Structural Lattice Reduction:
Generalized Worst-Case to Average-Case Reductions

Nicolas Gama [*] Malika Izabachene [†] Phong Q. Nguyen [‡] Xiang Xie [§]

April 23, 2014

Abstract

In lattice cryptography, worst-case to average-case reductions rely on two problems: Ajtai's SIS and Regev's LWE, which refer to a very small class of random lattices related to the group $G = \mathbb{Z}_q^n$. We generalize worst-case to average-case reductions to (almost) all integer lattices, by allowing G to be any (sufficiently large) finite abelian group. In particular, we obtain a partition of the set of full-rank integer lattices of large volume such that finding short vectors in a lattice chosen uniformly at random from any of the partition cells is as hard as finding short vectors in any integer lattice. Our main tool is a novel group generalization of lattice reduction, which we call structural lattice reduction: given a finite abelian group G and a lattice L, it finds a short basis of some lattice $\bar{L}$ such that $L \subseteq \bar{L}$ and $\bar{L}/L \simeq G$. Our group generalizations of SIS and LWE allow us to abstract lattice cryptography, yet preserve worst-case assumptions.

1 Introduction

A lattice is a discrete subgroup of $\mathbb{R}^m$, *e.g.* a subgroup of $\mathbb{Z}^m$. Nearly two decades after its introduction, lattice-based cryptography has emerged as a credible alternative to classical public-key cryptography based on factoring or discrete logarithm. It offers new properties (such as security based on worst-case assumptions) and new functionalities, such as noisy multilinear maps and fully-homomorphic encryption. The worst-case guarantees of lattice-based cryptography come from two major problems: the *short integer solution* (SIS) problem dating back to Ajtai's breakthrough work at STOC '96 [1], and the *learning with errors* (LWE) problem introduced by Regev at STOC '05 [31], and somewhat related to the Ajtai-Dwork cryptosystem [2]. These two average-case problems are provably as hard as solving certain lattice problems in the worst case, such as GapSVP (the decision version of the shortest vector problem in a lattice) and SIVP (finding short linearly independent lattice vectors).

The SIS problem can be defined as finding short (nonzero) vectors in a random lattice from a class $\mathcal{A}_{n,m,q}$ of m-dimensional integer lattices related to the finite abelian (homocyclic) group $G = \mathbb{Z}_q^n$, where n is the dimension of the worst-case lattice problem and q needs to be sufficiently large: any sequence $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$ chosen uniformly at random defines a lattice $\mathcal{L}_{\mathbf{g}} \in \mathcal{A}_{n,m,q}$ formed by all $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{Z}^m$ s.t. $\sum_{i=1}^m x_i g_i = 0$ in G; and SIS asks, given $\mathbf{g}$, to find a short (nonzero) $\mathbf{x} \in \mathcal{L}_{\mathbf{g}}$. The class $\mathcal{A}_{n,m,q}$ has an algebraic meaning: the distribution of $\mathcal{L}_{\mathbf{g}}$ turns out to be statistically close (for sufficiently large m) to the uniform distribution over the finite set $\mathcal{L}_{G,m}$ of all full-rank lattices $L \subseteq \mathbb{Z}^m$ such that $\mathbb{Z}^m/L \simeq G$. This suggests that Ajtai's lattices are very rare among all integer lattices: in fact, Nguyen and Shparlinski [24] recently showed that the set $\cup_{G \text{ cyclic}} \mathcal{L}_{G,m}$ of all full-rank integer lattices $L \subseteq \mathbb{Z}^m$ such that $\mathbb{Z}^m/L$ is cyclic (unlike $\mathbb{Z}_q^n$) has natural density $1/[\zeta(6) \prod_{k=4}^m \zeta(k)] \approx 85\%$ (for large m), which implies that Ajtai's classes $\mathcal{A}_{n,m,q}$ form a minority of lattices among all integer lattices.

This motivates the natural question of whether other classes of random lattices enjoy similar worst-case to average-case reductions: in particular, if we call GSIS the generalization of SIS to any finite abelian group G, does GSIS have similar properties as SIS for other finite abelian groups G than $G = \mathbb{Z}_q^n$? This would imply that the random lattices of the class $\mathcal{L}_{G,m}$ are also hard. Ajtai (in the proceedings version of [1]) and later Regev [30] noticed that the choice $G = \prod_{i=1}^n \mathbb{Z}_{q_i}$ where the q_i's are distinct prime numbers of similar bit-length also worked. Micciancio [18] showed that another special choice of G also worked: there, G is actually constructed by an algorithm [18, Lemma 2.11] given as input a very special

[*] UVSQ and CNRS, France

[†] EPF, France

[‡] INRIA, France and Tsinghua University, China

[§] Institute of Software, Chinese Academy of Sciences, China lattice (for which solving the closest vector problem is easy); if the input lattice is $\mathbb{Z}^n$, then $G = (\mathbb{Z}_q)^n$. However, all these choices of G are arguably very special, and it was unclear if the hardness properties held outside a small family of finite abelian groups.

