Can Binary LWE with polynomial modulus be solved by a slightly subexponential algorithm using only $m = O(n)$, or even $m = o(n)$, samples?

1.1 Our Contributions

We answer the above question affirmatively. For every fixed $a > 1/2$ and odd prime $q = n^{a+o(1)}$, we obtain an algorithm with expected time and space $\exp(O(n/\ln \ln n))$ using only $m = \lceil n/\ln \ln n \rceil$ samples. Thus, the slightly subexponential bound is achievable not only with a linear number of samples, but even with a sublinear number. Our result holds more generally for every fixed secret and error pair satisfying a bounded norm.

Definition 1.1 (Pointwise Narrow LWE). *Let $n, m \geq 1$ and $q \geq 2$ be integers, and let $B \geq 1$. Set $\mathcal{S}_B := \{(\mathbf{s}, \mathbf{e}) \in \mathbb{Z}^n \times \mathbb{Z}^m : \|\mathbf{s}\|_2^2 + \|\mathbf{e}\|_2^2 \leq B^2(n+m)\}$. The problem $\text{narrow-LWE}_{n,m,q,B}$ asks to distinguish $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{e} \bmod q)$ from $(\mathbf{A}, \mathbf{b})$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{b} \leftarrow \mathbb{Z}_q^m$ are independent and uniform, for every fixed $(\mathbf{s}, \mathbf{e}) \in \mathcal{S}_B$.*

The following theorem gives a more general tradeoff between the number of samples and the running time. It restates Theorem 4.8, proved in Section 4.

Theorem 1.2. *Fix constants $B \geq 1$, $c_+ > 0$, and $A > 0$. Let $m = m(n)$ and $q = q(n)$ satisfy $1 \leq m \leq c_+ n$, $q \leq n^A$, and $q^{n/(m+n)} > \sqrt{m+n}$, where q is an odd prime. Assume that $\tau := \frac{m}{m+n} \ln q \rightarrow \infty$. For all sufficiently large n , there is a randomized algorithm solving $\text{narrow-LWE}_{n,m,q,B}$ with completeness and soundness errors $O(n^{-10})$. Its expected running time and space are at most $\exp((8 \ln 2 + o(1)) \frac{n}{\ln \tau})$.*

The guarantee is pointwise in $(\mathbf{s}, \mathbf{e})$, so averaging over $\mathbf{s} \leftarrow \{0, 1\}^n$ and $\mathbf{e} \leftarrow \{0, 1\}^m$ gives the specialization for Binary LWE. For $q = n^{a+o(1)}$ with fixed $a > 1/2$, the theorem gives $\exp(o(n))$ expected time whenever $m = \omega(n/\ln n)$ and $m = o(n)$. In particular, $m = \lceil n/\ln \ln n \rceil$ yields the bound stated above. Our algorithm first generates short vectors in the lattice $\mathcal{L}_q^\perp([- \mathbf{A} \mid \mathbf{I}_n])$, and then applies the Aharonov–Regev Fourier distinguisher [2]. The main ingredient is a new instantiation of the auxiliary superlattices in the Gaussian–Wagner framework of [15]. These superlattices have small smoothing parameters even at exponentially small approximation error, while admitting exact shifted discrete Gaussian sampling within subexponential time. This yields a relation generator with the accuracy needed for LWE, without relying on sample amplification.

1.2 Technical Overview

Throughout this overview, assume the parameter regime of Theorem 1.2.

The Gaussian–Wagner framework. We first describe the construction of Ducas, Engelberts, and Loyer [15] for a systematic SIS matrix $\mathbf{B} = [\mathbf{B}_0 \mid \mathbf{I}_n] \in \mathbb{Z}_q^{n \times (k+n)}$. Partition the rows of $\mathbf{B}_0$ into blocks of sizes $b_1, \dots, b_r$, set $n_0 = 0$ and $n_i = \sum_{j \leq i} b_j$, and let $\mathbf{B}_i$ consist of its first n_i rows. The projected lattices are $\mathcal{L}_0 = \mathbb{Z}^k$ and

$$\mathcal{L}_i = \mathcal{L}_q^\perp([\mathbf{B}_i \mid \mathbf{I}_{n_i}]) = \{(\mathbf{u}, \mathbf{v}) \in \mathbb{Z}^k \times \mathbb{Z}^{n_i} : \mathbf{B}_i \mathbf{u} + \mathbf{v} = \mathbf{0} \bmod q\}.$$

Thus $\mathcal{L}_r = \mathcal{L}_q^\perp(\mathbf{B})$, and deleting the last b_i coordinates maps $\mathcal{L}_i$ onto $\mathcal{L}_{i-1}$. In round i , they choose $\mathcal{S}_i = \frac{q}{p_i} \cdot \mathbb{Z}^{b_i}$ and form $\mathcal{L}'_i = \mathcal{L}_i + (\{\mathbf{0}\} \times \mathcal{S}_i)$, where the zero vector occupies the first $k + n_{i-1}$ coordinates. This gives $\mathcal{L}'_i / \mathcal{L}_i \simeq \mathcal{S}_i / q\mathbb{Z}^{b_i} \simeq \mathbb{Z}_{p_i}^{b_i}$.

The lifting step has a concrete coordinate description. Let $\widehat{\mathbf{B}}_i$ denote the new block of b_i rows of $\mathbf{B}_0$. Given a current vector $\mathbf{x} = (\mathbf{u}, \mathbf{v}) \in \mathcal{L}_{i-1}$, choose an integer vector $\mathbf{r}$ satisfying $\mathbf{r} = -\widehat{\mathbf{B}}_i \mathbf{u} \bmod q$. Then $(\mathbf{x}, \mathbf{r}) \in \mathcal{L}_i$, but the new coordinates $\mathbf{r}$ may be much longer than the current Gaussian width. Instead of retaining this extension, sample $\mathbf{z} \leftarrow D_{\mathcal{S}_i, s_{i-1}, -\mathbf{r}}$ and output $\mathbf{x}' = (\mathbf{x}, \mathbf{r} + \mathbf{z}) \in \mathcal{L}'_i$. The old coordinates are unchanged, while the new block follows a centered discrete Gaussian on the coset $\mathbf{r} + \mathcal{S}_i$. When $s_{i-1} \geq \eta_\varepsilon(\mathcal{S}_i)$, the Gaussian masses of these cosets are nearly independent of $\mathbf{r}$, so lifting approximately preserves the Gaussian distribution on the whole lattice.

The combining step removes the auxiliary component. We bucket the lifted vectors $\mathbf{x}'$ according to their cosets in $\mathcal{L}'_i / \mathcal{L}_i$. Since $\mathcal{L}'_i / \mathcal{L}_i \simeq \mathcal{S}_i / q\mathbb{Z}^{b_i}$, there are $p_i^{b_i}$ buckets. Subtracting disjoint pairs $(\mathbf{x}'_1, \mathbf{x}'_2)$ from the same bucket gives vectors $\mathbf{x}'_1 - \mathbf{x}'_2$ in $\mathcal{L}_i$. Under the required smoothing conditions, these differences remain approximately Gaussian, with width increased from s_{i-1} to $s_i = \sqrt{2}s_{i-1}$. Starting with $3^r N$ independent samples from $D_{\mathbb{Z}^k, s_0}$ and assuming $N \geq p_i^{b_i}$ for every $i \in [r]$, round i receives $3^{r-i+1} N \geq 3p_i^{b_i}$ vectors and retains exactly $3^{r-i} N$ outputs. Thus the final list contains N vectors.

The coordinate construction $\mathcal{S}_i = \frac{q}{p_i} \cdot \mathbb{Z}^{b_i}$ admits efficient Gaussian sampling, but its smoothing parameter $\eta_\varepsilon(\mathcal{S}_i)$ is of order $(q/p_i)\sqrt{\ln(b_i/\varepsilon)}$. Thus, the lifting condition $s_{i-1} \geq \eta_\varepsilon(\mathcal{S}_i)$ requires $p_i = \Omega((q/s_{i-1})\sqrt{\ln(1/\varepsilon)})$. At the precision needed for the direct LWE application, $\varepsilon = 2^{-\tilde{\Omega}(n)}$, this forces $p_i \geq n^{\Omega(1)}$ even in the later rounds, since $s_{i-1} \leq q$. Consequently, the list size constraint $p_i^{b_i} \leq N$ limits each round to $b_i = O(\ln N / \ln n)$ coordinates. With only $r = O(\ln n)$ rounds, the requirement $\sum_{i=1}^r b_i = n$ forces $\ln N = \Omega(n)$. Thus, the required list size becomes exponential, and the original instantiation loses its subexponential complexity.

Auxiliary superlattices with smaller smoothing parameters. Therefore, a natural idea is to find a superlattice of $q\mathbb{Z}^b$ with both small smoothing parameters and sufficiently efficient Gaussian sampling. We retain the same lifting and combining operations, but replace the coordinate lattice on each new block. For a block dimension $b \geq 2$ and a prime P , sample $\mathbf{h} \leftarrow \mathbb{Z}_P^b \setminus \{\mathbf{0}\}$ and set $\mathcal{S}_{\mathbf{h}} := q(\mathcal{L}_P^\perp(\mathbf{h}))^* = q\mathbb{Z}^b + \frac{q}{P}\mathbf{h}\mathbb{Z}$, $|\mathcal{S}_{\mathbf{h}}/q\mathbb{Z}^b| = P$, where $\mathcal{L}_P^\perp(\mathbf{h}) = \{\mathbf{z} \in \mathbb{Z}^b : \langle \mathbf{h}, \mathbf{z} \rangle = 0 \bmod P\}$. In round i , we use $\mathcal{S}_i = \mathcal{S}_{\mathbf{h}_i}$ and $\mathcal{L}'_i = \mathcal{L}_i + (\{\mathbf{0}\} \times \mathcal{S}_i)$. The quotient now has P

classes, represented by $\frac{q}{P}a\mathbf{h}_i$ and labeled by $a \in \mathbb{Z}_P$. Here, P is the number of buckets in each round and determines the required list size N . Then, we obtain $\eta_\varepsilon(\mathcal{S}_{\mathbf{h}}) = O(\frac{q}{P^{1/b}}(1 + \sqrt{\frac{\ln(1+1/\varepsilon)}{b}}))$ with high probability in our parameter regime (Corollary 3.3). This follows from the large Euclidean minimum of its dual, which matches Minkowski's upper bound up to a constant factor, together with the Euclidean smoothing estimate (Lemma 2.9). Unlike the ℓ_∞ estimate used in [15], this estimate captures the Euclidean advantage of the new lattice and yields the division by $\sqrt{b}$ in the accuracy dependent term.

In particular, at $\varepsilon = e^{-n}$, our smoothing bound is $O((q/P^{1/b})(1 + \sqrt{n/b}))$, compared with $(q/P^{1/b})\sqrt{n}$ for the coordinate construction of comparable index. The dependence on n/b rather than n allows larger blocks to be processed in later rounds at the required Gaussian widths. With the block sizes chosen below, all lifting conditions $\eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i}) \leq s_{i-1}$ hold while $\sum_{i=1}^r b_i = n$ and $\ln P = o(n)$. This removes the auxiliary smoothing bottleneck of [15], allowing exponentially small ε while retaining subexponential complexity.⁷

Exact shifted Gaussian sampling. To use the new auxiliary lattices in the lifting step, we need exact samples from $D_{\mathcal{S}_{\mathbf{h}},s,\mathbf{c}}$. Following the coset sampling viewpoint of [16,8], we decompose $\mathcal{S}_{\mathbf{h}}$ into the P cosets $C_a = q\mathbb{Z}^b + \frac{q}{P}a\mathbf{h}$. The shifted Gaussian mass of each coset factors into a product of b one dimensional theta sums. We use certified evaluations of these sums and rejection sampling to select a coset with probability proportional to its mass. Conditioned on this coset, the coordinates can be sampled independently by scaling and shifting exact discrete Gaussian samples over $\mathbb{Z}$. Theorem 3.6 implements this procedure in the random bit model with expected time $P \text{poly}(n, \log q)$ per lift and polynomial expected auxiliary space in our parameter regime. Crucially, we do not need polynomial time sampling: the additional factor P per lift is compatible with the overall subexponential bound.

Applying the sampler to LWE. With the above sampler achieving exponentially small conditional approximation error, the remaining analysis largely follows the standard Gaussian Fourier approach of [2]; see also [31].

Given an input $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ and the norm bound B , we first use the above sampler to generate a list of short vectors $\mathbf{y}_i = (\mathbf{u}_i, \mathbf{v}_i) \in \mathcal{L}_q^\perp([- \mathbf{A} \mid \mathbf{I}_n])$, using only $\mathbf{A}$. This choice also makes the projected chain start from $\mathbb{Z}^m$, allowing the sampler to operate even when $m < n$. Each relation satisfies $\mathbf{v}_i = \mathbf{A}\mathbf{u}_i \bmod q$. For good choices of the matrix and auxiliary lattices, the output list is conditionally δ -similar to independent samples from $D_{\mathcal{L}_q^\perp([- \mathbf{A} \mid \mathbf{I}_n]),w}$, where $w = q/\tau$ and $\delta = e^{-O(n)}$.

For an LWE input $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e}$, we have $\langle \mathbf{b}, \mathbf{u}_i \rangle = \langle (\mathbf{e}, \mathbf{s}), \mathbf{y}_i \rangle \bmod q$. Under the norm promise $\|\mathbf{s}\|_2^2 + \|\mathbf{e}\|_2^2 \leq B^2(m+n)$, Poisson summation gives an

⁷ The analysis of Lift steps also needs the widths $s_i \geq \sqrt{2}\eta_\varepsilon(\mathcal{L}'_{i+1})$. A similar ℓ_2 -dual minimum argument bounds these parameters simultaneously, allowing the additional constraints to be satisfied in our parameter regime even for exponentially small ε (Corollary 3.13).

ideal Fourier bias $\mathbb{E}_{\mathbf{y} \leftarrow D_{\mathcal{L}_q^\perp([- \mathbf{A} | \mathbf{I}_n]), w}} \cos(\frac{2\pi}{q} \langle (\mathbf{e}, \mathbf{s}), \mathbf{y} \rangle)$ at least $e^{-\pi B^2(m+n)/\tau^2} / (1 + e^{-n})$, which dominates the sampler's approximation error. To detect this bias, we retain the first M vectors, compute $\langle \mathbf{b}, \mathbf{u}_i \rangle \bmod q$ for each one, and estimate the average $T = \frac{1}{M} \sum_{i=1}^M \cos(\frac{2\pi}{q} \langle \mathbf{b}, \mathbf{u}_i \rangle)$ to sufficient precision. The distinguisher outputs LWE if the computed average is at least $\frac{1}{2} e^{-\pi B^2(m+n)/\tau^2}$, and uniform otherwise.

A choice of $M = \exp(O(\ln n + (m+n)/\tau^2))$ suffices and is smaller than the sampler's output list, so the Fourier test does not change the overall complexity. Completeness follows from conditional similarity and a martingale bound. For soundness, after excluding some bad events, character orthogonality makes the cosine terms have mean zero and be pairwise uncorrelated under uniform $\mathbf{b}$; a second moment bound then controls false acceptance.

1.3 Organization of the Paper

Section 2 gives the required lattice and Gaussian preliminaries. Section 3 constructs and analyzes the auxiliary superlattices and the resulting Gaussian–Wagner sampler. Section 4 applies the sampler to the relation lattice and proves the LWE decision theorem.

2 Preliminaries

2.1 Notations

For every positive integer n , we write $[n] := \{1, \dots, n\}$. We write $\mathbb{Z}_q := \mathbb{Z}/q\mathbb{Z}$, and when q is prime we also write $\mathbb{F}_q$ for the finite field. We use $a = e^{\pm\delta}b$ as shorthand for $e^{-\delta}b \leq a \leq e^{\delta}b$. For a finite set S , $|S|$ denotes its cardinality, and $x \leftarrow S$ denotes sampling an element uniformly from S and assigning it to x . For a probability distribution $\mathcal{D}$, $x \leftarrow \mathcal{D}$ denotes sampling x according to $\mathcal{D}$. For an abelian group G and a subgroup $H \subseteq G$, we write G/H for the quotient group. For full rank lattices $\mathcal{L} \subseteq \mathcal{L}'$, we write $|\mathcal{L}'/\mathcal{L}|$ for their finite index. For a tuple $x = (x_1, \dots, x_n)$ and $i \in [n]$, we write $x_{-i} := (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$. The notation $x_{-\{i,j\}}$ is defined analogously.

Vectors are viewed as column vectors. We write $\mathbf{0}$ for the zero vector and zero matrix of the relevant dimensions, and $\mathbf{I}_n$ for the n -dimensional identity matrix. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we write $\langle \mathbf{x}, \mathbf{y} \rangle$ for their Euclidean inner product and $\|\mathbf{x}\|_2$ and $\|\mathbf{x}\|_\infty$ for their Euclidean and infinity norms. When there is no ambiguity, we write $\|\mathbf{x}\|$ in place of $\|\mathbf{x}\|_2$. We denote by $\mathcal{B}_n(r) := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq r\}$ the closed Euclidean ball of radius r centered at the origin.

All asymptotic statements are with respect to $n \rightarrow \infty$. We use the standard notation $O(\cdot)$, $\Omega(\cdot)$, $\Theta(\cdot)$, $o(\cdot)$, and $\omega(\cdot)$. We use $\ln$ for the natural logarithm and $\log$ for the binary logarithm.

In complexity bounds only, we slightly overload the logarithm notation as follows. For a rational number $x = u/v \in \mathbb{Q}$, written in lowest terms with $v > 0$, define $\log |x| := 1 + \lceil \log(1 + |u|) \rceil + \lceil \log(1 + v) \rceil$. For $x > 0$, we write

$\log x$ in place of $\log |x|$. For a rational vector $\mathbf{x} = (x_1, \dots, x_n)^\top \in \mathbb{Q}^n$, define $\log \|\mathbf{x}\|_\infty := \max_{j \in [n]} \log |x_j|$. For a rational matrix, $\log \|\mathbf{A}\|_\infty$ is defined analogously as the maximum encoding length of its entries. Thus, in complexity bounds, $\log x$, $\log \|\mathbf{x}\|_\infty$, and $\log \|\mathbf{A}\|_\infty$ denote encoding lengths rather than numerical logarithms.

2.2 Conditional Similarity

For two probability distributions P and Q on a countable set Ω , their statistical distance is $\Delta(P, Q) := \frac{1}{2} \sum_{\omega \in \Omega} |P(\omega) - Q(\omega)|$. Following [1, 15], we use a pointwise multiplicative notion of similarity. This notion is stronger than statistical distance and is particularly convenient for analyzing lists produced by bucket-and-combine algorithms.

Definition 2.1 (Similarity). *Let D be a probability distribution on a countable set Ω and let $\delta \geq 0$. A random variable $X \in \Omega$ is δ -similar to D if, for every $x \in \Omega$, $\Pr[X = x] = e^{\pm\delta} D(x)$.*

If X is δ -similar to D , then $\Delta(X, D) \leq e^\delta - 1$. In particular, for $0 \leq \delta \leq 1$, $\Delta(X, D) \leq 2\delta$.

Definition 2.2 (Conditional similarity). *Let D be a probability distribution on a countable set Ω and let $\delta \geq 0$. Random variables $X_1, \dots, X_N \in \Omega$ are conditionally δ -similar to independent samples from D if, for every $i \in [N]$ and every $x = (x_1, \dots, x_N) \in \Omega^N$ satisfying $\Pr[X_{-i} = x_{-i}] > 0$, $\Pr[X_i = x_i \mid X_{-i} = x_{-i}] = e^{\pm\delta} D(x_i)$.*

Conditional 0-similarity is equivalent to being independently distributed according to D . Conditional δ -similarity also implies that every individual marginal is δ -similar to D . We record the consequences of conditional similarity needed below, already observed in [1] and used in [15].

Lemma 2.3 ([1]). *Let $X_1, \dots, X_N$ be conditionally δ -similar to independent samples from a distribution D on a countable set Ω .*

1. *Every sublist of $X_1, \dots, X_N$ remains conditionally δ -similar to independent samples from D .*
2. *Let $\mathcal{F}$ be a sigma-algebra generated by arbitrary functions of $(X_j)_{j \neq i}$. Then, almost surely, $\Pr[X_i = x \mid \mathcal{F}] = e^{\pm\delta} D(x)$, $x \in \Omega$. Consequently, for every function $g : \Omega \rightarrow [-1, 1]$, $|\mathbb{E}[g(X_i) \mid \mathcal{F}] - \mathbb{E}_{X \leftarrow D}[g(X)]| \leq 2(e^\delta - 1)$ almost surely.*
3. *For every event $E \subseteq \Omega$, $\Pr[X_i \in E] \leq e^\delta D(E)$, where $D(E) := \Pr_{X \leftarrow D}[X \in E]$.*
4. *For distinct i, j and every event $E \subseteq \Omega^2$, $\Pr[(X_i, X_j) \in E] \leq e^{2\delta} (D \times D)(E)$.*

2.3 Basic Lattices

Lattices and duality. Let $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_k) \in \mathbb{R}^{n \times k}$ have linearly independent columns. The lattice generated by $\mathbf{B}$ is $\mathcal{L}(\mathbf{B}) := \mathbf{B}\mathbb{Z}^k = \left\{ \sum_{i=1}^k z_i \mathbf{b}_i : z_i \in \mathbb{Z} \right\}$. We say that $\mathbf{B}$ is a basis of $\mathcal{L}(\mathbf{B})$. The rank of $\mathcal{L}(\mathbf{B})$ is k , and its real span is $\text{span}_{\mathbb{R}}(\mathcal{L}(\mathbf{B})) = \text{span}_{\mathbb{R}}(\mathbf{B})$. When $k = n$, the lattice is full rank. The determinant of a rank- k lattice with basis $\mathbf{B}$ is $\det(\mathcal{L}(\mathbf{B})) := \sqrt{\det(\mathbf{B}^\top \mathbf{B})}$. For a full rank sublattice $\mathcal{L} \subseteq \mathcal{L}'$ in the same ambient space, $|\mathcal{L}'/\mathcal{L}| = \det(\mathcal{L})/\det(\mathcal{L}')$.

The dual lattice of $\mathcal{L}$ is $\mathcal{L}^* := \{\mathbf{y} \in \text{span}_{\mathbb{R}}(\mathcal{L}) : \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}, \forall \mathbf{x} \in \mathcal{L}\}$. For every nonzero real a , $(a\mathcal{L})^* = a^{-1}\mathcal{L}^*$, and the first Euclidean minimum of $\mathcal{L}$ is $\lambda_1(\mathcal{L}) := \min_{\mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}\}} \|\mathbf{x}\|_2$.

Projections and Primitive Sublattices For any $S \subseteq \mathbb{R}^n$, we write π_S for the projection onto $\text{span}_{\mathbb{R}}(S)$ and $\pi_S^\perp$ for the projection orthogonal to $\text{span}_{\mathbb{R}}(S)$.

We say $\mathcal{M}$ is a primitive sublattice of a lattice $\mathcal{L}$ if and only if $\mathcal{M} = \text{span}_{\mathbb{R}}(\mathcal{M}) \cap \mathcal{L}$. Primitivity is equivalent to the existence of a $\mathbb{Z}$ -module complement $\mathcal{N}$ with $\mathcal{L} = \mathcal{M} \oplus \mathcal{N}$, i.e., $\mathcal{M} + \mathcal{N} = \mathcal{L}$ and $\mathcal{M} \cap \mathcal{N} = \{\mathbf{0}\}$. Let $\mathcal{Q} = \pi_{\mathcal{M}}^\perp(\mathcal{L})$. The orthogonal projection $\pi_{\mathcal{M}}^\perp$ induces a bijection $\mathcal{N} \rightarrow \mathcal{Q}$, and both directions can be computed in polynomial time given the basis of $\mathcal{M}$ and $\mathcal{L}$. The representatives do not depend on the basis of $\mathcal{M}$ and $\mathcal{L}$.

q-ary lattices An n -dimensional lattice $\mathcal{L}$ is called q -ary lattice if $q\mathbb{Z}^n \subseteq \mathcal{L} \subseteq \mathbb{Z}^n$. The following two q -ary lattices are widely used in cryptography and will be used frequently throughout this work. For a wide matrix $\mathbf{A} \in \mathbb{Z}^{n \times m}$ where $m \geq n$, we define the following two m -dimensional q -ary lattices,

$$\begin{aligned} \mathcal{L}_q^\perp(\mathbf{A}) &:= \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = \mathbf{0} \bmod q\} \\ \mathcal{L}_q(\mathbf{A}) &:= \{\mathbf{y} \in \mathbb{Z}^m : \exists \mathbf{s} \in \mathbb{Z}_q^n \text{ s.t. } \mathbf{y} = \mathbf{A}^\top \mathbf{s} \bmod q\}. \end{aligned}$$

It is worth noting that $\mathcal{L}_q(\mathbf{A}) = \mathbf{A}^\top \mathbb{Z}^n + q\mathbb{Z}^m$ and $\mathcal{L}_q^\perp(\mathbf{A})$ and $\mathcal{L}_q(\mathbf{A})$ are duals up to q scaling, that is $\mathcal{L}_q(\mathbf{A}) = q \cdot (\mathcal{L}_q^\perp(\mathbf{A}))^*$ and $\mathcal{L}_q^\perp(\mathbf{A}) = q \cdot (\mathcal{L}_q(\mathbf{A}))^*$. Moreover, $\det(\mathcal{L}_q^\perp(\mathbf{A})) \leq q^n$ with equality if and only if $\mathbf{A}$ is surjective over $\mathbb{Z}_q$. In particular, when q is prime, equality holds if and only if $\mathbf{A}$ has full rank over $\mathbb{F}_q$. For a full rank $\mathbf{A} \in \mathbb{Z}^{n \times m}$ of the form $[\overline{\mathbf{A}} \mid \mathbf{I}_n]$, a basis of $\mathcal{L}_q^\perp(\mathbf{A})$ is $\begin{pmatrix} \mathbf{0} & \mathbf{I}_{m-n} \\ q\mathbf{I}_n & -\overline{\mathbf{A}} \end{pmatrix}$.

2.4 Discrete Gaussian Measures over Lattices

For $s > 0$, a center $\mathbf{c} \in \mathbb{R}^d$, and $\mathbf{x} \in \mathbb{R}^d$, define $\rho_{s,\mathbf{c}}(\mathbf{x}) := \exp\left(-\pi \frac{\|\mathbf{x}-\mathbf{c}\|_2^2}{s^2}\right)$. When $\mathbf{c} = \mathbf{0}$, we write ρ_s in place of $\rho_{s,\mathbf{0}}$. For a countable set $S \subseteq \mathbb{R}^d$, define $\rho_{s,\mathbf{c}}(S) := \sum_{\mathbf{x} \in S} \rho_{s,\mathbf{c}}(\mathbf{x})$.

Let $\mathcal{L} \subseteq \mathbb{R}^d$ be a lattice, let $\mathbf{t} \in \mathbb{R}^d$, and let $\mathbf{c} \in \mathbb{R}^d$. The discrete Gaussian distribution over the coset $\mathcal{L} + \mathbf{t}$, with parameter s and center $\mathbf{c}$, is defined by $\Pr_{X \leftarrow D_{\mathcal{L}+\mathbf{t},s,\mathbf{c}}}[X = \mathbf{x}] := \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(\mathcal{L}+\mathbf{t})}$, $\mathbf{x} \in \mathcal{L} + \mathbf{t}$. We write $D_{\mathcal{L},s,\mathbf{c}} := D_{\mathcal{L}+\mathbf{0},s,\mathbf{c}}$,

$D_{\mathcal{L}+\mathbf{t},s} := D_{\mathcal{L}+\mathbf{t},s,\mathbf{0}}$, and omit both the center and the parameter when they are respectively $\mathbf{0}$ and 1.

The Gaussian distribution respects scaling and translation. Namely, if $X \leftarrow D_{\mathcal{L}+\mathbf{t},s,\mathbf{c}}$, then $aX + \mathbf{z} \leftarrow D_{a\mathcal{L}+a\mathbf{t}+\mathbf{z},|a|s,a\mathbf{c}+\mathbf{z}}$ for every nonzero real a and every $\mathbf{z} \in \mathbb{R}^d$. Moreover, discrete Gaussians over Cartesian products factor into independent coordinates. In particular, $D_{\mathbb{Z}^d,s,\mathbf{c}} = \bigotimes_{j=1}^d D_{\mathbb{Z},s,c_j}$.

We use the following consequences of the Poisson summation formula [26, Lemma 2.8].

Lemma 2.4 ([26]). *Let $\mathcal{L} \subseteq \mathbb{R}^d$ be a full rank lattice. For every $s > 0$ and $\mathbf{c}, \mathbf{t} \in \mathbb{R}^d$, $\rho_{s,\mathbf{c}}(\mathcal{L}) = \frac{s^d}{\det(\mathcal{L})} \sum_{\mathbf{y} \in \mathcal{L}^*} e^{-\pi s^2 \|\mathbf{y}\|_2^2} e^{-2\pi i \langle \mathbf{y}, \mathbf{c} \rangle}$, and $\sum_{\mathbf{x} \in \mathcal{L}} \rho_s(\mathbf{x}) e^{2\pi i \langle \mathbf{t}, \mathbf{x} \rangle} = \frac{s^d}{\det(\mathcal{L})} \sum_{\mathbf{y} \in \mathcal{L}^*} e^{-\pi s^2 \|\mathbf{y} - \mathbf{t}\|_2^2}$. Consequently, for $X \leftarrow D_{\mathcal{L},s}$, $\mathbb{E}[e^{2\pi i \langle \mathbf{t}, X \rangle}] = \frac{\rho_{1/s}(\mathcal{L}^* - \mathbf{t})}{\rho_{1/s}(\mathcal{L}^*)}$.*

The following one-dimensional exact sampler is due to Brakerski, Langlois, Peikert, Regev, and Stehlé [9, Section 5.1].

Lemma 2.5 ([9]). *There is a randomized algorithm that, given $s \in \mathbb{Q}_{>0}$ and $c \in \mathbb{Q}$, returns an exact sample from $D_{\mathbb{Z},s,c}$ in expected time $\text{poly}(\log s, \log |c|)$.*

2.5 Bounds on Smoothing Parameters

Micciancio and Regev introduced the following lattice quantity.

Definition 2.6 (Smoothing parameter [26]). *For a lattice $\mathcal{L}$ and $\varepsilon > 0$, the smoothing parameter is $\eta_\varepsilon(\mathcal{L}) := \inf \{s > 0 : \rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon\}$.*

We will repeatedly use the following elementary properties. If $0 < \varepsilon_1 \leq \varepsilon_2$, then $\eta_{\varepsilon_1}(\mathcal{L}) \geq \eta_{\varepsilon_2}(\mathcal{L})$. For every nonzero real a , $\eta_\varepsilon(a\mathcal{L}) = |a|\eta_\varepsilon(\mathcal{L})$. In particular, isometric lattices have the same smoothing parameter. Poisson summation gives both a shifted-mass inequality and the usual smoothing estimate.

Lemma 2.7 ([26]). *Let $\mathcal{L} \subseteq \mathbb{R}^d$ be a full rank lattice. For every $s > 0$ and $\mathbf{c} \in \mathbb{R}^d$, $\rho_{s,\mathbf{c}}(\mathcal{L}) \leq \rho_s(\mathcal{L})$. Moreover, if $s \geq \eta_\varepsilon(\mathcal{L})$ for some $0 < \varepsilon < 1$, then $(1 - \varepsilon) \frac{s^d}{\det(\mathcal{L})} \leq \rho_{s,\mathbf{c}}(\mathcal{L}) \leq (1 + \varepsilon) \frac{s^d}{\det(\mathcal{L})}$, and hence $\frac{\rho_{s,\mathbf{c}}(\mathcal{L})}{\rho_s(\mathcal{L})} \in \left[\frac{1-\varepsilon}{1+\varepsilon}, 1 \right]$.*

The smoothing parameter of the integer lattice satisfies the following standard bound.

Lemma 2.8 ([26, Lemma 3.3]). *For every $d \geq 1$ and $\varepsilon > 0$, $\eta_\varepsilon(\mathbb{Z}^d) \leq \sqrt{\frac{\ln(2d(1+1/\varepsilon))}{\pi}}$.*

We next recall the relation between the first Euclidean minimum of a lattice and the smoothing parameter of its dual.

Lemma 2.9 ([14, Lemma 2.17]). *Let $\mathcal{L} \subseteq \mathbb{R}^d$ be a d -dimensional lattice. Then, for every $\varepsilon \in (0, 1)$, $\eta_\varepsilon(\mathcal{L}^*) \leq \frac{\sqrt{\ln(1+1/\varepsilon)/\pi} + \sqrt{d/(2\pi)}}{\lambda_1(\mathcal{L})}$.*

The smoothing parameter cannot increase under orthogonal projection.

Lemma 2.10. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a full rank lattice, let $V \subseteq \mathbb{R}^n$ be a linear subspace, and suppose that $\pi_V(\mathcal{L})$ is a lattice in V . Then, for every $\varepsilon > 0$, $\eta_\varepsilon(\pi_V(\mathcal{L})) \leq \eta_\varepsilon(\mathcal{L})$.*

Proof. Under the natural identification of V with its image in $\mathbb{R}^n$, $\pi_V(\mathcal{L})^* = \mathcal{L}^* \cap V$. Indeed, for every $\mathbf{y} \in V$ and $\mathbf{x} \in \mathcal{L}$, $\langle \mathbf{y}, \pi_V(\mathbf{x}) \rangle = \langle \mathbf{y}, \mathbf{x} \rangle$. Consequently, for every $s > 0$, $\rho_{1/s}(\pi_V(\mathcal{L})^* \setminus \{\mathbf{0}\}) \leq \rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\})$, which proves the claim. $\square$

2.6 Concentration Inequalities

We recall the two concentration inequalities used in the LWE application.