A similar question can be asked for LWE, which is known as a dual problem of SIS, and has been used extensively in lattice-based encryption. However, in order to define GLWE by analogy with GSIS, we need to change the usual definition of LWE based on linear algebra. Any finite abelian group G is isomorphic to its dual group $\hat{G}$ formed by its characters, *i.e.* homomorphisms from G to the torus $\mathbb{R}/\mathbb{Z}$. We define search-GLWE as the problem of learning a character $\hat{s} \in \hat{G}$ chosen uniformly at random, given noisy evaluations of $\hat{s}$ at (public) random points $g_1, \ldots, g_m \in G$, namely one is given g_i and a "Gaussian" perturbation of $\hat{s}(g_i)$ for all $1 \leq i \leq m$. By analogy with LWE, decisional-GLWE is defined as the problem of distinguishing the previous "Gaussian" perturbations of $\hat{s}(g_i)$ from random elements in $\mathbb{R}/\mathbb{Z}$. If $G = (\mathbb{Z}_q)^n$, it can be checked that GLWE is simply LWE. If $G = \mathbb{Z}_p$ where p is a large prime number, then search-GLWE is a randomized version of Boneh-Venkatesan's *Hidden Number Problem* [8] (introduced to study the bit-security of Diffie-Hellman key exchange, but also used in side-channel attacks on discrete-log based signatures [23]), which asks to recover a secret number $s \in \mathbb{Z}_p$, given $t_1, \ldots, t_m$ chosen uniformly at random from $\mathbb{Z}_p$ and approximations of $st_i \bmod p$ for each $1 \leq i \leq m$. Here, randomized means that the approximations given are "Gaussian" perturbations of $st_i \bmod p$. Thus, GLWE allows to view LWE and the Hidden Number Problem as a single problem, instantiated with different groups. Alternatively, GLWE can be viewed as a lattice problem: solving a randomized version of bounded distance decoding (with "Gaussian" errors) for the dual lattice of $\mathcal{L}_{\mathbf{g}}$.

Our results. We show that the worst-case to average-case reductions for SIS and LWE (search and decisional) can be generalized to GSIS and GLWE, provided that G is any sufficiently large finite abelian group, *e.g.* of order $n^{\Omega(\max(n, \mathrm{rank}(G)))}$ if n is the dimension of the worst-case lattice problem and rank (G) denotes the minimal size of a generating set for G. For GSIS and search-GLWE, our reductions are direct from worst-case lattice problems. On the other hand, we transfer all hardness results on decisional-LWE to decisional-GLWE, by reducing decisional-LWE to decisional-GLWE (under similar size constraints on G): we do so by generalizing the modulus-dimension switching technique of Brakerski *et al.* [9] (which was inspired by previous work on fully-homomorphic encryption).

Our reductions are based on a new tool, which we call structural lattice reduction, and which might be of independent interest: Becker *et al.* [7] recently used it to design new exponential-space algorithms for lattice problems. In lattice reduction, one is given a full-rank lattice $L \subseteq \mathbb{Z}^n$ and wants to find a short basis of L. In our structural lattice reduction, one is further given a finite abelian group G of rank $\leq n$, and wants to find a short basis of some overlattice $\bar{L}$ of L such that $\bar{L}/L \simeq G$ effectively, *i.e.* the map φ in the short exact sequence $0 \to L \xrightarrow{\mathrm{id}} \bar{L} \xrightarrow{\varphi} G \to 0$ is efficiently computable. Our starting point is that previous worst-case to average-case reductions (*e.g.* [14, 9]) implicitly used a trivial case of structural lattice reduction: if B is a short basis of a full-rank lattice $L \subseteq \mathbb{Z}^n$ and q is an integer, then $q^{-1}B$ is a short basis of the lattice $\bar{L} = q^{-1}L$ such that $\bar{L}/L \simeq \mathbb{Z}_q^n$, which explains the importance of $\mathbb{Z}_q^n$ in SIS and LWE.

Our GSIS reduction shows that in some sense all integer lattices are hard. Indeed, the set of full-rank lattices $L \subseteq \mathbb{Z}^m$ (of sufficiently large co-volume $\geq n^{\Omega(m)}$) can be partitioned based on the finite abelian group $\mathbb{Z}^m/L$, and the reduction implies that each partition cell $\mathcal{L}_{G,m}$ has this worst-case to average-case property: finding short vectors in a lattice chosen uniformly at random from $\mathcal{L}_{G,m}$ is as hard as finding short vectors in any integer lattice of dimension n.

Consider the special case $G = \mathbb{Z}_p$ for a large prime p. Then our GSIS reduction provides the first hardness results for the random lattices in $\mathcal{L}_{\mathbb{Z}_p,m}$ used in many experiments [12, 10] to benchmark lattice reduction algorithms, as well as in Darmstadt's SVP internet challenges. And our GLWE reduction provides a general hardness result for the hidden number problem: previously, [9, Cor 3.4] established the hardness for the hidden number problem when the large prime p is replaced by q^n where q is smooth.

Finally, our generalizations of SIS and LWE allow us to abstract (the many) lattice-based schemes based on SIS and/or LWE, where the role of $G = (\mathbb{Z}_q)^n$ was not very explicit in most descriptions (typically based on linear algebra). We believe such an abstraction can have several benefits. First, it can clarify analyses and designs: the El Gamal cryptosystem is arguably better described with an arbitrary group G, rather than by focusing on the historical choice $G = \mathbb{Z}_p^*$; comparisons and analogies with "traditional" public-key cryptography based on factoring or discrete logarithm will be easier. Second, it opens up the possibility of obtaining more efficient schemes using different choices of G than $G = (\mathbb{Z}_q)^n$.

We do not claim that there are better choices than $G = (\mathbb{Z}_q)^n$, but such a topic is worth investigating, which we leave to future work. Many factors influence efficiency: trapdoor generation, hashing, efficiency of the security reduction, *etc.* For instance, hashing onto $\mathbb{Z}_p$ can sometimes be more efficient than onto $(\mathbb{Z}_q)^n$ for large n, which could be useful in certain settings, like digital signatures.

Related work. Baumslag *et al.* also introduced in [6] group generalizations of LWE, targeting non-commutative groups, but did not obtain any hardness result.

Open problems. The recent reduction of Brakerski *et al.* [9] proves the hardness of decisional-LWE for a wide range of parameters, without establishing a direct search-to-decision equivalence for all these parameters. Similarly, our strongest hardness result for decisional-GLWE bypasses the one for search-GLWE. It is unknown if there is a direct search-to-decision equivalence for GLWE, valid for all sufficiently large finite abelian groups. It would also be interesting to study if structural lattice reduction can be adapted to the ring setting, in order to obtain more hardness results based on worst-case assumptions for ideal lattices. Finally, our GSIS and GLWE reductions require the order of G to be sufficiently large compared to the worst-case lattice dimension, and it is interesting to reduce as much as possible this size constraint: in particular, the case $G = \mathbb{Z}_2^n$ for GLWE corresponds essentially to LPN, whose hardness is well-known to be open; here, the order 2^n does not grow quickly enough with respect to the rank n to be covered by our reduction. On the other hand, Micciancio and Peikert [20] recently showed how to decrease q for SIS.

Roadmap. The paper is organized as follows. In Sect. 2, we recall background on lattices. In Sect. 3, we discuss factor groups of integer lattices, and introduce our group generalizations of SIS and LWE. In Sect. 4, we introduce structural lattice reduction, which will be used in all our reductions. We show hardness of GSIS in Sect. 5, and search-GLWE in Sect. 6, by generalizing SIS/LWE reductions. In Sect. 7, we show hardness of decisional-GLWE. Detailed missing proofs can be found in appendix, as well as an example of abstracting lattice cryptography.

2 Background and Notation

$\mathbb{Z}_q$ denotes the group $\mathbb{Z}/q\mathbb{Z}$. We use row notation for vectors and matrices. I_n denotes the $n \times n$ identity matrix. A function $\text{negl}(n)$ is *negligible*, if it vanishes faster than the inverse of any polynomial in n. For an $n \times m$ matrix B, $\|B\|$ denotes the length of its longest row vector, *i.e.* $\|B\| = \max_{1 \le i \le n} \|\mathbf{b}_i\|$.

Lattices. A *lattice* L is a discrete subgroup of $\mathbb{R}^m$: it is of the form $L(B) = \{\sum_{i=1}^n \alpha_i \mathbf{b}_i,\ \alpha_i \in \mathbb{Z}\}$ for some set $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of linearly independent vectors, called a *basis*. When $L \subseteq \mathbb{Z}^m$, L is an *integer lattice*. The *dimension* of L, denoted by $\mathsf{dim}(L)$, is the dimension n of $\mathsf{span}(L)$. Bases of L are related by multiplication with integer matrix of determinant ± 1. The *(co)-volume* $\mathsf{vol}(L)$ is the volume $\sqrt{\det(BB^t)}$ of any basis B of L. For $1 \le i \le \mathsf{dim}(L)$, $\lambda_i(L)$ is the i-th minimum of L, *i.e.* the smallest radius of the 0-centered ball containing at least i linearly independent lattice vectors. The *dual lattice* $L^\times$ is the set of all $\mathbf{u} \in \mathsf{span}(L)$ s.t. $\langle \mathbf{u}, \mathbf{v} \rangle \in \mathbb{Z}$ for all $\mathbf{v} \in L$. If B is a basis of L, its dual basis is $B^\times = (BB^t)^{-1}B$, which is a basis of $L^\times$. For a factor $\gamma = \gamma(n) \ge 1$, GapSVP$_\gamma$ asks, given a basis B of an n-dimensional lattice L and a number $d \in \mathbb{R}_{ge0}$, to decide if $\lambda_1(L) \le d$ or $\lambda_1(L) > \gamma d$. ApproxSIVP$_\gamma$ asks to compute a full-rank family of lattice vectors of norm $\le \gamma \lambda_n(L)$.

Gram-Schmidt Orthogonalization (GSO). Let $B = (\mathbf{b}_1, ..., \mathbf{b}_n)$ be a basis of a lattice. The GSO of B is the unique decomposition as $B = \mu \cdot D \cdot Q$, where μ is a lower triangular matrix with unit diagonal, D is a positive diagonal matrix, and Q has orthonormal rows. We let $B^* = DQ$ whose i-th row $\mathbf{b}_i^*$ is $\pi_i(\mathbf{b}_i)$, where π_i denotes the orthogonal projection of $\mathbf{b}_i$ over $\mathsf{span}\{\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}\}^\perp$. We use the notation $B_{[i,j]}$ for the projected block $[\pi_i(\mathbf{b}_i), \ldots, \pi_i(\mathbf{b}_j)]$. Let $B^\times$ be the dual basis of B and $(B^\times)^*$ denotes its GSO matrix, then $\|(\mathbf{b}_i^\times)^*\| \cdot \|\mathbf{b}_{n-i+1}^*\| = 1$ for $1 \le i \le n$.

(Explicit) Finite abelian groups. Any finite abelian group G is isomorphic to a product $\prod_{i=1}^k \mathbb{Z}_{q_i}$ of cyclic groups. We call *rank* of G the minimal number of cyclic groups in such decompositions: this should not be confused with the rank of an abelian group, which would be zero here. We say that G is *explicit* if one knows integers $q_1, \ldots, q_k$ and an isomorphism $\prod_{i=1}^k \mathbb{Z}_{q_i} \to G$ computable in polynomial time: we will assume that k is the rank and each q_{i+1} divides q_i, because from an arbitrary decomposition, one can always derive the rank and such q_i's in polynomial time. The isomorphism induces k generators $e_1, \ldots, e_k \in G$ s.t. $G = \langle e_1 \rangle \oplus \cdots \oplus \langle e_k \rangle$ and each e_i has order q_i. If the inverse of the isomorphism is also computable in polynomial time, we say that G is *fully-explicit.*

Overlattices and exact sequences. When a lattice $\bar{L} \in \mathbb{R}^n$ contains a sublattice L of the same dimension n, $\bar{L}$ is an *overlattice* of L. Then the quotient $\bar{L}/L$ is a finite abelian group of rank $\leq n$ and order $\mathrm{vol}(L)/\mathrm{vol}(\bar{L})$. Then $0 \to L \overset{\mathrm{id}}{\to} \bar{L} \overset{\varphi}{\to} G \to 0$ is a short exact sequence for some φ, *i.e.* $\varphi : \bar{L} \to G$ is a surjective morphism s.t $\ker \varphi = L$. In other words, φ represents the isomorphism $\bar{L}/L \simeq G$.