Lemma 2.11 (Chebyshev's inequality). *Let X be a real random variable with finite variance. Then, for every $t > 0$, $\Pr[|X - \mathbb{E}[X]| \geq t] \leq \frac{\text{Var}(X)}{t^2}$.*

A sequence $D_1, \dots, D_M$ adapted to a filtration $\mathcal{F}_0 \subseteq \dots \subseteq \mathcal{F}_M$ is called a martingale-difference sequence if $\mathbb{E}[D_i \mid \mathcal{F}_{i-1}] = 0$ for every $i \in [M]$.

Lemma 2.12 (Azuma–Hoeffding inequality). *Let $D_1, \dots, D_M$ be a martingale-difference sequence with respect to a filtration $(\mathcal{F}_i)_{i=0}^M$. Suppose that $|D_i| \leq c_i$ almost surely. Then, for every $t > 0$, $\Pr\left[\sum_{i=1}^M D_i \leq -t\right] \leq \exp\left(-\frac{t^2}{2\sum_{i=1}^M c_i^2}\right)$.*

3 A new instantiation of Gaussian-Wagner in [15]

This section develops the Gaussian–Wagner sampler used in our LWE algorithm. We first construct one-check auxiliary superlattices with small smoothing parameters and give a shifted Gaussian sampler for them. We then incorporate these lattices into the lifting-and-combining framework [15] and choose parameters that achieve exponentially small conditional approximation error within subexponential time.

3.1 The random superlattice of $q\mathbb{Z}^b$ and its Gaussian sampler

Random one-check superlattices Let P be prime, let $b \geq 2$, and choose $\mathbf{h} \leftarrow \mathbb{Z}_P^b \setminus \{\mathbf{0}\}$. Define $\mathcal{L}_P^\perp(\mathbf{h}) = \{\mathbf{z} \in \mathbb{Z}^b : \langle \mathbf{h}, \mathbf{z} \rangle = 0 \bmod P\}$, and $\mathcal{S}_\mathbf{h} = q \cdot (\mathcal{L}_P^\perp(\mathbf{h}))^* = q\mathbb{Z}^b + \frac{q}{P}\mathbf{h}\mathbb{Z}$. We have the following facts: $|\mathbb{Z}^b / \mathcal{L}_P^\perp(\mathbf{h})| = P$, $|\mathcal{S}_\mathbf{h} / q\mathbb{Z}^b| = P$, and $\mathcal{S}_\mathbf{h}^* = q^{-1} \cdot \mathcal{L}_P^\perp(\mathbf{h})$.

Next, we introduce some basic facts about $\mathcal{L}_P^\perp(\mathbf{h})$ and $\mathcal{S}_\mathbf{h}$.

Fact 3.1 *For every $r \geq 0$ and $b \geq 1$, $|\mathbb{Z}^b \cap \mathcal{B}_b(r)| \leq \lceil \sqrt{2\pi e} \left(\frac{r}{\sqrt{b}} + \frac{1}{2} \right) \rceil^b$.*

Proof. The unit cubes centered at the integer points are disjoint and lie in $\mathcal{B}_b(r + \sqrt{b}/2)$. Applying the standard bound $\text{vol}(\mathcal{B}_b(1)) \leq (2\pi e/b)^{b/2}$ proves the claim.

The next lemma is a Euclidean-norm analogue of the counting-and-union-bound argument used in [17, Lemma 5.3], which proves a high-probability lower bound on the ℓ_∞ -minimum of a random q -ary lattice.

Lemma 3.2 (Lower bound on $\lambda_1(\mathcal{L}_P^\perp(\mathbf{h}))$). *Let $\gamma_0 = 1/(32\sqrt{2\pi e})$, $b > 1$ and $P > 24^b$ be a prime. Then $\Pr_{\mathbf{h}}[\lambda_1(\mathcal{L}_P^\perp(\mathbf{h})) < \gamma_0\sqrt{b}P^{1/b}] \leq 6^{-b}$.*

Proof. Set $r = \gamma_0\sqrt{b}P^{1/b}$. By Fact 3.1, $|\mathbb{Z}^b \cap \mathcal{B}_b(r)| \leq [P^{1/b}/32 + \sqrt{2\pi e}/2]^b \leq [P^{1/b}/32 + P^{1/b}/8]^b \leq P/6^b$. The condition $r < P$ excludes every nonzero point of $P\mathbb{Z}^b$. For a fixed $\mathbf{z} \not\equiv 0 \pmod{P}$, $\Pr_{\mathbf{h} \neq \mathbf{0}}[\langle \mathbf{h}, \mathbf{z} \rangle = 0] = \frac{P^{b-1}-1}{P^b-1} < \frac{1}{P}$. A union bound over nonzero integer points in the ball proves the claim. $\square$

Applying Lemma 2.9 to the scaled lattice $\mathcal{S}_{\mathbf{h}} = q \cdot (\mathcal{L}_P^\perp(\mathbf{h}))^*$ gives the following corollary that we will use below.

Corollary 3.3 (Upper Bound on $\eta_\varepsilon(\mathcal{S}_{\mathbf{h}})$). *Under the conditions of Lemma 3.2, except with probability 6^{-b} , $\eta_\varepsilon(\mathcal{S}_{\mathbf{h}}) \leq 32\sqrt{e} \cdot \frac{q}{P^{1/b}} \left(\sqrt{\frac{2\ln(1+1/\varepsilon)}{b}} + 1 \right)$.*

Proof. Since $(q\mathcal{L})^* = q^{-1}\mathcal{L}^*$ and $\mathcal{S}_{\mathbf{h}} = q \cdot (\mathcal{L}_P^\perp(\mathbf{h}))^*$, the claim follows by combining Lemmas 2.9 and 3.2. $\square$

Next we show that shifted discrete Gaussians over $\mathcal{S}_{\mathbf{h}}$ can be sampled exactly in expected time $P \text{poly}(\text{bits})$, where bits denotes the bit length of the sampler input.

Exact Gaussian sampling on $\mathcal{S}_{\mathbf{h}}$ Write $\mathbf{h} = (h_1, \dots, h_b)^\top \in \mathbb{Z}_P^b \setminus \{\mathbf{0}\}$. Recall the disjoint coset decomposition $\mathcal{S}_{\mathbf{h}} = \bigsqcup_{a \in \mathbb{Z}_P} C_a$, $C_a := q\mathbb{Z}^b + \frac{q}{P}a\mathbf{h}$. For a width $s \in \mathbb{Q}_{>0}$ and a center $\mathbf{c} = (c_1, \dots, c_b)^\top \in \mathbb{Q}^b$, define the Gaussian mass of the a th coset by

$$W_a(\mathbf{c}) := \rho_{s,\mathbf{c}}(C_a) = \prod_{j=1}^b \rho_{s,c_j} \left(q\mathbb{Z} + \frac{q}{P}ah_j \right). \quad (1)$$

Conditioned on the coset index a , the coordinates are independent.

It is useful to rewrite the factors in Equation (1) as one-dimensional theta sums. Set $\sigma := s/q$ and $\mu_{a,j} := c_j/q - ah_j/P$. For $\mu \in \mathbb{R}$ and $\sigma > 0$, write $\Theta(\mu, \sigma) := \sum_{z \in \mathbb{Z}} \exp\left(-\pi \frac{(z-\mu)^2}{\sigma^2}\right)$. Then we have $W_a(\mathbf{c}) = \prod_{j=1}^b \Theta(\mu_{a,j}, \sigma)$.

For $\Theta(\mu, \sigma)$, we use the following standard fact. When $\sigma < 1$, the defining series decays rapidly, and when $\sigma \geq 1$, the same holds for its Poisson-dual series. The proof is deferred to Appendix C

Lemma 3.4. *Let $\mu \in \mathbb{Q}$ and $\sigma \in \mathbb{Q}_{>0}$. Then, for every integer $t \geq 1$, a certified dyadic interval of width at most 2^{-t} containing $\log \Theta(\mu, \sigma)$ can be computed in $\text{poly}(\log |\sigma|, t, \log |\mu|)$ time and space.*

Rather than explicitly normalizing all P computable-real weights, we use rejection sampling, which is also used in [17, 9]. The following elementary lemma is the key point.

Lemma 3.5. *Let $W_{\max}(\mathbf{c}) := \max_{a \in \mathbb{Z}_P} W_a(\mathbf{c})$. Assume $q, P \leq 2^{\text{poly}(n)}$, $0 < s < q$, $\frac{q}{s} \leq n^{O(1)}$. One can compute, in a streaming pass over $a \in \mathbb{Z}_P$, a positive dyadic rational $M(\mathbf{c})$ such that for every $a \in \mathbb{Z}_P$, $W_a(\mathbf{c}) \leq M(\mathbf{c})$ and $M(\mathbf{c}) < 4W_{\max}(\mathbf{c})$. The computation takes $P \text{poly}(n, b, \log q, \log |s|, \log \|\mathbf{c}\|_\infty)$ time and polynomial auxiliary space.*

Proof. Note that $W_a(\mathbf{c}) = \prod_{j=1}^b \Theta(\mu_{a,j}, \sigma)$, where $\sigma = s/q$, $\mu_{a,j} = c_j/q - ah_j/P$. Set $t_0 := \lceil \log(4b) \rceil$. For every $a \in \mathbb{Z}_P$ and $j \in [b]$, apply Lemma 3.4 to compute an interval of width at most $2^{-t_0} \leq 1/(4b)$ containing $\log \Theta(\mu_{a,j}, \sigma)$. Summing these b intervals gives $\log W_a(\mathbf{c}) \in [\underline{\omega}_a, \bar{\omega}_a]$ with $\bar{\omega}_a - \underline{\omega}_a \leq \frac{1}{4}$.

Under the encoding convention fixed in the preliminaries, $\log |\sigma| = O(\log |s| + \log q)$ and $\log |\mu_{a,j}| = O(\log q + \log P + \log \|\mathbf{c}\|_\infty)$. Since $\log P = \text{poly}(n)$ and $q/s \leq n^{O(1)}$, every interval can be computed in $\text{poly}(n, b, \log q, \log |s|, \log \|\mathbf{c}\|_\infty)$ time and polynomial space.

Define $m(\mathbf{c}) := \lceil \max_{a \in \mathbb{Z}_P} \bar{\omega}_a \rceil$, $M(\mathbf{c}) := 2^{m(\mathbf{c})}$. Clearly $M(\mathbf{c}) \geq W_a(\mathbf{c})$ for every a . Moreover, $\max_a \bar{\omega}_a \leq \log W_{\max}(\mathbf{c}) + \frac{1}{4}$, and therefore $M(\mathbf{c}) < 2^{5/4} W_{\max}(\mathbf{c}) < 4W_{\max}(\mathbf{c})$. The computation is streaming: after processing a value of a , only the current maximum of $\bar{\omega}_a$ needs to be retained. $\square$

For a proposed $a \in \mathbb{Z}_P$, set $p_a := \frac{W_a(\mathbf{c})}{M(\mathbf{c})}$. Although p_a is generally irrational, Lemma 3.4 gives certified dyadic intervals of arbitrary precision for $\log p_a = \log W_a(\mathbf{c}) - m(\mathbf{c})$ and hence, by directed exponentiation, for p_a itself. To sample a Bernoulli random variable of parameter p_a exactly, interpret a sequence of independent unbiased bits as the binary expansion of a lazy uniform random variable $U \in [0, 1]$. At precision t , maintain certified intervals $U \in I_t$, $p_a \in J_t$, each of width $O(2^{-t})$. Accept if I_t lies strictly below J_t , reject if I_t lies strictly above J_t , and otherwise increase the precision. No incorrect decision is ever made. Moreover, if the comparison remains unresolved at precision t , then $|U - p_a| = O(2^{-t})$, which occurs with probability $O(2^{-t})$. Since certified precision t can be computed in time polynomial in t and the input length, the comparison halts almost surely and has expected polynomial running time. Thus no rounded probability is ever used.

The sampler is now particularly simple.

The conditional one-dimensional samples follow by scaling. Namely, if $z_j \leftarrow D_{\mathbb{Z}, s/q, c_j/q - ah_j/P}$ then $qz_j + \frac{q}{P}ah_j \leftarrow D_{q\mathbb{Z} + (q/P)ah_j, s, c_j}$.

Theorem 3.6 (Exact Gaussian sampler on $\mathcal{S}_{\mathbf{h}}$). *Let $\mathbf{h} \in \mathbb{Z}_P^b \setminus \{\mathbf{0}\}$, $\mathbf{c} \in \mathbb{Q}^b$, and $s \in \mathbb{Q}_{>0}$, and assume $q, P \leq 2^{\text{poly}(n)}$, $0 < s < q$, $\frac{q}{s} \leq n^{O(1)}$. Then Algorithm 1 is a randomized algorithm whose output distribution is exactly $D_{\mathcal{S}_{\mathbf{h}}, s, \mathbf{c}}$. Its expected running time is $P \text{poly}(n, b, \log q, \log |s|, \log \|\mathbf{c}\|_\infty)$, and its expected auxiliary space is polynomial in the same parameters.*

Proof. Fix $a \in \mathbb{Z}_P$ and $\mathbf{x} = (x_1, \dots, x_b)^\top \in C_a = q\mathbb{Z}^b + \frac{q}{P}a\mathbf{h}$. In one proposal iteration, the probability of proposing a , accepting it, and then outputting $\mathbf{x}$ is $\frac{1}{P} \cdot \frac{W_a(\mathbf{c})}{M(\mathbf{c})} \cdot \prod_{j=1}^b \frac{\rho_{s, c_j}(x_j)}{\rho_{s, c_j}(q\mathbb{Z} + \frac{q}{P}ah_j)} = \frac{\rho_{s, \mathbf{c}}(\mathbf{x})}{P \cdot M(\mathbf{c})}$. The denominator is independent of $\mathbf{x}$. Therefore, conditioned on producing an output, one proposal iteration outputs

Algorithm 1 OneCheckDGS($\mathbf{h}, q, P, s, \mathbf{c}$)

Require: $\mathbf{h} \in \mathbb{Z}_P^b \setminus \{\mathbf{0}\}$, $\mathbf{c} \in \mathbb{Q}^b$, and $s \in \mathbb{Q}_{>0}$

Ensure: A vector $\mathbf{x} \in \mathcal{S}_{\mathbf{h}}$

1: Compute $M(\mathbf{c})$ from Lemma 3.5

2: **repeat**

3: Sample $a \leftarrow \mathbb{Z}_P$ uniformly

4: With probability $W_a(\mathbf{c})/M(\mathbf{c})$, accept a ; otherwise reject it

5: **until** a is accepted

6: **for** $j = 1, \dots, b$ **do**

7: Sample $x_j \leftarrow D_{q\mathbb{Z} + (q/P)ah_j, s, c_j}$

▷ By Lemma 2.5

8: **return** $\mathbf{x} = (x_1, \dots, x_b)^\top$

exactly the discrete Gaussian $D_{\mathcal{S}_{\mathbf{h}}, s, \mathbf{c}}$. Repeating rejected proposals does not change this conditional law.

The probability of accepting a proposed coset index is $p_{\text{acc}} = \frac{\sum_{a \in \mathbb{Z}_P} W_a(\mathbf{c})}{PM(\mathbf{c})} \geq \frac{W_{\max}(\mathbf{c})}{P \cdot M(\mathbf{c})} > \frac{1}{4P}$, where the last inequality follows from Lemma 3.5. Hence, the expected number of iterations is at most $4P$.

The preprocessing step of computing $M(\mathbf{c})$ costs $P \text{poly}(n, b, \log q, \log |s|, \log \|\mathbf{c}\|_\infty)$. Each proposal requires one exact uniform sample from $\mathbb{Z}_P$ and one lazy Bernoulli comparison. Uniform sampling from $\mathbb{Z}_P$ is obtained by rejection from $\lceil \log P \rceil$ random bits, and the expected cost of the Bernoulli comparison is polynomial by Lemma 3.4 and the preceding discussion. Once a coset is accepted, the b conditional coordinates are sampled in expected polynomial time by Lemma 2.5 and scaling. Combining these bounds with the expected number of iterations proves the claimed time and space bounds. $\square$

Remark 3.7 (Relation to prior Gaussian samplers). The quotient-class decomposition $\mathcal{S}_{\mathbf{h}} = \bigsqcup_{a \in \mathbb{Z}_P} q\mathbb{Z}^b + \frac{q}{P}a\mathbf{h}$ is an instance of the general lattice-extension sampling principle of Espitau, Wallet, and Yu [16]. It is also the rank-one specialization of the Construction-A codeword sampler of Bollauf, Lie, and Ling [8], for the one-dimensional code $\{a\mathbf{h} : a \in \mathbb{Z}_P\}$. Thus, the coset decomposition itself is prior art. In the present setting the acceptance probability is at least $1/(4P)$ rather than a constant, which is exactly sufficient for the $P \text{poly}(n, b, \log q, \log |s|, \log \|\mathbf{c}\|_\infty)$ cost per lift and the corresponding $P^2 \text{poly}(\cdot)$ Gaussian–Wagner budget in Algorithm 4.

3.2 The new instantiation of the DEL framework

The lattice chain. Following Ducas, Engelberts, and Loyer [15], we view the Gaussian–Wagner procedure as a walk through a sequence of projected lattices and auxiliary superlattices. The only difference is that we replace their coordinate auxiliary lattices by the random one-check superlattices introduced in the last subsection.

Let $\mathbf{A}' = [\mathbf{A} \mid \mathbf{I}_n] \in \mathbb{Z}_q^{n \times (m+n)}$, $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$. Fix positive integers $b_1, \dots, b_r$ satisfying $\sum_{i=1}^r b_i = n$, define $n_0 = 0$, and $n_i := \sum_{j=1}^i b_j$, $i \in [r]$. For each $i \in [r]$, let $\mathbf{A}_i$ denote the first n_i rows of $\mathbf{A}$, and set $\mathbf{A}'_i := [\mathbf{A}_i \mid \mathbf{I}_{n_i}] \in \mathbb{Z}_q^{n_i \times (m+n_i)}$. Define

$$\mathcal{L}_0 := \mathbb{Z}^m, \quad \mathcal{L}_i := \mathcal{L}_q^\perp(\mathbf{A}'_i) = \{(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^m \times \mathbb{Z}^{n_i} : \mathbf{A}_i \mathbf{x} + \mathbf{y} \equiv \mathbf{0} \pmod{q}\}. \quad (2)$$

Thus $\mathcal{L}_i$ is the lattice corresponding to the first n_i parity checks, and $\mathcal{L}_r = \mathcal{L}_q^\perp(\mathbf{A}')$.

For round i , let P_i be prime and sample $\mathbf{h}_i \leftarrow \mathbb{Z}_{P_i}^{b_i} \setminus \{\mathbf{0}\}$ uniformly. Define the one-check lattice on the newly introduced block by $\mathcal{S}_{\mathbf{h}_i} := q\mathbb{Z}^{b_i} + \frac{q}{P_i} \mathbf{h}_i \mathbb{Z}$ and embed it into the last b_i coordinates as $\widehat{\mathcal{S}}_i := \left\{ \begin{pmatrix} \mathbf{0} \\ \mathbf{z} \end{pmatrix} : \mathbf{z} \in \mathcal{S}_{\mathbf{h}_i} \right\} \subseteq \mathbb{R}^{m+n_{i-1}} \times \mathbb{R}^{b_i}$. The auxiliary superlattice used in round i is

$$\mathcal{L}'_i := \mathcal{L}_i + \widehat{\mathcal{S}}_i \subseteq \mathbb{R}^{m+n_i}. \quad (3)$$

Let $\pi_i : \mathbb{R}^{m+n_i} \rightarrow \mathbb{R}^{m+n_{i-1}}$ denote the orthogonal projection that deletes the last b_i coordinates. Then we have $\pi_i(\mathcal{L}'_i) = \mathcal{L}_{i-1}$, $\ker(\pi_i|_{\mathcal{L}'_i}) = \widehat{\mathcal{S}}_i$, $\mathcal{L}'_i/\mathcal{L}_i \simeq \mathcal{S}_{\mathbf{h}_i}/q\mathbb{Z}^{b_i} \simeq \mathbb{Z}_{P_i}$. In particular, $\widehat{\mathcal{S}}_i$ is primitive in $\mathcal{L}'_i$, and $\mathcal{L}'_i/\mathcal{L}_i$ has exactly P_i cosets.

The i th round first lifts every vector from $\mathcal{L}_{i-1}$ to $\mathcal{L}'_i$ along the primitive sublattice $\widehat{\mathcal{S}}_i$, and then combines pairs of lifted vectors lying in the same coset modulo $\mathcal{L}_i$. Their differences belong to $\mathcal{L}_i$. Hence the algorithm follows the zigzag

$$\mathbb{Z}^m = \mathcal{L}_0 \xleftarrow{\pi_1} \mathcal{L}'_1 \supseteq \mathcal{L}_1 \xleftarrow{\pi_2} \mathcal{L}'_2 \supseteq \dots \xleftarrow{\pi_r} \mathcal{L}'_r \supseteq \mathcal{L}_r = \mathcal{L}_q^\perp(\mathbf{A}').$$

discrete Gaussian lifting. Algorithm 2 lifts a vector from a projected lattice $\mathcal{Q}$ to the ambient lattice $\mathcal{L}'$ while preserving, up to a small multiplicative loss, the corresponding discrete Gaussian distribution. In round i of our Gaussian–Wagner sampler, it lifts vectors from $\mathcal{L}_{i-1}$ to $\mathcal{L}'_i$ along the primitive sublattice $\widehat{\mathcal{S}}_i$.

The procedure can be viewed as a blockwise reinterpretation of the GPV sampling algorithm [17]: rather than reducing a problem in dimension n to instances in dimensions 1 and $n-1$, the ambient space is decomposed into arbitrary complementary subspaces, see also [16, Algorithm 2]. For its analysis, we use the following result from [15, Lemma 3.1], itself a variant of [16, Theorem 1].

Lemma 3.8 (Analysis of DGLift [15]). *Let $\mathcal{L}' \subseteq \mathbb{R}^n$ be a lattice and let $\mathcal{Q} = \pi_{\mathcal{M}}^\perp(\mathcal{L}')$ for a primitive sublattice $\mathcal{M} \subseteq \mathcal{L}'$. Let $s > 0$ be a real such that a randomized algorithm $\mathcal{A}$ exists that, given $\mathbf{c} \in \text{span}_{\mathbb{R}}(\mathcal{M})$, returns a sample from $\mathcal{D}_{\mathcal{M}, s, \mathbf{c}}$. Then $\text{DGLift}(\mathcal{Q}, \mathcal{L}', s, \cdot)$ (Algorithm 2) is a randomized algorithm that, given a vector $\mathbf{x} \in \mathcal{Q}$, outputs a vector $\mathbf{x}' \in \mathcal{L}'$. It uses one query to $\mathcal{A}$, and all other operations run in polynomial time.*

Moreover, for any reals $\delta \geq 0$ and $0 < \varepsilon \leq \frac{1}{2}$ satisfying $s \geq \eta_\varepsilon(\mathcal{M})$, if $X_1, \dots, X_N \in \mathcal{Q}$ are conditionally δ -similar to independent samples from $\mathcal{D}_{\mathcal{Q}, s}$, then the distribution of $X'_1, \dots, X'_N$, where $X'_i := \text{DGLift}(\mathcal{Q}, \mathcal{L}', s, X_i)$ is conditionally $(\delta + 3\varepsilon)$ -similar to independent samples from $\mathcal{D}_{\mathcal{L}', s}$.

We will use the above lemma with $\mathcal{M} := \mathcal{S}_{\mathbf{h}}$ defined in Section 3.1, and the algorithm $\mathcal{A}$ will be implemented by Theorem 3.6. In round i , we apply the lemma with $\mathcal{M} = \widehat{\mathcal{S}}_i$, $\mathcal{Q} = \mathcal{L}_{i-1}$, $\mathcal{L}' = \mathcal{L}'_i$. Since $\widehat{\mathcal{S}}_i$ is isometric to $\mathcal{S}_{\mathbf{h}_i}$, the required shifted sampler is implemented by Theorem 3.6.

Algorithm 2 DGLift($\mathcal{Q}, \mathcal{L}', s, \mathbf{x}$)

Require: Lattices $\mathcal{Q}, \mathcal{L}'$ satisfying $\mathcal{Q} = \pi_{\mathcal{M}}^{\perp}(\mathcal{L}')$ for a primitive sublattice $\mathcal{M} \subseteq \mathcal{L}'$, $s > 0$, and $\mathbf{x} \in \mathcal{Q}$

Ensure: $\mathbf{x}' \in \mathcal{L}'$ such that $\pi_{\mathcal{M}}^{\perp}(\mathbf{x}') = \mathbf{x}$

1: Compute the complement $\mathcal{N}$ of $\mathcal{M}$

2: Compute the unique $\mathbf{y} \in \mathcal{N}$ satisfying $\pi_{\mathcal{M}}^{\perp}(\mathbf{y}) = \mathbf{x}$

▷ Lifting

3: Sample $\mathbf{z} \leftarrow D_{\mathcal{M}, s, \mathbf{x} - \mathbf{y}}$

▷ Sampling

4: **return** $\mathbf{x}' = \mathbf{y} + \mathbf{z}$

Combining to a Sublattice We next use a combining procedure that maps discrete Gaussian samples from a lattice $\mathcal{L}'$ to samples supported on a full rank sublattice $\mathcal{L} \subseteq \mathcal{L}'$ in Algorithm 3. The procedure partitions its input vectors according to their cosets modulo $\mathcal{L}$ and takes differences of appropriately paired vectors in the same coset. Such differences lie in $\mathcal{L}$, and the discrete Gaussian convolution lemma [27, 25] suggests that their distribution should be close to $D_{\mathcal{L}, \sqrt{2}s}$. Algorithm 3 is the combining procedure of [15, Algorithm 3] and has the following distributional guarantee from [15, Lemma 3.3].

Lemma 3.9 (Analysis of BucketAndCombine [15]). *Let $\mathcal{L} \subseteq \mathcal{L}'$ be full rank lattices in $\mathbb{R}^n$. Let $N \geq 3|\mathcal{L}'/\mathcal{L}|$ be a positive integer, and let $\delta \geq 0$, $0 < \varepsilon \leq \frac{1}{2}$, and $s \geq \sqrt{2}\eta_{\varepsilon}(\mathcal{L}')$ be reals. If the input list consists of N random variables on $\mathcal{L}'$ that are conditionally δ -similar to independent samples from $D_{\mathcal{L}', s}$, then Algorithm 3 returns a list of $\lfloor N/3 \rfloor$ vectors from $\mathcal{L}$ that are conditionally $(4\delta + 3\varepsilon)$ -similar to independent samples from $D_{\mathcal{L}, \sqrt{2}s}$.*

A New Instantiation of the Wagner Gaussian Sampler We are now ready to present our Gaussian sampler below. Applying the two procedures along the lattice chain gives Algorithm 4. Each round restores one coordinate block, increases the Gaussian width by a factor of $\sqrt{2}$, and retains one third of the list. Thus, r rounds transform $3^r N$ initial samples on $\mathbb{Z}^m$ into N vectors on the final lattice, provided the smoothing and quotient size conditions below hold.

Theorem 3.10 (Analysis of WagnerGaussianSampler). *Let $0 \leq \varepsilon \leq 1/2$ and $r \geq 1$ be an integer. If the input parameters satisfy $N \geq P_i$, $s_0 \geq \sqrt{\ln(2m+4)}/\pi$ and $2^{\frac{i-1}{2}} s_0 \geq \max(\eta_{\varepsilon}(\mathcal{S}_{\mathbf{h}_i}), \sqrt{2}\eta_{\varepsilon}(\mathcal{L}'_i))$ for all $i \in [r]$, then Algorithm 4 outputs a list of size N consisting of vectors that are conditionally $4^r 5\varepsilon$ -similar to independent samples from $D_{\mathcal{L}_q^{\perp}(\mathbf{A}'), 2^{\frac{r}{2}} s_0}$.*

The proof is similar to that of [15, Theorem 3.2] and is deferred to Appendix B.

Algorithm 3 BucketAndCombine($\mathcal{L}', \mathcal{L}, L$)

Require: Full rank lattices $\mathcal{L} \subseteq \mathcal{L}'$ and a list L with N vectors $\{\mathbf{x}_i\}_{i \in [N]}$ in $\mathcal{L}'$

Ensure: A list L_{out} with $\lfloor N/3 \rfloor$ vectors in $\mathcal{L}$

- 1: Initialize empty lists $L(\mathbf{c})$ for each coset $\mathbf{c} \in \mathcal{L}'/\mathcal{L}$
 - 2: **for** $i = 1, \dots, N$ **do** $\triangleright$ Bucketing
 - 3: Let $\mathbf{c}_i := \mathbf{x}_i \bmod \mathcal{L}$
 - 4: Append $\mathbf{x}_i$ to $L(\mathbf{c}_i)$
 - 5: Initialize an empty list L_{out}
 - 6: **for** $i = 1, \dots, N$ **do** $\triangleright$ Combining
 - 7: **if** $|L(\mathbf{c}_i)| \geq 2$ and $|L_{out}| < \lfloor N/3 \rfloor$ **then**
 - 8: Let $\mathbf{x}_1, \mathbf{x}_2$ be the first two elements in $L(\mathbf{c}_i)$
 - 9: Append $\mathbf{y} := \mathbf{x}_1 - \mathbf{x}_2$ to L_{out}
 - 10: Remove $\mathbf{x}_1$ and $\mathbf{x}_2$ from $L(\mathbf{c}_i)$
 - 11: **return** L_{out}
-

Algorithm 4 WagnerGaussianSampler($\mathbf{A}$)

Require: Integers n, m, q ;

Full rank matrix $\mathbf{A}' = [\mathbf{A} \mid \mathbf{I}_n] \in \mathbb{Z}_q^{n \times (m+n)}$;

Integer parameters $N, r, (P_i)_{i \in [r]}, (b_i)_{i \in [r]}$ with $\sum_{i \in [r]} b_i = n$;

Real parameter $s_0 > 0$.

Ensure: A list L_r with N vectors in $\mathcal{L}_q^\perp(\mathbf{A}')$

- 1: Let $\mathcal{L}_0 := \mathbb{Z}^m$.
 - 2: Initialize L_0 with $3^r N$ independent samples from $D_{\mathcal{L}_0, s_0}$.
 - 3: Sample independent $\mathbf{h}_i \leftarrow \mathbb{Z}_{P_i}^{b_i} \setminus \{\mathbf{0}\}$ and construct $\mathcal{L}_i$ and $\mathcal{L}'_i$
 - 4: **for** $i = 1, \dots, r$ **do**
 - 5: Let $\mathcal{L}_i$ be as defined in Equation (2).
 - 6: Let $\mathcal{L}'_i$ be as defined in Equation (3).
 - 7: $L'_{i-1} \leftarrow \emptyset$.
 - 8: **for** $\mathbf{x} \in L_{i-1}$ **do** $\triangleright$ Lift to $\mathcal{L}'_i$
 - 9: Sample $\mathbf{x}' \leftarrow \text{DGLift}(\mathcal{L}_{i-1}, \mathcal{L}'_i, s_{i-1}, \mathbf{x})$.
 - 10: Append $\mathbf{x}'$ to L'_{i-1} .
 - 11: $L_i \leftarrow \text{BucketAndCombine}(\mathcal{L}'_i, \mathcal{L}_i, L'_{i-1})$ $\triangleright$ Combine into $\mathcal{L}_i$
 - 12: $s_i \leftarrow \sqrt{2} s_{i-1}$.
 - 13: **return** L_r .
-

3.3 Putting it all together: a Gaussian–Wagner sampler with exponentially small ε

We now choose the parameters of Algorithm 4. There are two different smoothing constraints. The smoothing parameter of the auxiliary lattice $\mathcal{S}_{\mathbf{h}_i}$ determines how many coordinates can be treated in round i , and is therefore directly coupled with the quotient index P_i . By contrast, the smoothing parameters of the projected relation lattices $\mathcal{L}_i$ can be bounded uniformly over all prefixes and paid for once in the initial width s_0 . We first prove the latter statement using the Euclidean minimum of the dual lattice. This is the same counting argument as the one used for $\mathcal{S}_{\mathbf{h}}$, but applied to random systematic q -ary lattices.

Lemma 3.11. *Let $m, n \geq 1$, and let q be prime such that $q^{m/(m+n)} > 24$, $q > q^{m/(m+n)}\sqrt{m+n}$. Define $\gamma_0 := 1/(32\sqrt{2\pi e})$. Choose $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ uniformly and define $\mathbf{A}' := [\mathbf{A} \mid \mathbf{I}_n] \in \mathbb{Z}_q^{n \times (m+n)}$.*

Then $\Pr_{\mathbf{A}}[\lambda_1(\mathcal{L}_q^\perp(\mathbf{A}')^*) < \gamma_0\sqrt{m+n}q^{-n/(m+n)}] \leq 6^{-(m+n)}$.

Proof. We have $q\mathcal{L}_q^\perp(\mathbf{A}')^* = [\mathbf{A} \mid \mathbf{I}_n]^\top \mathbb{Z}_q^n + q\mathbb{Z}^{m+n}$. Set $r := \gamma_0\sqrt{m+n}q^{m/(m+n)}$.