Lattice reduction. Gentry *et al.* [14] introduced the *basis length* of a lattice L as $\mathrm{bl}(L) = \min_{\text{basis } B} \|B^*\|$, where $\|B^*\| = \max_{i \in [1,n]} \|\mathbf{b}_i^*\|$. Then: $\lambda_n(L) \geq \mathrm{bl}(L) \geq \lambda_n(L)/\sqrt{n}$, $\mathrm{bl}(L) \geq \lambda_1(L)$, and $\mathrm{bl}(L) \geq \mathrm{vol}(L)^{1/n}$. Lattice reduction can find bases B with small $\|B^*\|$. For instance, a basis B is LLL-reduced [16] with factor $\varepsilon_{_{\mathrm{LLL}}} \geq 0$ if its GSO μ satisfies $|\mu_{i,j}| \leq \frac{1}{2}$ for all $1 \leq j < i$ and every 2×2 block $B_{[i,i+1]}$ satisfies Lovász's condition: $\|\mathbf{b}_i^*\|^2 \leq (1 + \varepsilon_{_{\mathrm{LLL}}})(\left\|\mathbf{b}_{i+1}^*\right\|^2 + \mu_{i+1,i} \|\mathbf{b}_i^*\|^2)$. Then it is folklore that: $\|B^*\| \leq \left((1 + \varepsilon_{_{\mathrm{LLL}}})\sqrt{4/3}\right)^{(n-1)/2} \mathrm{bl}(L)$. Given as input $\varepsilon_{_{\mathrm{LLL}}} > 0$ and a basis B of a lattice $L \subseteq \mathbb{Z}^n$, the LLL algorithm [16] outputs an LLL-reduced basis of factor $\varepsilon_{_{\mathrm{LLL}}}$ in time polynomial in $1/\varepsilon_{_{\mathrm{LLL}}}$ and $\mathrm{size}(B)$. Usually, one selects $\varepsilon_{_{\mathrm{LLL}}}$ s.t. $(1 + \varepsilon_{_{\mathrm{LLL}}})\sqrt{4/3} = \sqrt{2}$ or $\varepsilon_{_{\mathrm{LLL}}} = 1/\mathrm{poly}(n)$.

2.1 Gaussian Measures

The statistical distance between two distributions $\mathcal{P}$ and $\mathcal{Q}$ over a domain X is defined as $\Delta(\mathcal{P}, \mathcal{Q}) = \frac{1}{2}\int_{a \in X} |\mathcal{P}(a) - \mathcal{Q}(a)| da$ or $\frac{1}{2}\sum_{a \in X} |\mathcal{P}(a) - \mathcal{Q}(a)|$ when X is discrete. Two distributions $\mathcal{P}$ and $\mathcal{Q}$ are (statistically) ϵ-indistinguishable if $\Delta(\mathcal{P}, \mathcal{Q}) < \epsilon$. We write $\mathbf{y} \leftarrow \mathcal{P}$ (resp. $\leftarrow_\varepsilon \mathcal{P}$) to denote a sample from the distribution $\mathcal{P}$ (resp. a distribution ϵ-indistinguishable from $\mathcal{P}$). And the symbol $\leftarrow_\approx$ means $\leftarrow_\varepsilon$ for some negligible function ε.

Gaussian Distributions. The *Gaussian Distribution* (over $\mathbb{R}^n$) $\mathcal{D}_{\mathbb{R}^n,\sigma,\mathbf{c}}$ centered at $\mathbf{c} \in \mathbb{R}^n$ of parameter $\sigma \in \mathbb{R}_{ge0}$ is defined by a density function proportional to $\rho_{\mathbb{R}^n,\sigma,\mathbf{c}}(\mathbf{x}) = \exp\left(-\pi\|\mathbf{x} - \mathbf{c}\|^2/\sigma^2\right)$. If $\mathbf{c}$ is omitted, then $\mathbf{c} = 0$. For any countable subset $C \subseteq \mathbb{R}^n$ (*e.g.* a lattice L or a coset $\mathbf{x} + L$), $\rho_{\mathbb{R}^n,\sigma,\mathbf{c}}(C)$ denotes $\sum_{\mathbf{u} \in C} \rho_{\mathbb{R}^n,\sigma,\mathbf{c}}(\mathbf{u})$. The *discrete Gaussian distribution* $\mathcal{D}_{C,\sigma,\mathbf{c}}$ over a lattice or a coset $C \subset \mathbb{R}^n$ is defined by $\mathcal{D}_{C,\sigma,\mathbf{c}}(\mathbf{x}) = \rho_{\mathbb{R}^n,\sigma,\mathbf{c}}(\mathbf{x})/\rho_{\mathbb{R}^n,\sigma,\mathbf{c}}(C)$ where $\mathbf{x} \in C$. It is known how to sample efficiently the discrete Gaussian distribution over lattices to within negligible distance [14, 28], or even exactly [9]:

Lemma 2.1 *There is a polynomial-time algorithm which, given a basis B of an n-dimensional lattice L, $\mathbf{c} \in \mathbb{R}^n$, and a parameter $\sigma \geq \|B^*\| \cdot \sqrt{\ln(2n+4)/\pi}$, outputs a sample with distribution $\mathcal{D}_{L,\sigma,\mathbf{c}}$.*