The assumption $q > q^{m/(m+n)}\sqrt{m+n}$ implies $r < q$, so the ball $\mathcal{B}_{m+n}(r)$ contains no nonzero vector of $q\mathbb{Z}^{m+n}$. By Fact 3.1, we have $|\mathbb{Z}^{m+n} \cap \mathcal{B}_{m+n}(r)| \leq [\sqrt{2\pi e} \left(\frac{r}{\sqrt{m+n}} + \frac{1}{2} \right)]^{m+n} \leq [q^{\frac{m/(m+n)}{32}} + \frac{\sqrt{2\pi e}}{2}]^{m+n} \leq \frac{q^m}{6^{m+n}}$, where the last inequality uses $q^{m/(m+n)} > 24$.

Fix $\mathbf{z} = (\mathbf{x}, \mathbf{y})^\top \in \mathbb{Z}^m \times \mathbb{Z}^n$ that is nonzero modulo q . If $\mathbf{y} = \mathbf{0} \bmod q$, then $\mathbf{z} \notin q\mathcal{L}_q^\perp(\mathbf{A}')^*$ unless $\mathbf{x} = \mathbf{0} \bmod q$. Otherwise, $\mathbf{y} \bmod q$ determines a nonzero vector $\mathbf{u} \in \mathbb{Z}_q^n$, and $\Pr_{\mathbf{A}}[\mathbf{A}^\top \mathbf{u} = \mathbf{x} \bmod q] = q^{-m}$, because the m columns of $\mathbf{A}$ are independent and uniform. Therefore, $\Pr_{\mathbf{A}}[\mathbf{z} \in q\mathcal{L}_q^\perp(\mathbf{A}')^*] \leq q^{-m}$. A union bound over the nonzero integer points in $\mathcal{B}_{m+n}(r)$ gives $\Pr_{\mathbf{A}}[\lambda_1(q\mathcal{L}_q^\perp(\mathbf{A}')^*) < r] \leq \frac{q^m}{6^{m+n}} \cdot q^{-m} = 6^{-(m+n)}$. Dividing by q proves the claim. $\square$

Applying Lemma 2.9 gives the following smoothing bound.

Corollary 3.12. *Under the conditions of Lemma 3.11, except with probability $6^{-(m+n)}$, it holds that $\eta_\varepsilon(\mathcal{L}_q^\perp(\mathbf{A}')) \leq 32\sqrt{e}q^{n/(n+m)} \left(1 + \sqrt{\frac{2\ln(1+1/\varepsilon)}{m+n}} \right)$.*

Proof. Combining Lemma 3.11 and Lemma 2.9 gives the claimed bound. $\square$

For $i \in \{0, \dots, r\}$, recall the definition of $\mathcal{L}_i$ in Equation (2) and let $\Pi_i : \mathbb{R}^{m+n} \rightarrow \mathbb{R}^{m+n_i}$ be the coordinate projection deleting the last $n - n_i$ coordinates. Then $\Pi_i(\mathcal{L}_r) = \mathcal{L}_i$. Indeed, projecting a vector in $\mathcal{L}_r$ clearly gives a vector in $\mathcal{L}_i$. Conversely, if $(\mathbf{x}, \mathbf{y}) \in \mathcal{L}_i$, then one can choose the remaining $n - n_i$ coordinates so that the remaining equations are satisfied modulo q , thereby extending $(\mathbf{x}, \mathbf{y})$ to an element of $\mathcal{L}_r$. We therefore obtain a single simultaneous smoothing estimate for the whole chain.

Corollary 3.13. *Let $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ be uniform, and assume $q^{m/(m+n)} > 24$, $q > q^{m/(m+n)}\sqrt{m+n}$. Then, except with probability $6^{-(m+n)}$, simultaneously for every $i \in \{0, \dots, r\}$, $\eta_\varepsilon(\mathcal{L}_i) \leq 32\sqrt{e} q^{n/(m+n)} \left(1 + \sqrt{\frac{2 \ln(1+1/\varepsilon)}{m+n}}\right)$.*

Proof. Apply Corollary 3.12 to $\mathcal{L}_r = \mathcal{L}_q^\perp([\mathbf{A} \mid \mathbf{I}_n])$. This gives the stated bound for $\mathcal{L}_r$ except with probability $6^{-(m+n)}$. Equations $\Pi_i(\mathcal{L}_r) = \mathcal{L}_i$ and Lemma 2.10 then give the same bound for every $\mathcal{L}_i$. $\square$

This explains precisely why the smoothing cost of the projected relation lattices can be paid once. We choose s_0 above the single bound for $\mathcal{L}_r$. Since every $\mathcal{L}_i$ is a projection of $\mathcal{L}_r$ and the Gaussian width only increases from round to round, no additional round-dependent cost is incurred by the projected chain.

Choice of parameters. Fix constants $c_+, A, C_\varepsilon > 0$, and assume

$$1 \leq m \leq c_+ n, \quad q \leq n^A, \quad q > q^{m/(m+n)}\sqrt{m+n}, \quad \ln(1/\varepsilon) \leq C_\varepsilon n, \quad (4)$$

where q is prime, and $0 < \varepsilon \leq 1/2$. Define and assume that

$$\tau := \frac{m}{m+n} \ln q, \quad \Gamma := \frac{\tau}{\ln f}, \quad f \rightarrow \infty, \quad \Gamma \rightarrow \infty, \quad \ln \Gamma = o(f^2). \quad (5)$$

In particular, $\tau \rightarrow \infty$, and hence $q^{m/(m+n)} = e^\tau > 24$ for all sufficiently large n .

The parameter f specifies the target output width q/f . We choose $s_\star$ from the global smoothing bound, take the number of rounds r to be as large as this lower bound permits, and adjust s_0 so that the final width is exactly q/f . Hence set $s_\star := 128\sqrt{e} q^{n/(m+n)} \left(1 + \sqrt{\frac{2 \ln(1+3/\varepsilon)}{m+n}}\right)$, $r := \left\lfloor 2 \log \left(\frac{q}{fs_\star}\right) \right\rfloor$, $s_0 := \frac{q}{f2^{r/2}}$. For $i \in \{0, \dots, r\}$, define $s_i := 2^{i/2} s_0$. For $i \in [r]$, the Gaussian width used in round i is s_{i-1} , and $s_r = 2^{r/2} s_0 = \frac{q}{f}$.

Define $t_i := 2 \ln \left(\frac{q}{s_{i-1}}\right) + C_0$ for $i \in [r]$, $T := \sum_{i=1}^r \frac{1}{t_i}$, and $R := \lceil \exp \left(\frac{4n}{T}\right) \rceil$, where $C_0 > 0$ is a sufficiently large absolute constant. Choose a prime P satisfying $R \leq P \leq 2R$. Such a prime exists by Bertrand's theorem. Set $P_i := P$, $i \in [r]$, $N := 2P$.

Define the ideal block sizes by $\tilde{b}_i := \frac{n}{T t_i}$, $i \in [r]$. Since $\sum_{i=1}^r \tilde{b}_i = n$, set $\ell := n - \sum_{i=1}^r \lfloor \tilde{b}_i \rfloor \in \{0, \dots, r-1\}$, and define $b_i := \begin{cases} \lfloor \tilde{b}_i \rfloor + 1, & i \leq \ell, \\ \lfloor \tilde{b}_i \rfloor, & i > \ell. \end{cases}$ Then $\sum_{i=1}^r b_i = n$, $|b_i - \tilde{b}_i| \leq 1$. For later use, define $x_i := \frac{\ln P}{b_i}$, $i \in [r]$.

Theorem 3.14 (Gaussian–Wagner sampling with exponentially small error). *Assume (4) and (5), and choose $r, s_0, P, N, (P_i)_{i \in [r]}, (b_i)_{i \in [r]}$ as above. For all sufficiently large n , except with probability at most $\exp \left(-\Omega \left(\frac{n}{\tau \ln \Gamma}\right)\right)$ over uniform $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ and independent $\mathbf{h}_i \leftarrow \mathbb{Z}_P^{b_i} \setminus \{\mathbf{0}\}$, Algorithm 4 returns a list of $N = 2P$ vectors in $\mathcal{L}_r = \mathcal{L}_q^\perp([\mathbf{A} \mid \mathbf{I}_n])$ that are conditionally $\delta := 5 \cdot 4^r \varepsilon$ -similar to independent samples from $D_{\mathcal{L}_r, q/f}$.*

The expected running time and space are $P^2 \text{poly}(n, \log q)$ and $P \text{poly}(n, \log q)$, respectively, where $\ln P = (4 \ln 2 + o(1))n / \ln \Gamma$. In particular, taking $\varepsilon = e^{-cn}$ for any fixed $c > 0$ gives $\delta = e^{-(c-o(1))n}$.

We prove Theorem 3.14 in three steps. Lemma 3.15 establishes the basic parameter relations, Lemmas 3.16 and 3.17 verify the smoothing conditions, and the final argument combines Theorem 3.10 with the per-lift cost of Theorem 3.6.

Lemma 3.15 (Basic parameter relations). *For all sufficiently large n , $s_\star \leq s_0 < \sqrt{2}s_\star$, $s_r = q/f$, and $r = (2/\ln 2 + o(1))\tau$. Moreover, $t_{i+1} = t_i - \ln 2$ for $i \in [r-1]$, $t_1 = 2\tau + O(1)$, and $t_r = 2\ln f + O(1)$. Thus $T = (1/\ln 2 + o(1))\ln \Gamma$, $\ln P = (4 \ln 2 + o(1))\frac{n}{\ln \Gamma}$.*

Uniformly for every $i \in [r]$,

$$b_i = (1 + o(1))\frac{n}{Tt_i}, \text{ and } \min_{i \in [r]} b_i = \Omega\left(\frac{n}{\tau \ln \Gamma}\right) \rightarrow \infty. \quad (6)$$

Furthermore, $\ln P/b_i =: x_i \geq 2t_i = 4 \ln(q/s_{i-1}) + 2C_0$, thus $P > 24^{b_i}$, $i \in [r]$.

Proof. By $\ln(1/\varepsilon) \leq C_\varepsilon n$, the parenthesized factor in $s_\star$ is bounded above and below by positive constants. Hence $s_\star = \Theta(q^{n/(m+n)})$. It follows that $\ln\left(\frac{q}{fs_\star}\right) = \tau - \ln f + O(1)$. Since $\Gamma = \tau / \ln f \rightarrow \infty$, we have $\ln f = o(\tau)$, and hence $\ln\left(\frac{q}{fs_\star}\right) = (1 - o(1))\tau \rightarrow \infty$. This proves the bound on r .

By the definition of r , $2^{r/2} \leq \frac{q}{fs_\star} < 2^{(r+1)/2}$. Substituting $s_0 = q/(f2^{r/2})$ gives $s_\star \leq s_0 < \sqrt{2}s_\star$. The identity $s_r = q/f$ follows directly.

Since $s_0 = \Theta(s_\star)$ and $s_\star = \Theta(q^{n/(m+n)})$, we have $\ln\left(\frac{q}{s_0}\right) = \tau + O(1)$, and hence $t_1 = 2\tau + O(1)$. Also, $s_{r-1} = \frac{q}{\sqrt{2}f}$, so $t_r = 2\ln(\sqrt{2}f) + C_0 = 2\ln f + O(1)$. The recurrence $t_{i+1} = t_i - \ln 2$ follows from $s_i = \sqrt{2}s_{i-1}$.

Since $(t_i)_{i \in [r]}$ is an arithmetic progression with common difference $-\ln 2$, an integral comparison⁸ gives $T = \frac{1}{\ln 2} \ln\left(\frac{t_1}{t_r}\right) + O\left(\frac{1}{t_r}\right)$. By $t_1 = 2\tau + O(1)$, $t_r = 2\ln f + O(1)$, we have $\frac{t_1}{t_r} = (1 + o(1))\frac{\tau}{\ln f} = (1 + o(1))\Gamma$. Since $\Gamma \rightarrow \infty$, this proves the bound on T . The claimed estimate for P follows from $\ln P = \frac{4n}{T} + O(1)$.

The ideal block sizes sum to n . Since $t_i \leq t_1 = \Theta(\tau)$ and $T = \Theta(\ln \Gamma)$, $\min_i \tilde{b}_i = \Omega\left(\frac{n}{\tau \ln \Gamma}\right)$. Moreover, $q \leq n^A$ implies $\tau = O(\ln n)$ and $\ln \Gamma = O(\ln \ln n)$. Therefore, $\frac{n}{\tau \ln \Gamma} \rightarrow \infty$. The additive rounding error is consequently a uniform relative $o(1)$ error, proving (6).

For all sufficiently large n , $b_i \leq 2\tilde{b}_i = \frac{2n}{Tt_i}$. Since $\ln P \geq 4n/T$, $x_i = \frac{\ln P}{b_i} \geq \frac{4n/T}{2n/(Tt_i)} = 2t_i$. Finally, $t_i \geq t_r \rightarrow \infty$, and hence $x_i > \ln 24$ for every i and all sufficiently large n . Equivalently, $P > 24^{b_i}$. $\square$

⁸ Let $g(x) := (t_1 - (x-1)\ln 2)^{-1}$ for $x \in [1, r]$. Then g is increasing and $g(i) = 1/t_i$. Hence $g(i) \leq \int_i^{i+1} g(x) dx \leq g(i+1)$ for every $i \in [r-1]$. Summing gives $\frac{1}{\ln 2} \ln\left(\frac{t_1}{t_r}\right) + \frac{1}{t_1} \leq T \leq \frac{1}{\ln 2} \ln\left(\frac{t_1}{t_r}\right) + \frac{1}{t_r}$.

We use the following primitive case consequence of the modularity bound for the smoothing parameter; see [16, Section 3.1.2]. A proof is given in Appendix A.

Lemma 3.16. *For every $i \in [r]$, $\varepsilon \in (0, 1)$, $\eta_\varepsilon(\mathcal{L}'_i) \leq \max\{\eta_{\varepsilon/3}(\widehat{\mathcal{S}}_i), \eta_{\varepsilon/3}(\mathcal{L}_{i-1})\}$.*

We use the above lemma to give bounds on the smoothing parameters of all rounds.

Lemma 3.17. *Except with probability at most $6^{-(m+n)} + \sum_{i=1}^r 6^{-b_i} \leq \exp(-\Omega(\frac{n}{\tau \ln \Gamma}))$, every round $i \in [r]$ satisfies $s_{i-1} \geq \eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i})$, and $s_{i-1} \geq \sqrt{2}\eta_\varepsilon(\mathcal{L}'_i)$.*

Proof. By Corollary 3.13, except with probability $6^{-(m+n)}$, simultaneously for every $i \in \{0, \dots, r\}$, $\eta_{\varepsilon/3}(\mathcal{L}_i) \leq 32\sqrt{e}q^{n/d} \left(1 + \sqrt{\frac{2\ln(1+3/\varepsilon)}{m+n}}\right) = \frac{s_*}{4}$. Since $s_{i-1} \geq s_0 \geq s_*$, $\eta_{\varepsilon/3}(\mathcal{L}_{i-1}) \leq \frac{s_{i-1}}{4}$, $i \in [r]$.

By Corollary 3.3, $\eta_{\varepsilon/3}(\mathcal{S}_{\mathbf{h}_i}) \leq 32\sqrt{e}qe^{-x_i} \left(1 + \sqrt{\frac{2\ln(1+3/\varepsilon)}{b_i}}\right)$, except with probability 6^{-b_i} . The conditions of that corollary follow from Lemma 3.15.

By $\ln(1/\varepsilon) \leq C_\varepsilon n$ and $b_i = (1 + o(1))\frac{n}{Tt_i}$, we have $\frac{\ln(1+3/\varepsilon)}{b_i} = O(Tt_i)$. Furthermore, $x_i \geq 2t_i$, and $\frac{q}{s_{i-1}} = \exp(\frac{t_i - C_0}{2})$. Dividing the above bound of $\eta_{\varepsilon/3}(\mathcal{S}_{\mathbf{h}_i})$ by s_{i-1} gives $\eta_{\varepsilon/3}(\mathcal{S}_{\mathbf{h}_i})/s_{i-1} \leq O(e^{-3t_i/2 - C_0/2} (1 + \sqrt{Tt_i}))$. For sufficiently large t , the right-hand side is decreasing in t . Since $t_i \geq t_r = 2\ln f + O(1)$, it is uniformly bounded by $O(\frac{1 + \sqrt{\ln \Gamma \ln f}}{f^3})$. The condition $\ln \Gamma = o(f^2)$ implies $\frac{1 + \sqrt{\ln \Gamma \ln f}}{f^3} = o(1)$. Hence, for all sufficiently large n , $\eta_{\varepsilon/3}(\mathcal{S}_{\mathbf{h}_i}) \leq \frac{s_{i-1}}{4}$, $i \in [r]$.