Reciprocally, on can construct a short lattice basis from short discrete Gaussian samples and an arbitrary basis:

Proposition 2.2 *(Cor. of [30, Lemma 14]) Let $\varepsilon > 0$ and $L(B)$ be an n-dimensional lattice. Given a set of $m = O(n)$ independent Gaussian samples $(\mathbf{y}_i \leftarrow_\varepsilon \mathcal{D}_{L,s_i})$ s.t. $\sqrt{2}\eta_\varepsilon(L) \leq s_i \leq \sigma$, $1 \leq i \leq m$, one can compute in polynomial time a basis C of L s.t. $\|C^*\| \leq \sqrt{n/2\pi} \cdot \max_i s_i$.*

Modular Distributions and Smoothing Parameter. The continuous distribution $\mathcal{D}_{\mathbb{R}^n,\sigma,c}$ and discrete distribution $\mathcal{D}_{\bar{L},\sigma,c}$ over an overlattice $\bar{L} \supseteq L$ can be projected modulo L. Thus $\mathcal{D}_{\mathbb{R}^n/L,\sigma,c}$ (resp. $\mathcal{D}_{\bar{L}/L,\sigma,c}$) has a density function $\mathcal{D}_{\mathbb{R}^n,\sigma,\mathbf{c}}(\mathbf{x} + L)$ for $\mathbf{x} \in \mathbb{R}^n/L$ (resp. $\bar{L}/L$). Both $\mathcal{D}_{\mathbb{R}^n/L,\sigma}$ and $\mathcal{D}_{\bar{L}/L,\sigma}$ converge (uniformly) to the uniform distribution when σ increases. This is quantified by the *smoothing parameter* $\eta_\varepsilon(L)$ (where $\varepsilon > 0$) introduced by Micciancio and Regev [21] as the minimal $\sigma > 0$ s.t. $\rho_{\mathbb{R}^n,\frac{1}{\sigma}}(L^\times \setminus \{0\}) \leq \varepsilon$, *i.e.* $\left\|\mathcal{D}_{\mathbb{R}^n/L,\sigma}(\mathbf{x} + L) - \frac{1}{\mathrm{vol}(L)}\right\|_\infty \leq \frac{\varepsilon}{\mathrm{vol}(L)}$ by Poisson's summation formula, which proves:

Lemma 2.3 (see Cor 2.8 of [14]) *If $\bar{L}$ is an overlattice of L, $\varepsilon \in (0, 1/2)$, $\sigma \geq \eta_\varepsilon(L)$ and $\mathbf{c} \in \mathbb{R}^n$, then $\mathcal{D}_{\bar{L}/L,\sigma,\mathbf{c}+L}$ is within statistical distance at most 2ε from the uniform distribution over $\bar{L}/L$.*

For any n-dim basis B, $\eta_\varepsilon(L(B)) \leq \eta_\varepsilon(L(B^*)) \leq \eta_\varepsilon(\mathbb{Z}^n) \cdot \|B^*\|$ where $\eta_\varepsilon(\mathbb{Z}^n) \leq \sqrt{\log\left(2n \cdot (1+\frac{1}{\varepsilon})\right)/\pi}$. In particular, $\eta_\varepsilon(L) \leq \eta_\varepsilon(\mathbb{Z}^n) \cdot \mathrm{bl}(L)$. Finally, we give a technical lemma on the convolution of a discrete Gaussian inside a dot product, proved in App. A.2, analogous to [31, 28].

Lemma 2.4 (Dot product convolution) *Let $\mathbb{K}$ be $\mathbb{R}$ or $\mathbb{T} = \mathbb{R}/\mathbb{Z}$. Let $c \in \mathbb{R}$, $\mathbf{u} \in \mathbb{R}^n$, $\alpha, \sigma \in \mathbb{R}_{ge0}$, $\varepsilon \in (0, 1/2)$, and $\mathbf{z} + L$ be an arbitrary coset of an n-dimensional lattice $L \subseteq \mathbb{R}^n$. Assume that $\left(\frac{1}{\sigma^2} + \frac{\|\mathbf{u}\|^2}{\alpha^2}\right)^{-1/2} \geq \eta_\varepsilon(L)$. Then the distribution $\mathcal{D}_{\mathbb{K},\alpha,c+\langle\mathbf{u},\mathbf{v}\rangle}$ where $\mathbf{v} \leftarrow \mathcal{D}_{\mathbf{z}+L,\sigma}$ is within statistical distance $\leq 4\varepsilon$ from $\mathcal{D}_{\mathbb{K},\sqrt{\alpha^2+\sigma^2\|\mathbf{u}\|^2},c}$. This still holds when $\mathbb{K}$ is a discrete subgroup $\frac{1}{N}\mathbb{Z}$ of $\mathbb{R}$ or $\frac{1}{N}\mathbb{Z}/\mathbb{Z}$ of $\mathbb{T}$ if $\alpha \geq \eta_\varepsilon(\frac{1}{N}\mathbb{Z})$.*

2.2 SIS and LWE

Let $G = \mathbb{Z}_q^n$. Ajtai's SIS(m, n, q, β) problem [1] asks, given $\mathbf{g} = (g_1, \ldots, g_m)$ chosen uniformly at random from G^m, to find a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ s.t. $\sum_{i=1}^m x_i g_i = 0$ and $\|\mathbf{x}\| \leq \beta$. Such an $\mathbf{x}$ exists if $\beta \geq \sqrt{m} q^{n/m}$. Ajtai [1] proved that SIS (with suitable parameters) is at least as hard as approximating SIVP in the worst case for dimension n to within some polynomial factor: in the best reduction known [14], the factor is $\tilde{O}(n)$. In Regev's LWE(m, n, q, β) problem [31], one picks $\mathbf{s} \in G$ and $(g_1, \ldots, g_m) \in G^m$ uniformly at random. Let A be the $n \times m$ matrix whose i-th column is g_i. LWE asks to recover $\mathbf{s} \in G$, given as input $(A, \mathbf{t} = \mathbf{s}A + \mathbf{e})$ where $\mathbf{e} \in \mathbb{Z}_q^m$ is chosen with distribution $\mathcal{D}_{\mathbb{Z}^m,\beta q}$. In [31], the distribution of $\mathbf{e}$ was slightly different, but Peikert's convolution sampler [28] allows to use this distribution instead.

3 Lattice Factor Groups and Generalizations of SIS and LWE

In this section, we present our group generalizations of SIS and LWE, which are related to factor groups of integer lattices.

3.1 Lattice Factor Groups

If L is a full-rank lattice $\subseteq \mathbb{Z}^m$, its factor group $\mathbb{Z}^m/L$ is a finite abelian group of order $\mathrm{vol}(L)$. For any finite abelian group G, denote by $\mathcal{L}_{G,m}$ the (finite) set of full-rank lattices $L \subseteq \mathbb{Z}^m$ such that $\mathbb{Z}^m/L \simeq G$. The following elementary characterization of $\mathcal{L}_{G,m}$ is a consequence of [26]:

Theorem 3.1 *Let G be a finite abelian group and L be a full-rank lattice in $\mathbb{Z}^m$. Then $L \in \mathcal{L}_{G,m}$ if and only if G has rank $\leq m$ and there exists $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$ such that the g_i's generate G and $L = \mathcal{L}_{\mathbf{g}}$ where $\mathcal{L}_{\mathbf{g}} = \{(x_1, \ldots, x_m) \in \mathbb{Z}^m \text{ s.t. } \sum_{i=1}^m x_i g_i = 0 \text{ in } G\}$.*

Given G, Alg. 1 shows how to sample efficiently lattices from the uniform distribution over $\mathcal{L}_{G,m}$, and its correctness follows from (the trivial) Lemma 3.2. Previously, efficient sampling was only known for $G = \mathbb{Z}_p$ where p is a large prime (see [15]).

Algorithm 1 Sampling lattices of given factor group

```
Input: Integer $m \geq 1$ and a finite abelian group $G = \mathbb{Z}_{q_1} \times \cdots \times \mathbb{Z}_{q_k}$ such that $1 \leq k \leq m$.
Output: A random lattice from the uniform distribution over $\mathcal{L}_{G,m}$.
1: Generate elements $g_1, \ldots, g_m$ uniformly at random from $G$ until the $g_i$'s generate $G$.
2: Return the lattice $\mathcal{L}_{\mathbf{g}}$ where $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$.
```

Lemma 3.2 *Let G be a finite abelian group. Let $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$ be such that the g_i's generate G. Let $\mathbf{h} = (h_1, \ldots, h_m) \in G^m$. Then $\mathcal{L}_{\mathbf{g}} = \mathcal{L}_{\mathbf{h}}$ if and only if there is an automorphism ψ of G such that $h_i = \psi(g_i)$ for all $1 \leq i \leq m$. In such a case, ψ is uniquely determined.*

We note that several implementations of lattice-based cryptography (such as [13]) implicitly used lattices in $\mathcal{L}_{G,m}$ for some large cyclic group G. Recently, Nguyen and Shparlinski [24] showed that such lattices are dominant: the set $\cup_{G \text{ cyclic}} \mathcal{L}_{G,m}$ of all full-rank integer lattices $L \subseteq \mathbb{Z}^m$ such that $\mathbb{Z}^m/L$ is cyclic has natural density $1/[\zeta(6) \prod_{k=4}^m \zeta(k)] \approx 85\%$ (for large m).

3.2 The Group-SIS Problem (GSIS)

We introduce the *Group-SIS* problem (GSIS), which is a natural generalization of SIS to arbitrary finite abelian groups. The GSIS parameters are $m \geq 1$, a finite abelian group G and a bound $\beta \in \mathbb{R}_{ge0}$. One picks a sequence $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$ uniformly at random. GSIS(G, m, β) asks to find a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ s.t. $\sum_{i=1}^m x_i g_i = 0$ and $\|\mathbf{x}\| \leq \beta$. In other words, GSIS asks to find short vectors in random relation lattices $\mathcal{L}_{\mathbf{g}} = \{\mathbf{x} \in \mathbb{Z}^m \text{s.t.} \sum_{i=1}^m x_i g_i = 0\}$. For instance, GSIS$(\mathbb{Z}_q^n, m, \beta)$ is SIS, and GSIS$(\mathbb{Z}_q, m, \beta)$ is finding short vectors in random m-dimensional co-cyclic lattices of volume q. If $\#G$ denotes the order of G, the existence of a GSIS-solution is guaranteed if $\beta \geq \sqrt{m}(\#G)^{1/m}$.

GSIS is connected to $\mathcal{L}_{G,m}$ as follows. It is known [17, 25] that as soon as $m \geq n+2 \log \log \#G+5$ (resp. $m > 2 \log \#G + 2$), $g_1, \ldots, g_m$ generate the whole group G with probability $\geq 1/e$ (resp. $\geq 1 - 1/\#G$), in which case $\mathbb{Z}^m / \mathcal{L}_{\mathbf{g}} \simeq G$. In particular, if $m > 2 \log \#G + 2$, then the distribution of GSIS lattices $\mathcal{L}_{\mathbf{g}}$ is statistically close to the uniform distribution over $\mathcal{L}_{G,m}$, because it is statistically close to the distribution produced by Alg. 1, in which case, solving GSIS is equivalent to finding short vectors in random lattices from $\mathcal{L}_{G,m}$.