Since $\widehat{\mathcal{S}}_i$ is isometric to $\mathcal{S}_{\mathbf{h}_i}$, Lemma 3.16, $\eta_{\varepsilon/3}(\mathcal{L}_{i-1}) \leq \frac{s_{i-1}}{4}$, and $\eta_{\varepsilon/3}(\mathcal{S}_{\mathbf{h}_i}) \leq \frac{s_{i-1}}{4}$, $i \in [r]$ give $\eta_\varepsilon(\mathcal{L}'_i) \leq \frac{s_{i-1}}{4}$, $i \in [r]$. Thus, $\eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i}) \leq \frac{s_{i-1}}{4}$. This proves the two smoothing conditions.

Finally, by $b_i = (1 + o(1))\frac{n}{Tt_i}$, $\sum_{i=1}^r 6^{-b_i} \leq r \exp(-\Omega(\frac{n}{\tau \ln \Gamma}))$. Since $r = O(\ln n)$ and $n/(\tau \ln \Gamma) \gg \ln r$, the polynomial factor is absorbed into the exponent, proving the claimed bound. $\square$

With the smoothing conditions established, we can now apply Theorem 3.10. The remaining argument tracks the output width and list size and bounds the total cost of the Gaussian lifts.

Proof (Proof of Theorem 3.14). By Lemma 3.15, we have $s_0 \geq s_*$. The assumption $q > q^{m/(m+n)}\sqrt{m+n}$ is equivalent to $q^{n/(m+n)} > \sqrt{m+n}$. It follows from the definition of s_* that $s_0 \geq \sqrt{\frac{\ln(2m+4)}{\pi}}$ for all sufficiently large n . Hence the initial list can be filled with exact independent samples from $D_{\mathbb{Z}^m, s_0}$.

For every round $i \in [r]$, the quotient $\mathcal{L}'_i/\mathcal{L}_i$ has exactly $P_i = P$ cosets. Since $N = 2P > P_i$, the list entering round i has size $3^{r-i+1}N \geq 3P_i$. Thus Algorithm 3 returns exactly one third as many vectors in every round, and the final list has size N .

On the event of Lemma 3.17, for every $i \in [r]$, we have $2^{(i-1)/2}s_0 = s_{i-1} \geq \max\{\eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i}), \sqrt{2}\eta_\varepsilon(\mathcal{L}'_i)\}$. Therefore, all the smoothing conditions of Theorem 3.10

hold. The output list is conditionally $5 \cdot 4^r \varepsilon$ -similar to independent samples from $D_{\mathcal{L}_r, 2^{r/2} s_0}$. By the definition of s_0 , $2^{r/2} s_0 = \frac{q}{f}$, which proves the distributional claim.

It remains to analyze the complexity. For every $i \in [r]$, $0 < s_{i-1} \leq s_{r-1} = \frac{q}{\sqrt{2}f} < q$, and $\frac{q}{s_{i-1}} \leq \frac{q}{s_0} = e^{\tau+O(1)} \leq n^{O(1)}$, where the last inequality follows from $q \leq n^A$. Hence the sampler of Theorem 3.6 implements in expected time $P \text{poly}(n, \log q)$ for every $i \in [r]$.

The complement maps used by DGLift have polynomial bit complexity. Moreover, after the combining step, the difference of two lifted vectors lies in the integral lattice $\mathcal{L}_i$, so denominators introduced by the superlattices do not multiply from one round to the next. Since $r = O(\tau) = O(\ln n)$, all vectors and all centers supplied to the Gaussian samplers have polynomial bit-length.

The total number of calls to DGLift is $\sum_{i=1}^r |L_{i-1}| = O(3^r N)$. Since $r = O(\ln n)$ and $N = 2P$, the total expected running time is $P^2 \text{poly}(n, \log q)$. The largest stored list contains $3^r N = P \text{poly}(n)$ vectors, and the coset-weight computation of Theorem 3.6 is streaming. Hence the expected space is $P \text{poly}(n, \log q)$.

By Lemma 3.15, $T = (\frac{1}{\ln 2} + o(1)) \ln \Gamma$. Since $R = \lceil e^{4n/T} \rceil$ and $R \leq P \leq 2R$, we obtain $\ln P = \frac{4n}{T} + O(1) = (4 \ln 2 + o(1)) \frac{n}{\ln T}$. A suitable prime $P \in [R, 2R]$ can be found by enumerating candidates and applying a polynomial-time primality test; its cost is at most $P \text{poly}(n)$ and is absorbed by the preceding running-time bound.

Furthermore, $r = O(\tau) = O(\ln n)$. Therefore, for every fixed $c > 0$, we have $5 \cdot 4^r e^{-cn} = e^{-cn+O(\ln n)} = e^{-(c-o(1))n}$. $\square$

Remark 3.18 (Limitation of the parameterization). For comparison, outside the asymptotic regime of Equation (5), suppose that $\tau = O(1)$ while retaining the same construction and its sufficient smoothing conditions. By the definition of s_* and the assumption $\ln(1/\varepsilon) = O(n)$, $s_* = \Theta(q^{n/(m+n)})$. Consequently, we have $\frac{q}{s_*} = \Theta(q^{m/(m+n)}) = e^{O(1)}$. Since $s_0 \geq s_*$ and $f > 1$, every valid parameter choice has only $r = O(1)$ Wagner rounds. Because $\sum_{i=1}^r b_i = n$, there exists some $i \in [r]$ such that $b_i = \Omega(n)$. The smoothing analysis of $\mathcal{S}_{\mathbf{h}_i}$ in Lemma 3.2 requires $P > 24^{b_i}$ and hence $\ln P > b_i \ln 24 = \Omega(n)$. Thus the resulting running time is exponential. This proves that the condition $\tau \rightarrow \infty$, equivalently $m = \omega(n/\ln q)$ in the sublinear regime, is asymptotically sharp for the present parameterization.

4 Application to sublinear sample Binary LWE

As mentioned in Section 1, our algorithm does not amplify the input LWE samples. Instead, it generates a list of short vectors which are conditionally similar to the Gaussian distribution over $\mathcal{L}_q^\perp(\mathbf{A}')$. The right-hand side of the LWE instance is used only in the final Fourier test. Therefore, the conditional noise-independence problem arising in sample amplification does not occur here.

The short vectors generated in Section 3 are not independent. For completeness, we use the conditional-expectation consequence of Lemma 2.3. In the uniform case, all Fourier phases depend on the same uniform vector, so we first

exclude zero and signed residue collisions and then apply exact character orthogonality.

4.1 Pointwise LWE and its relation lattice

Let q be an odd prime, let $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, and fix integer vectors $\mathbf{s} \in \mathbb{Z}^n$, $\mathbf{e} \in \mathbb{Z}^m$. The pointwise LWE distribution associated with $(\mathbf{s}, \mathbf{e})$ is $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e} \bmod q$. The uniform distribution has $\mathbf{b} \leftarrow \mathbb{Z}_q^m$ uniform and independent of $\mathbf{A}$. And the pointwise narrow LWE additionally needs $\|\mathbf{s}\|^2 + \|\mathbf{e}\|^2 \leq O(m+n)$, see Definition 1.1.

For a matrix $\mathbf{A}' = [-\mathbf{A} \mid \mathbf{I}_n] \in \mathbb{Z}_q^{n \times (m+n)}$, define its q -ary lattice $\mathcal{L}_q^\perp(\mathbf{A}') = \left\{ \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \in \mathbb{Z}^{m+n} : -\mathbf{A}\mathbf{u} + \mathbf{v} \equiv \mathbf{0} \bmod q \right\}$. The matrix $\mathbf{A}'$ is systematic and has rank n for every $\mathbf{A}$. Since $-\mathbf{A}$ is uniform whenever $\mathbf{A}$ is uniform, Theorem 3.14 applies directly to $\mathbf{A}'$. The reason for using the systematic form is the following identity.

Lemma 4.1. *Let $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e}$. For every $\mathbf{y} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \in \mathcal{L}_q^\perp(\mathbf{A}')$, we have $\langle \mathbf{b}, \mathbf{u} \rangle = \langle \begin{pmatrix} \mathbf{e} \\ \mathbf{s} \end{pmatrix}, \mathbf{y} \rangle \bmod q$.*

Proof. Since $\mathbf{y} \in \mathcal{L}_q^\perp(\mathbf{A}')$, we have $\mathbf{v} = \mathbf{A}\mathbf{u} \bmod q$. Therefore, $\langle \mathbf{b}, \mathbf{u} \rangle = \langle \mathbf{A}^\top \mathbf{s} + \mathbf{e}, \mathbf{u} \rangle = \langle \mathbf{s}, \mathbf{A}\mathbf{u} \rangle + \langle \mathbf{e}, \mathbf{u} \rangle = \langle \mathbf{s}, \mathbf{v} \rangle + \langle \mathbf{e}, \mathbf{u} \rangle \bmod q$. $\square$

4.2 Instantiation of the Gaussian–Wagner sampler

We use the following parameterization of Algorithm 4. Set $\tau := \frac{m}{m+n} \ln q$, $f := \tau$, $\varepsilon_0 := e^{-n}$, and $w := \frac{q}{f}$. We work in the regime $\tau \rightarrow \infty$. Thus we have $\Gamma = \frac{\tau}{\ln f} = \frac{\tau}{\ln \tau}$. Hence $\Gamma \rightarrow \infty$ and $\ln \Gamma = (1 + o(1)) \ln \tau$, so the conditions of Theorem 3.14 are satisfied.

Let P and r be the parameters selected in Section 3.3, and write $N_{\text{gw}} := 2P$ and $\delta_{\text{gw}} := 5 \cdot 4^r \varepsilon_0$. Moreover, by Lemma 3.15, we have $\ln P = (4 \ln 2 + o(1)) \frac{n}{\ln \tau}$ and $r = (\frac{2}{\ln 2} + o(1))\tau$. Since $q \leq n^{O(1)}$, we have $\tau = O(\ln n)$, and therefore $\delta_{\text{gw}} = e^{-n+O(\tau)} = e^{-(1-o(1))n}$.

By Theorem 3.14, except with probability $\xi_n := \exp(-\Omega(\frac{n}{\tau \ln \tau}))$, Algorithm 4, applied to $\mathbf{A}'$, returns N_{gw} vectors in $\mathcal{L}_q^\perp(\mathbf{A}')$ that are conditionally δ_{gw} -similar to independent samples from $D_{\mathcal{L}_q^\perp(\mathbf{A}'), w}$. On the same good event, Corollary 3.13 and the parameter selection in Section 3.3 give $w \geq \eta_{\varepsilon_0}(\mathcal{L}_q^\perp(\mathbf{A}'))$. In the remainder, let $\mathcal{G}$ denote this good event.

4.3 Fourier bias of an ideal relation

The distinguisher uses the first Fourier coefficient of the relation distribution. This is the Fourier statistic underlying the dual test of Aharonov and Regev [2] and the final bias test in BKW-type algorithms [22,21].

Lemma 4.2. Let $\mathcal{L} \subseteq \mathbb{Z}^n$ be a full-rank lattice, let $\mathbf{x} \in \mathbb{Z}^n$, and let $\mathbf{y} \leftarrow D_{\mathcal{L}, w}$. Then $\mathbb{E} \exp(\frac{2\pi i}{q} \langle \mathbf{x}, \mathbf{y} \rangle) = \frac{\sum_{\mathbf{z} \in \mathcal{L}^*} e^{-\pi w^2 \|\mathbf{z} - \mathbf{x}/q\|^2}}{\sum_{\mathbf{z} \in \mathcal{L}^*} e^{-\pi w^2 \|\mathbf{z}\|^2}}$. In particular, this quantity is real and nonnegative. If $w \geq \eta_\varepsilon(\mathcal{L})$, then $\mathbb{E} \cos(\frac{2\pi}{q} \langle \mathbf{x}, \mathbf{y} \rangle) \geq \frac{e^{-\pi w^2 \|\mathbf{x}\|^2/q^2}}{1+\varepsilon}$.

Proof. Apply Poisson summation to $\mathbf{y} \mapsto e^{-\pi \|\mathbf{y}\|^2/w^2} e^{2\pi i \langle \mathbf{x}/q, \mathbf{y} \rangle}$ over $\mathcal{L}$. The common factor $w^n/\det(\mathcal{L})$ cancels between the numerator and denominator, giving the first equation. The inequality retains the $\mathbf{z} = \mathbf{0}$ term in the numerator and uses $\rho_{1/w}(\mathcal{L}^*) \leq 1 + \varepsilon$ in the denominator. $\square$

Choice of parameters for LWE. Now we set the choices of parameters for the LWE distinguisher.

Set $\mathbf{x} := \begin{pmatrix} \mathbf{e} \\ \mathbf{s} \end{pmatrix} \in \mathbb{Z}^{m+n}$ and suppose that $\|\mathbf{e}\|^2 + \|\mathbf{s}\|^2 \leq B^2(m+n)$ for a fixed constant $B \geq 1$. By $w \geq \eta_{\varepsilon_0}(\mathcal{L}_q^\perp(\mathbf{A}'))$ and Lemma 4.2, the ideal Fourier bias satisfies

$$\beta_F := \mathbb{E}_{\mathbf{y} \leftarrow D_{\mathcal{L}_q^\perp(\mathbf{A}'), w}} \cos(\frac{2\pi}{q} \langle \mathbf{x}, \mathbf{y} \rangle) \geq \frac{e^{-\pi B^2(m+n)/\tau^2}}{1 + \varepsilon_0}. \quad (7)$$

Define

$$\mu := \frac{1}{4} e^{-\pi B^2(m+n)/\tau^2}, \quad M := \lceil n^{20} \mu^{-2} \rceil. \quad (8)$$

Since $m+n = O(n)$ and $\tau := \frac{m}{m+n} \ln q \rightarrow \infty$, we have $\ln M = O(\ln n + \frac{n}{\tau^2}) = o(\frac{n}{\ln \tau}) = o(\ln P)$. Hence $M < P < N_{\text{gw}}$, for all sufficiently large n . Furthermore, $\delta_{\text{gw}} = o(\mu)$, because $\ln(1/\delta_{\text{gw}}) = (1 - o(1))n$ whereas $\ln(1/\mu) = O(n/\tau^2) = o(n)$.

4.4 The LWE distinguisher

Algorithm 5 GaussianWagnerLWE($\mathbf{A}, \mathbf{b}, B$)

Require: $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{b} \in \mathbb{Z}_q^m$, and a norm bound $B \geq 1$

Ensure: One of LWE and Uniform

- 1: Set $\tau \leftarrow \frac{m}{m+n} \ln q$, $f \leftarrow \tau$, $\varepsilon_0 \leftarrow e^{-n}$, and $\mathbf{A}' \leftarrow [-\mathbf{A} \mid \mathbf{I}_n]$.
 - 2: Set $\mu \leftarrow \frac{1}{4} e^{-\pi B^2(m+n)/\tau^2}$ and $M \leftarrow \lceil n^{20} \mu^{-2} \rceil$.
 - 3: Run Algorithm 4 on $\mathbf{A}'$ using the parameters of Section 4.2.
 - 4: Retain its first M outputs $\mathbf{y}_i = \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix} \in \mathcal{L}_q^\perp(\mathbf{A}')$, $i \in [M]$.
 - 5: **for** $i = 1, \dots, M$ **do**
 - 6: Compute a certified approximation $\tilde{Z}_i$ satisfying $|\tilde{Z}_i - \cos(\frac{2\pi}{q} \langle \mathbf{b}, \mathbf{u}_i \rangle)| \leq \frac{\mu}{10}$.
 - 7: Set $\tilde{T} \leftarrow \frac{1}{M} \sum_{i=1}^M \tilde{Z}_i$.
 - 8: **if** $\tilde{T} \geq 2\mu$ **then**
 - 9: **return** LWE.
 - 10: **else**
 - 11: **return** Uniform.
-

The cosine evaluations are carried out using directed interval arithmetic. Since $\ln(1/\mu) = O(\frac{n}{\tau^2})$, the numerical part takes $M \text{poly}(n, \log q)$ bit operations and is asymptotically dominated by the Gaussian Wagner sampler.

4.5 Completeness

Lemma 4.3 (Completeness on the good event). *Fix $\mathbf{s}, \mathbf{e}$ satisfying $\|\mathbf{e}\|^2 + \|\mathbf{s}\|^2 \leq B^2(m+n)$, and suppose $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e} \bmod q$. Conditioned on $\mathcal{G}$, Algorithm 5 outputs Uniform with probability $\exp(-\Omega(n^{20}))$.*

Proof. For $i \in [M]$, set $Z_i := \cos(\frac{2\pi}{q} \langle \mathbf{b}, \mathbf{u}_i \rangle)$, $T := \frac{1}{M} \sum_{i=1}^M Z_i$. By Lemma 4.1, $Z_i = \cos(\frac{2\pi}{q} \langle \mathbf{x}, \mathbf{y}_i \rangle)$.