Finally, we note that to establish hardness of GSIS, it suffices to focus on low-rank groups G. Indeed, if $G' = G \times H$ for some finite abelian group G, H, then GSIS over G can trivially be reduced to GSIS over G', by "projecting" G' to G.

3.3 The Group-LWE Problem (GLWE)

We introduce the *Group-LWE* problem (GLWE), which similarly generalizes LWE. GLWE uses the torus $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ and a finite abelian group G. Let $\hat{G}$ be the dual group of homomorphisms from G to $\mathbb{T}$: it is isomorphic to G but not canonically. If G is explicit, we have $G = \langle e_1 \rangle \oplus \cdots \oplus \langle e_k \rangle$ where each e_i has order q_i. Then $\hat{G}$ is generated by the characters $\hat{e}_1, \ldots, \hat{e}_k$ defined as $\hat{e}_i(\sum_{j=1}^k \alpha_j e_j) = \alpha_i / q_i \mod 1$ for any $0 \leq \alpha_j < q_j$.

Let $\mathcal{S}$ be a known distribution over $\hat{G}$. Search-GLWE is the problem of learning a character $\hat{s} \in \hat{G}$ picked from $\mathcal{S}$, given noisy evaluations of $\hat{s}$ at (public) random points $a_1, \ldots, a_m \in G$, namely one is given (for all i's) a_i and a "Gaussian" perturbation of $\hat{s}(a_i)$. Like LWE, several noise distributions are possible. As in [31], we focus on the continuous distribution where $\hat{s}(a)$ is shifted by an error $e \leftarrow \mathcal{D}_{\mathbb{R},\alpha}$. These distributions need to be discretized in order to have a finite representation. In App. B.4, we present discrete versions of GLWE and show that they are at least as hard as the continuous version for some suitable parameters, which explains why we only consider the continuous GLWE problem in the rest of the article:

Definition 3.3 *Let G be an explicit finite abelian group: in particular, $G = \langle e_1 \rangle \oplus \cdots \oplus \langle e_k \rangle$. Let $\alpha > 0$ and $\hat{s} \in \hat{G}$.*

- *$A_{G,\alpha}(\hat{s})$ is the distribution over $G \times \mathbb{T}$ defined by choosing $a \in G$ uniformly at random, setting $b \leftarrow \mathcal{D}_{\mathbb{T},\alpha,\hat{s}(a)}$, and outputting $(a, b) \in G \times \mathbb{T}$.*
- *Search-*GLWE$_{G,\alpha}(\mathcal{S})$ *asks to find $\hat{s}$ from $A_{G,\alpha}(\hat{s})$ for a fixed $\hat{s}$ sampled from $\mathcal{S}$ given arbitrarily many independent samples. By finding $\hat{s}$, we mean finding integers s_i's s.t. $\hat{s} = \sum_{i=1}^k s_i \hat{e}_i$.*
- *Decisional-*GLWE$_{G,\alpha}(\mathcal{S})$ *asks to distinguish $A_{G,\alpha}(\hat{s})$ from the uniform distribution over $G \times \mathbb{T}$ for a fixed $\hat{s}$ sampled from $\mathcal{S}$ given arbitrarily many independent samples.*
- *For $0 < \alpha < 1$, (Search) Decisional-*GLWE$_{G,\leq\alpha}(\mathcal{S})$ *is the problem of solving (Search) Decisional-*GLWE$_{G,\beta}(\mathcal{S})$ *for any $\beta \leq \alpha$ respectively,* i.e. *when the noise parameter is unknown yet $\leq \alpha$, by analogy with LWE.*

*Search-*GLWE$_{G,m,\alpha}(\mathcal{S})$ *and Decisional-*GLWE$_{G,m,\alpha}(\mathcal{S})$ *denote the variants where the algorithms are given a bounded number of samples $m \in \mathbb{N}$. If $\mathcal{S}$ is omitted, we mean the uniform distribution over $\hat{G}$.*

If $G = \mathbb{Z}_q^n$, the canonical representation of G and $\hat{G}$ shows that GLWE is equivalent to the fractional version of Regev's original LWE. If $G = \mathbb{Z}_p$ for some prime number p, then $\hat{G}$ can be defined by multiplications: $\hat{\mathbf{s}}$ is the homomorphism that maps any $t \in \mathbb{Z}_p$ to $ts/p \mod 1$. Thus, GLWE can be viewed as a randomized version of Boneh-Venkatesan's *Hidden Number Problem* [8]: recover a secret number $s \bmod p$, given approximations of $st_i \bmod p$ for many random integers t_i's.

By analogy with LWE (see [31, 9]), there is a folklore reduction from (Search) Decisional-$\mathrm{GLWE}_{G,\leq\alpha}(\mathcal{S})$ to (Search) Decisional-$\mathrm{GLWE}_{G,\alpha}(\mathcal{S})$, respectively.

Lemma 3.4 *(Adapted from [9, Lemma 2.13]) Let $\mathcal{A}$ be an algorithm for Decisional-*$\mathrm{GLWE}_{G,m,\alpha}(\mathcal{S})$ *(resp. Search) with advantage at least $\varepsilon > 0$. Then there exists an algorithm $\mathcal{B}$ for Decisional-*$\mathrm{GLWE}_{G,m',\leq\alpha}(\mathcal{S})$ *(resp. Search) using oracle access to $\mathcal{A}$ and with advantage $\geq 1/3$, where both m' and its running time are* poly$(m, 1/\varepsilon, \log \#G)$.