Conditioned on $\mathcal{G}$, the retained sublist remains conditionally δ_{gw} -similar to independent samples from $D_{\mathcal{L}_q^\perp(\mathbf{A}'), w}$ by Lemma 2.3. Let $\mathcal{F}_{i-1} := \sigma(Z_1, \dots, Z_{i-1})$. By Lemmas 2.3 and Lemma 4.2, $\mathbb{E}[Z_i \mid \mathcal{F}_{i-1}] \geq \beta_F - 2(e^{\delta_{\text{gw}}} - 1)$, and Equations (7) and (8) give $\beta_F \geq \frac{4\mu}{1+\varepsilon_0}$.

Together with $\delta_{\text{gw}} = o(\mu)$, this implies, for all sufficiently large n , $\mathbb{E}[Z_i \mid \mathcal{F}_{i-1}] \geq 3\mu$, $i \in [M]$.

Let $D_i := Z_i - \mathbb{E}[Z_i \mid \mathcal{F}_{i-1}]$, $S_k := \sum_{i=1}^k D_i$. Then D_i is $\mathcal{F}_i$ -measurable and $\mathbb{E}[D_i \mid \mathcal{F}_{i-1}] = 0$, so $(S_k, \mathcal{F}_k)_{k=0}^M$ is a martingale. Moreover, $Z_i \in [-1, 1]$, and therefore $\mathbb{E}[Z_i \mid \mathcal{F}_{i-1}] \in [-1, 1]$, $|D_i| \leq 2$.

If Algorithm 5 outputs Uniform, then the certified approximation guarantee implies $T < \frac{21}{10}\mu$. By $\mathbb{E}[Z_i \mid \mathcal{F}_{i-1}] \geq 3\mu$, this event implies $S_M = MT - \sum_{i=1}^M \mathbb{E}[Z_i \mid \mathcal{F}_{i-1}] < -\frac{9}{10}M\mu$. The one-sided Azuma–Hoeffding inequality therefore gives $\Pr[T < \frac{21}{10}\mu] \leq \exp(-\frac{(9M\mu/10)^2}{8M}) = \exp(-\Omega(M\mu^2)) = \exp(-\Omega(n^{20}))$. $\square$

4.6 Soundness

All Fourier phases use the same uniform vector $\mathbf{b}$, so they are not independent. We first exclude residue collisions in the short vector list and then use character orthogonality. For a residue class $\mathbf{c} \in \mathcal{L}_q^\perp(\mathbf{A}')/q\mathbb{Z}^{m+n}$, define $p_{\mathbf{c}} := \Pr_{\mathbf{y} \leftarrow D_{\mathcal{L}_q^\perp(\mathbf{A}'), w}}[\mathbf{y} \in \mathbf{c} + q\mathbb{Z}^{m+n}]$ and $p_{\max} := \max_{\mathbf{c}} p_{\mathbf{c}}$.

Lemma 4.4. *We have $p_{\max} \leq \exp(-(1 - o(1))(m+n)\tau)$.*

Proof. Since $\mathbf{A}'$ has rank n , $\det(\mathcal{L}_q^\perp(\mathbf{A}')) = q^n$. For any representative $\mathbf{c} \in \mathcal{L}_q^\perp(\mathbf{A}')$, the standard shifted-mass inequality gives $\rho_w(\mathbf{c} + q\mathbb{Z}^{m+n}) \leq \rho_w(q\mathbb{Z}^{m+n}) = \rho_{1/f}(\mathbb{Z}^{m+n})$ by $w = q/f$. Since $f = \tau \rightarrow \infty$, it holds that $\rho_{1/f}(\mathbb{Z}) = 1 + 2\sum_{z \geq 1} e^{-\pi f^2 z^2} = 1 + O(e^{-\pi f^2})$ and hence $\rho_{1/f}(\mathbb{Z}^{m+n}) \leq \exp(O((m+n)e^{-\pi f^2}))$.

On the other hand, by Poisson summation, it holds that $\rho_w(\mathcal{L}_q^\perp(\mathbf{A}')) = \frac{w^{m+n}}{q^n} \rho_{1/w}(\mathcal{L}_q^\perp(\mathbf{A}')^*) \geq \frac{w^d}{q^n}$. It gives $p_{\max} \leq \rho_{1/f}(\mathbb{Z}^{m+n}) \frac{q^n}{w^{m+n}} = \rho_{1/f}(\mathbb{Z}^{m+n}) \frac{f^d}{q^m}$. By definition, $m \ln q = (m+n)\tau$. Therefore, $\ln p_{\max} \leq O((m+n)e^{-\pi \tau^2}) + (m+n) \ln \tau - (m+n)\tau = -(1 - o(1))(m+n)\tau$. $\square$

Let $\mathcal{C}$ denote the event that the retained vectors are nonzero and pairwise distinct up to sign modulo q : $\forall i, j, \mathbf{y}_i \not\equiv \mathbf{0} \pmod{q}$ and $\mathbf{y}_i \not\equiv \pm \mathbf{y}_j \pmod{q}$ ($i \neq j$).

Lemma 4.5. *Conditioned on $\mathcal{G}$, we have $\Pr[-\mathcal{C}] \leq e^{\delta_{\text{gw}}} M p_{\max} (1 + 2e^{\delta_{\text{gw}}} M) = \exp(-\Omega((m+n)\tau))$.*

Proof. For one ideal sample from $D_{\mathcal{L}_q^\perp(\mathbf{A}'), w}$, the probability of the zero residue is at most $p_{\max}$. For two independent ideal samples, $\Pr[\mathbf{y}_i \equiv \mathbf{y}_j \pmod{q\mathbb{Z}^{m+n}}] = \sum_{\mathbf{c}} p_{\mathbf{c}}^2 \leq p_{\max}$, and similarly, $\Pr[\mathbf{y}_i \equiv -\mathbf{y}_j \pmod{q\mathbb{Z}^{m+n}}] = \sum_{\mathbf{c}} p_{\mathbf{c}} p_{-\mathbf{c}} \leq p_{\max}$. Lemma 2.3 increases the one-coordinate bound by at most $e^{\delta_{\text{gw}}}$ and each two-coordinate bound by at most $e^{2\delta_{\text{gw}}}$. A union bound proves the first inequality.

Finally, $\ln M = O(\ln n + \frac{n}{\tau^2}) = o((m+n)\tau)$, which proves the claimed asymptotic bound. $\square$

Lemma 4.6. *Fix $\mathbf{A}$ and a short vector list satisfying $\mathcal{C} = \{ \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix} \in \mathcal{L}_q^\perp(\mathbf{A}') \}_{i \in [M]}$.*

Let $\mathbf{b} \leftarrow \mathbb{Z}_q^m$ be uniform. For $i \in [M]$, set $Z_i := \cos(\frac{2\pi}{q} \langle \mathbf{b}, \mathbf{u}_i \rangle)$, $T := \frac{1}{M} \sum_{i=1}^M Z_i$. Then $\mathbb{E}_{\mathbf{b}}[Z_i] = 0$, $\mathbb{E}_{\mathbf{b}}[Z_i^2] = \frac{1}{2}$, $\mathbb{E}_{\mathbf{b}}[Z_i Z_j] = 0$ ($i \neq j$). Consequently, $\mathbb{E}_{\mathbf{b}}[T] = 0$, $\text{Var}_{\mathbf{b}}(T) = \frac{1}{2M}$.

Proof. By the definition of $\mathbf{A}$, the event $\mathcal{C}$ implies $\mathbf{u}_i \neq \mathbf{0}$, $\mathbf{u}_i \neq \pm \mathbf{u}_j$ in $\mathbb{Z}_q^m$. Every nontrivial additive character of $\mathbb{Z}_q^m$ has mean zero, hence $\mathbb{E}_{\mathbf{b}}[Z_i] = 0$. Since q is odd, $2\mathbf{u}_i \neq \mathbf{0}$, and $Z_i^2 = \frac{1}{2} + \frac{1}{2} \cos(\frac{4\pi}{q} \langle \mathbf{b}, \mathbf{u}_i \rangle)$ has expectation $1/2$. Finally, $Z_i Z_j = \frac{1}{2} \cos(\frac{2\pi}{q} \langle \mathbf{b}, \mathbf{u}_i - \mathbf{u}_j \rangle) + \frac{1}{2} \cos(\frac{2\pi}{q} \langle \mathbf{b}, \mathbf{u}_i + \mathbf{u}_j \rangle)$, so $\mathbb{E}_{\mathbf{b}}[Z_i Z_j] = 0$ ($i \neq j$). $\square$

Lemma 4.7. *Suppose $\mathbf{b} \leftarrow \mathbb{Z}_q^m$ is uniform and independent of $\mathbf{A}$. Conditioned on $\mathcal{G}$, Algorithm 5 outputs LWE with probability $O(n^{-20}) + \exp(-\Omega((m+n)\tau))$.*

Proof. Condition on $\mathcal{C}$ and on the entire short vector list. If the algorithm outputs LWE , the certified approximation guarantee implies $T \geq \frac{19}{10}\mu$. By Lemma 4.6 and Lemma 2.11, $\Pr_{\mathbf{b}}[T \geq \frac{19}{10}\mu] \leq \frac{1}{2M(19\mu/10)^2} = O(n^{-20})$. The bound is uniform over every fixed short vector list satisfying $\mathcal{C}$. Averaging over the list and adding the collision probability from Lemma 4.5 proves the claim. $\square$

4.7 The decision theorem

Theorem 4.8. *Fix constants $B \geq 1$, $c_+ > 0$, $A > 0$. For all sufficiently large n , let $m = m(n)$ and an odd prime $q = q(n)$ satisfy*

$$1 \leq m \leq c_+ n, \quad q \leq n^A, \quad q > q^{m/(m+n)} \sqrt{m+n}. \quad (9)$$

Assume that $\tau := \frac{m}{m+n} \ln q \rightarrow \infty$.

There is a randomized algorithm solving narrow-LWE $_{n,m,q,B}$ with completeness and soundness errors $O(n^{-10})$. Its expected running time is $\exp((8 \ln 2 + o(1)) \frac{n}{\ln \tau})$ and its expected space is $\exp((4 \ln 2 + o(1)) \frac{n}{\ln \tau})$.

Proof. Run Algorithm 5. In the LWE case, Lemma 4.3 gives error $\exp(-\Omega(n^{20}))$ conditioned on $\mathcal{G}$. In the uniform case, Lemma 4.7 gives an error $O(n^{-20}) + \exp(-\Omega((m+n)\tau))$ conditioned on $\mathcal{G}$.

The failure probability of $\mathcal{G}$ is $\xi_n = \exp(-\Omega(\frac{n}{\tau \ln \tau}))$. Since $\tau \leq \ln q \leq A \ln n$, $\frac{n}{\tau \ln \tau} \geq \Omega(\frac{n}{\ln n \ln \ln n}) = \omega(\ln n)$. Hence $\xi_n = n^{-\omega(1)}$, and adding it to the two conditional error bounds gives completeness and soundness errors $O(n^{-10})$.

By Theorem 3.14, the Gaussian–Wagner sampler has the expected time $P^2 \text{poly}(n, \log q)$ and the expected space $P \text{poly}(n, \log q)$.

The Fourier test uses only $M < P$ outputs and costs $M \text{poly}(n, \log q)$ additional bit operations. $\ln P = (4 \ln 2 + o(1)) \frac{n}{\ln \tau}$ by the parameterization in Section 4.2. Therefore, it gives the claimed time and space bounds.

It remains to verify the specialization. Suppose $q = n^{a+o(1)}$, $m = o(n)$, and $m = (n/\ln n)g(n)$. Then $\frac{m}{m+n} = (1 + o(1)) \frac{m}{n}$, so $\tau = \frac{m}{m+n} \ln q = (a + o(1))g(n)$. Moreover, $\frac{q}{q^{m/(m+n)} \sqrt{m+n}} = n^{a-1/2-o(1)} \rightarrow \infty$, where we use $a > 1/2$ and $g(n) = o(\ln n)$. Thus (9) holds for all sufficiently large n . Substituting $\ln \tau = (1 + o(1)) \ln g(n)$ into the general complexity bound proves the stated specialization.

Finally, for $m = n/\ln \ln n$, $g(n) = \frac{\ln n}{\ln \ln n}$, $\ln g(n) = (1 + o(1)) \ln \ln n$, which proves the last assertion. $\square$

Remark 4.9 (Sample–time tradeoff). Suppose that $q = n^{a+o(1)}$ is an odd prime for a fixed constant $a > 1/2$, and let $m = (n/\ln n)g(n)$, where $g(n) \rightarrow \infty$ and $g(n) = o(\ln n)$. Then $m = o(n)$, $\tau = (a + o(1))g(n)$, and $\frac{q}{q^{m/(m+n)} \sqrt{m+n}} = n^{a-1/2-o(1)} \rightarrow \infty$. Taking any fixed $A > a$, Equation (9) holds for all sufficiently large n , and Theorem 4.8 gives expected running time $\exp\left((8 \ln 2 + o(1)) \frac{n}{\ln g(n)}\right)$. This is subexponential whenever $m = \omega(n/\ln n)$ and $m = o(n)$. In particular, for $m = \lceil n/\ln \ln n \rceil$, $g(n) = (1 + o(1)) \ln n/\ln \ln n$ and $\ln g(n) = (1 + o(1)) \ln \ln n$, giving expected running time $\exp\left((8 \ln 2 + o(1)) \frac{n}{\ln \ln n}\right)$.

References

1. Aggarwal, D., Li, Z., Stephens-Davidowitz, N.: A $2^{n/2}$ -time algorithm for $\sqrt{n}$ -SVP and $\sqrt{n}$ -hermite SVP, and an improved time-approximation tradeoff for (H) SVP. In: Canteaut, A., Standaert, F.X. (eds.) *Advances in Cryptology – EUROCRYPT 2021, Part I*. Lecture Notes in Computer Science, vol. 12696, pp. 467–497. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77870-5_17
2. Aharonov, D., Regev, O.: Lattice problems in $\text{NP} \cap \text{coNP}$. *Journal of the ACM* **52**(5), 749–765 (2005). <https://doi.org/10.1145/1089023.1089025>
3. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pp. 99–108. ACM (1996). <https://doi.org/10.1145/237814.237838>
4. Albrecht, M.R.: On dual lattice attacks against small-secret LWE and parameter choices in HELib and SEAL. In: *Advances in Cryptology – EUROCRYPT 2017, Part II*. Lecture Notes in Computer Science, vol. 10211, pp. 103–129. Springer (2017). https://doi.org/10.1007/978-3-319-56614-6_4, <https://eprint.iacr.org/2017/047>

5. Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: Automata, Languages and Programming – ICALP 2011, Part I. Lecture Notes in Computer Science, vol. 6755, pp. 403–415. Springer (2011). https://doi.org/10.1007/978-3-642-22006-7_34, <https://eccc.weizmann.ac.il/report/2010/066/>
6. Bai, S., Galbraith, S.D.: Lattice decoding attacks on binary LWE. In: Information Security and Privacy – ACISP 2014. Lecture Notes in Computer Science, vol. 8544, pp. 322–337. Springer (2014). https://doi.org/10.1007/978-3-319-08344-5_21, <https://eprint.iacr.org/2013/839>
7. Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM* **50**(4), 506–519 (2003). <https://doi.org/10.1145/792538.792543>
8. Bollauf, M.F., Lie, M., Ling, C.: On gaussian sampling for q -ary lattices and linear codes with lee weight. In: Advances in Cryptology – CRYPTO 2025, Part I. Lecture Notes in Computer Science, vol. 16000, pp. 321–352. Springer (2025). https://doi.org/10.1007/978-3-032-01855-7_11, full version: Cryptology ePrint Archive, Report 2025/087
9. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC). pp. 575–584. ACM (2013). <https://doi.org/10.1145/2488608.2488680>
10. Brakerski, Z., Vaikuntanathan, V.: Efficient Fully Homomorphic Encryption from (Standard) LWE . *SIAM Journal on Computing* **43**(2), 831–871 (Jan 2014). <https://doi.org/10.1137/120868669>
11. Buchmann, J., Göpfert, F., Player, R., Wunderer, T.: On the hardness of LWE with binary error: Revisiting the hybrid lattice-reduction and meet-in-the-middle attack. In: Progress in Cryptology – AFRICACRYPT 2016. Lecture Notes in Computer Science, vol. 9646, pp. 24–43. Springer (2016). https://doi.org/10.1007/978-3-319-31517-1_2, <https://eprint.iacr.org/2016/089>
12. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: Fast Fully Homomorphic Encryption Over the Torus. *Journal of Cryptology* **33**(1), 34–91 (Jan 2020). <https://doi.org/10.1007/s00145-019-09319-x>
13. Choudhuri, A.R., Jain, A., Jin, Z.: SNARGs for P from LWE. In: 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS). pp. 68–79. IEEE (2022). <https://doi.org/10.1109/FOCS52979.2021.00016>, <https://eprint.iacr.org/2021/808>
14. Dadush, D., Regev, O., Stephens-Davidowitz, N.: On the closest vector problem with a distance guarantee. In: 2014 IEEE 29th Conference on Computational Complexity (CCC). pp. 98–109 (2014). <https://doi.org/10.1109/CCC.2014.18>, full version
- 15.ucas, L., Engelberts, L., Loyer, J.: Wagner’s algorithm provably runs in subexponential time for sis^∞ . In: Advances in Cryptology – CRYPTO 2025, Part I. Lecture Notes in Computer Science, vol. 16000, pp. 353–384. Springer (2025). https://doi.org/10.1007/978-3-032-01855-7_12, full version: Cryptology ePrint Archive, Report 2025/575
16. Espitau, T., Wallet, A., Yu, Y.: On gaussian sampling, smoothing parameter and application to signatures. Cryptology ePrint Archive, Report 2023/1654 (2023), major revision dated January 2026
17. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC). pp. 197–206. ACM (2008). <https://doi.org/10.1145/1374376.1374407>

18. Glabush, L., Longa, P., Naehrig, M., Peikert, C., Stebila, D., Virdia, F.: FrodoKEM: A CCA-secure learning with errors key encapsulation mechanism. *IACR Communications in Cryptology* **2**(3) (2025). <https://doi.org/10.62056/ayivom2hd>
19. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Attribute-based encryption for circuits. In: *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. pp. 545–554. ACM (2013). <https://doi.org/10.1145/2488608.2488677>, <https://eprint.iacr.org/2013/337>
20. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Predicate encryption for circuits from LWE. In: *Advances in Cryptology – CRYPTO 2015, Part II. Lecture Notes in Computer Science*, vol. 9216, pp. 503–523. Springer (2015). https://doi.org/10.1007/978-3-662-48000-7_25, <https://eprint.iacr.org/2015/029>
21. Herold, G., Kirshanova, E., May, A.: On the asymptotic complexity of solving LWE. *Designs, Codes and Cryptography* **86**(1), 55–83 (2018). <https://doi.org/10.1007/s10623-016-0326-0>
22. Kirchner, P., Fouque, P.A.: An improved BKW algorithm for LWE with applications to cryptography and lattices. In: *Advances in Cryptology – CRYPTO 2015, Part I. Lecture Notes in Computer Science*, vol. 9215, pp. 43–62. Springer (2015). https://doi.org/10.1007/978-3-662-47989-6_3, full version: *Cryptology ePrint Archive*, Report 2015/552
23. Lyubashevsky, V.: The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In: *APPROX-RANDOM 2005. Lecture Notes in Computer Science*, vol. 3624, pp. 378–389. Springer (2005). https://doi.org/10.1007/11538462_32
24. Micciancio, D.: On the hardness of learning with errors with binary secrets. *Theory of Computing* **14**(13), 1–17 (2018). <https://doi.org/10.4086/toc.2018.v014a013>, <https://theoryofcomputing.org/articles/v014a013/>
25. Micciancio, D., Peikert, C.: Hardness of SIS and LWE with small parameters. In: *Advances in Cryptology – CRYPTO 2013, Part I. Lecture Notes in Computer Science*, vol. 8042, pp. 21–39. Springer (2013). https://doi.org/10.1007/978-3-642-40041-4_2
26. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing* **37**(1), 267–302 (2007). <https://doi.org/10.1137/S0097539705447360>
27. Peikert, C.: Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity* **17**(2), 300–351 (May 2008). <https://doi.org/10.1007/s00037-008-0251-3>
28. Peikert, C.: A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science* **10**(4), 283–424 (2016). <https://doi.org/10.1561/0400000074>, <https://eprint.iacr.org/2015/939>
29. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: *Advances in Cryptology – CRYPTO 2019, Part I. Lecture Notes in Computer Science*, vol. 11692, pp. 89–114. Springer (2019). https://doi.org/10.1007/978-3-030-26948-7_4, <https://eprint.iacr.org/2019/158>
30. Peikert, C., Vaikuntanathan, V., Waters, B.: A Framework for Efficient and Composable Oblivious Transfer. In: Wagner, D. (ed.) *Advances in Cryptology – CRYPTO 2008*. pp. 554–571. Springer, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_31
31. Pouly, A., Shen, Y.: Provable dual attacks on learning with errors. In: *Advances in Cryptology – EUROCRYPT 2024, Part VII. Lecture Notes in Computer Science*, vol. 14657, pp. 256–285. Springer (2024). https://doi.org/10.1007/978-3-031-58174-5_15

- 1007/978-3-031-58754-2_10, updated full version: Cryptology ePrint Archive, Report 2023/1508
32. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC). pp. 84–93. ACM (2005). <https://doi.org/10.1145/1060590.1060603>
 33. Regev, O.: The learning with errors problem (invited survey). In: 2010 IEEE 25th Annual Conference on Computational Complexity. pp. 191–204. IEEE (2010). <https://doi.org/10.1109/CCC.2010.26>, <https://cims.nyu.edu/~regev/papers/lwesurvey.pdf>
 34. Sun, C., Tibouchi, M., Abe, M.: Revisiting the hardness of binary error LWE. In: Information Security and Privacy – ACISP 2020. Lecture Notes in Computer Science, vol. 12248, pp. 425–444. Springer (2020). https://doi.org/10.1007/978-3-030-55304-3_22, <https://eprint.iacr.org/2020/666>
 35. Wagner, D.: A generalized birthday problem. In: Advances in Cryptology – CRYPTO 2002. Lecture Notes in Computer Science, vol. 2442, pp. 288–304. Springer (2002). https://doi.org/10.1007/3-540-45708-9_19

A Proof of modularity bound for the smoothing parameter

Lemma A.1 (Smoothing under a primitive extension). *Let $\mathcal{M}$ be a primitive sublattice of a lattice $\mathcal{L}'$, and let $\mathcal{Q} := \pi_{\mathcal{M}}^{\perp}(\mathcal{L}')$. Then, for every $0 < \varepsilon \leq 1$, $\eta_{\varepsilon}(\mathcal{L}') \leq \max\{\eta_{\varepsilon/3}(\mathcal{M}), \eta_{\varepsilon/3}(\mathcal{Q})\}$. In particular, for every $i \in [r]$, $\eta_{\varepsilon}(\mathcal{L}'_i) \leq \max\{\eta_{\varepsilon/3}(\widehat{\mathcal{S}}_i), \eta_{\varepsilon/3}(\mathcal{L}_{i-1})\}$.*

Proof. Let $u \geq \max\{\eta_{\varepsilon/3}(\mathcal{M}), \eta_{\varepsilon/3}(\mathcal{Q})\}$. Since $\mathcal{M}$ is primitive in $\mathcal{L}'$, restriction of integer linear functionals from $\mathcal{L}'$ to $\mathcal{M}$ is surjective. Equivalently, $\pi_{\mathcal{M}}((\mathcal{L}')^*) = \mathcal{M}^*$, $\ker(\pi_{\mathcal{M}}|_{(\mathcal{L}')^*}) = \mathcal{Q}^*$.

For every $\mathbf{y} \in \mathcal{M}^*$, choose $\tilde{\mathbf{y}} \in (\mathcal{L}')^*$ satisfying $\pi_{\mathcal{M}}(\tilde{\mathbf{y}}) = \mathbf{y}$, and write $\tilde{\mathbf{y}} = \mathbf{y} + \mathbf{c}_{\mathbf{y}}$, $\mathbf{c}_{\mathbf{y}} \in \text{span}_{\mathbb{R}}(\mathcal{Q})$. The fiber of $\pi_{\mathcal{M}}$ over $\mathbf{y}$ is $\tilde{\mathbf{y}} + \mathcal{Q}^*$. By orthogonality, $\rho_{\frac{1}{u}}(\tilde{\mathbf{y}} + \mathcal{Q}^*) = \rho_{\frac{1}{u}}(\mathbf{y})\rho_{\frac{1}{u}}(\mathcal{Q}^* + \mathbf{c}_{\mathbf{y}})$. The Gaussian mass of a lattice coset is at most that of the centered lattice, so $\rho_{\frac{1}{u}}(\mathcal{Q}^* + \mathbf{c}_{\mathbf{y}}) \leq \rho_{\frac{1}{u}}(\mathcal{Q}^*)$. Summing over $\mathbf{y} \in \mathcal{M}^*$ gives $\rho_{\frac{1}{u}}((\mathcal{L}')^*) \leq \rho_{\frac{1}{u}}(\mathcal{M}^*)\rho_{\frac{1}{u}}(\mathcal{Q}^*)$. By the definition of u , $\rho_{\frac{1}{u}}(\mathcal{M}^*) \leq 1 + \frac{\varepsilon}{3}$, $\rho_{\frac{1}{u}}(\mathcal{Q}^*) \leq 1 + \frac{\varepsilon}{3}$. Consequently, $\rho_{\frac{1}{u}}((\mathcal{L}')^* \setminus \{\mathbf{0}\}) \leq (1 + \frac{\varepsilon}{3})^2 - 1 = \frac{2\varepsilon}{3} + \frac{\varepsilon^2}{9} \leq \varepsilon$, where the last inequality uses $\varepsilon \leq 1$. This proves the first claim. The second follows from (3). $\square$

B Correctness of the Gaussian Wagner sampler

We first analyze one iteration of Algorithm 4. This is the analogue of [15, Theorem 3.1] for the one-check auxiliary lattice.

Theorem B.1 (Analysis of one round of WagnerGaussianSampler). *Let $\delta \geq 0$, let $0 < \varepsilon \leq 1/2$, and fix a round $i \in [r]$. Assume that exact sampling from $D_{\widehat{\mathcal{S}}_i, s_{i-1}, \mathbf{c}}$ is available for every center $\mathbf{c}$ used by Algorithm 2. Suppose that $s_{i-1} \geq$*

$\max \{ \eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i}), \sqrt{2}\eta_\varepsilon(\mathcal{L}'_i) \}$, and that L_{i-1} consists of $N \geq 3P_i$ random variables in $\mathcal{L}_{i-1}$ that are conditionally δ -similar to independent samples from $D_{\mathcal{L}_{i-1}, s_{i-1}}$. Then round i of Algorithm 4 outputs a list L_i of $\lfloor N/3 \rfloor$ vectors in $\mathcal{L}_i$ that are conditionally $4\delta + 15\varepsilon$ -similar to independent samples from $D_{\mathcal{L}_i, s_i}$, where $s_i = \sqrt{2}s_{i-1}$.

Proof. By (3), $\widehat{\mathcal{S}}_i$ is primitive in $\mathcal{L}'_i$ and $\pi_{\widehat{\mathcal{S}}_i}^\perp(\mathcal{L}'_i) = \mathcal{L}_{i-1}$. Moreover, $\widehat{\mathcal{S}}_i$ is isometric to $\mathcal{S}_{\mathbf{h}_i}$, and therefore $\eta_\varepsilon(\widehat{\mathcal{S}}_i) = \eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i})$. The condition $s_{i-1} \geq \eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i})$ therefore allows us to apply Lemma 3.8. After lifting all elements of L_{i-1} , the list L'_{i-1} consists of N random variables in $\mathcal{L}'_i$ that are conditionally $\delta + 3\varepsilon$ -similar to independent samples from $D_{\mathcal{L}'_i, s_{i-1}}$.

By $|\mathcal{L}'_i/\mathcal{L}_i| = P_i$, $N \geq 3P_i$ guarantees that Algorithm 3 outputs $\lfloor N/3 \rfloor$ vectors. Since $s_{i-1} \geq \sqrt{2}\eta_\varepsilon(\mathcal{L}'_i)$, Lemma 3.9, applied with input similarity parameter $\delta + 3\varepsilon$, shows that these vectors are conditionally $4(\delta + 3\varepsilon) + 3\varepsilon = 4\delta + 15\varepsilon$ -similar to independent samples from $D_{\mathcal{L}_i, \sqrt{2}s_{i-1}}$. $\square$

Proof (Proof of Theorem 3.10). The initial list L_0 consists of $3^r N$ exact independent samples from $D_{\mathcal{L}_0, s_0}$, and is therefore conditionally 0-similar to independent samples from this distribution.

For $i \in \{0, \dots, r\}$, define $\delta_i := 5(4^i - 1)\varepsilon$. We prove by induction on i that L_i consists of exactly $3^{r-i}N$ vectors in $\mathcal{L}_i$ that are conditionally δ_i -similar to independent samples from $D_{\mathcal{L}_i, s_i}$, where $s_i = 2^{i/2}s_0$.

The claim for $i = 0$ follows from the initialization, since $\delta_0 = 0$. Now fix $i \in [r]$ and assume that the claim holds for $i - 1$. The input list to round i has size $|L_{i-1}| = 3^{r-i+1}N \geq 3N \geq 3P_i$. Moreover, $s_{i-1} = 2^{(i-1)/2}s_0 \geq \max \{ \eta_\varepsilon(\mathcal{S}_{\mathbf{h}_i}), \sqrt{2}\eta_\varepsilon(\mathcal{L}'_i) \}$. Thus Theorem B.1, applied with $\delta = \delta_{i-1}$, shows that L_i consists of $\frac{|L_{i-1}|}{3} = 3^{r-i}N$ vectors that are conditionally $4\delta_{i-1} + 15\varepsilon$ -similar to independent samples from $D_{\mathcal{L}_i, s_i}$. Finally, $4\delta_{i-1} + 15\varepsilon = 4 \cdot 5(4^{i-1} - 1)\varepsilon + 15\varepsilon = \delta_i$. This proves the induction claim.

Taking $i = r$, the final list has size N and consists of vectors in $\mathcal{L}_r = \mathcal{L}_q^\perp(\mathbf{A}')$ that are conditionally $\delta_r = 5(4^r - 1)\varepsilon \leq 5 \cdot 4^r \varepsilon$ -similar to independent samples from $D_{\mathcal{L}_r, s_r}$, where $s_r = 2^{r/2}s_0$. $\square$

C Proof for calculating $\Theta(\mu, \sigma)$

Proof (Proof of Lemma 3.4). Reduce μ modulo $\mathbb{Z}$ to $\bar{\mu} \in [-1/2, 1/2]$, and set $J := t + 3$.

First assume that $0 < \sigma < 1$. Factoring out the term corresponding to the nearest integer gives $\Theta(\bar{\mu}, \sigma) = \exp\left(-\pi \frac{\bar{\mu}^2}{\sigma^2}\right) \left(\sum_{z \in \mathbb{Z}} \exp\left(-\pi \frac{z^2 - 2z\bar{\mu}}{\sigma^2}\right) \right)$. Since $|\bar{\mu}| \leq 1/2$, for every $z \in \mathbb{Z}$, $z^2 - 2z\bar{\mu} \geq |z|(|z| - 1)$. The parenthesized sum is therefore between 1 and $1 + 2 \sum_{k \geq 1} e^{-\pi k(k-1)} < 4$, and the contribution of the terms with $|z| \geq J$ is at most $2 \sum_{k \geq J} \exp\left(-\pi \frac{k(k-1)}{\sigma^2}\right) \leq 4 \exp\left(-\pi \frac{J(J-1)}{\sigma^2}\right) < 2^{-t-4}$. Thus, the truncated sum together with the certified tail bound gives a certified approximation to the logarithm. More explicitly, $\log \Theta(\bar{\mu}, \sigma) = -\frac{\pi \bar{\mu}^2}{\sigma^2 \ln 2} +$

$\log \left(\sum_{z \in \mathbb{Z}} \exp \left(-\pi \frac{z^2 - 2z\bar{\mu}}{\sigma^2} \right) \right)$. The second term is the logarithm of a number in $[1, 4]$. Since $\forall x > 0, \ln(1+x) < x$, no precision is lost from taking the logarithm.

Now assume that $\sigma \geq 1$. By the Poisson summation formula, $\Theta(\bar{\mu}, \sigma) = \sigma \left(1 + 2 \sum_{k \geq 1} \exp(-\pi \sigma^2 k^2) \cos(2\pi k \bar{\mu}) \right)$. The expression in parentheses is bounded away from zero, since $2 \sum_{k \geq 1} e^{-\pi \sigma^2 k^2} \leq 2 \sum_{k \geq 1} e^{-\pi k^2} < \frac{1}{10}$. Hence it belongs to $(\frac{9}{10}, \frac{11}{10})$. Moreover, the contribution of the terms with $k \geq J$ is at most $2 \sum_{k \geq J} e^{-\pi \sigma^2 k^2} \leq 4e^{-\pi \sigma^2 J^2} < 2^{-t-4}$. It follows that the truncated Poisson series gives a certified approximation to $\log \Theta(\bar{\mu}, \sigma)$.

In either case, only $O(t)$ terms are evaluated. Standard directed interval algorithms for π , exponentiation, cosine, and logarithms, with polynomially many guard bits, therefore give an interval of width at most 2^{-t} in $\text{poly}(\log |\sigma|, t, \log |\mu|)$ time and space. $\square$