Proof. (Sketch, see App. B.3 for a detailed proof). Like in LWE, the basic idea is to add noises in small increments to the distribution obtained from the challenger, and feed it to the oracle solving the Decisional-$\mathrm{GLWE}_{G,\alpha}(\mathcal{S})$ (resp. Search) and estimate the behavior of the oracle. □

4 Structural Lattice Reduction

4.1 Overview

A basic result (following from structure theorems of finitely-generated modules over principal ideal domains) states that for any full-rank sublattice L of a full-rank lattice $\bar{L} \subseteq \mathbb{R}^n$, there exists a basis $\bar{B} = (\bar{\mathbf{b}}_1, \ldots, \bar{\mathbf{b}}_n)$ of $\bar{L}$ and integers $q_1, \ldots, q_n \geq 1$ such that $q_1 \geq q_2 \geq \cdots \geq q_n \geq 1$ and $B = (q_1\bar{\mathbf{b}}_1, \ldots, q_n\bar{\mathbf{b}}_n)$ is a basis of L. The q_i's can be made unique by selecting them as powers of prime numbers, or by requiring each q_{i+1} to divide q_i, in which case $q_1, \ldots, q_n$ are the *elementary divisors* of the pair $(\bar{L}, L)$: for instance, if $\bar{L} = \mathbb{Z}^n$ and L is a full-rank integer lattice, the q_i's are the diagonal coefficients of the Smith normal form of L.

In this section, we introduce a lattice reduction converse to the previous structure theorem, which we call *structural lattice reduction*. Lattice reduction asks to find a short basis of a given full-rank lattice $L \subseteq \mathbb{Z}^n$. In structural lattice reduction, one is further given a finite abelian group G of rank $\leq n$, and wants to find a *short* basis of some overlattice $\bar{L}$ of L such that $\bar{L}/L \simeq G$ effectively. More precisely, given a basis B of a full-rank lattice $L \subseteq \mathbb{Z}^n$, a suitable bound $\sigma > 0$ and integers $q_1 \geq \cdots \geq q_k$ defining $G = \mathbb{Z}_{q_1} \times \cdots \times \mathbb{Z}_{q_k}$, one asks to compute a basis $\bar{B}$ of an overlattice $\bar{L} \supseteq L$ such that $\|\bar{B}^*\| \leq \sigma$ and $B = (q_1\bar{\mathbf{b}}_1, \ldots, q_k\bar{\mathbf{b}}_k, \bar{\mathbf{b}}_{k+1}, \ldots, \bar{\mathbf{b}}_n)$ is a basis of L. Interestingly, we do not require the input basis B to have integer or rational coefficients, as long as its Gram-Schmidt coefficients are known with enough precision. Indeed, our structural reduction algorithm can simply focus on finding the rational transformation matrix between $\bar{B}$ and B.

Previous worst-case to average-case reductions implicitly used the homocyclic group $G = \mathbb{Z}_q^n$, thus $\bar{L} = L/q$. It follows that finding a basis $\bar{B}$ of $\bar{L}$ with small $\|\bar{B}^*\|$ is the same as finding the basis $B = q\bar{B}$ of L with small $\|B^*\|$, which is just classical lattice reduction. However, we obtain new problems and applications by considering different choices of G.

In the trivial case $G = \mathbb{Z}_q^n$, we note that $\bar{B} = q^{-1}B$ implies that $\|\bar{B}^*\| = \|B^*\|/q$ where the factor q is exactly $\#G^{1/n}$: this suggests that in the general case, we might hope to reduce $\|\bar{B}^*\|$ by a factor close to $\#G^{1/n}$, compared to $\|B^*\|$.

Another trivial case of structural lattice reduction is $G = \mathbb{Z}_{q_1} \times \cdots \times \mathbb{Z}_{q_n}$ where the q_i's are distinct positive integers of similar bit-length. If $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ is a basis of $L \subseteq \mathbb{Z}^n$, then $\bar{B} = (q_1^{-1}\mathbf{b}_1, \ldots, q_n^{-1}\mathbf{b}_n)$ generates an overlattice $\bar{L}$ such that $\bar{B}^* = (q_1^{-1}\mathbf{b}_1^*, \ldots, q_n^{-1}\mathbf{b}_n^*)$, and therefore $\|\bar{B}^*\| \leq \|B^*\|/\min_{i=1}^n q_i$. The factor $\min_{i=1}^n q_i$ is close to $\#G^{1/n}$ if the q_i's have similar bit-length. But if the q_i's are unbalanced, such as when $\min_{i=1}^n q_i = 1$, then the bound is much weaker. In particular, the case $G = \mathbb{Z}_p$ for some large prime p looks challenging, as the trivial choice $\bar{B} = (p^{-1}\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n)$ looks useless: $\bar{L}/L \simeq G$ but $\|\bar{B}^*\|$ is likely to be essentially as big as $\|B^*\|$, because for a typical reduced basis, the first $\|\mathbf{b}_i^*\|$'s have the same size.

4.2 Co-cyclic Lattice Reduction

As a warm-up, we solve structural lattice reduction when the target group G is cyclic of order q, which we call *co-cyclic lattice reduction*. Let $\bar{B}$ be a solution of structural reduction on $(L(B), G, \sigma)$. Then $C = (q\bar{b}_1, \bar{b}_2, \ldots, \bar{b}_n)$ is a basis of L satisfying $\|\mathbf{c}_1\| \leq q\sigma$ and $\|\mathbf{c}_i^*\| \leq \sigma$ for all $i \geq 2$.

To find such a basis $\bar{B}$, we first show how to transform B to ensure $\|\mathbf{b}_i^*\| \leq \sigma$ for all $i \geq 2$, using a polynomial-time algorithm which we call *unbalanced reduction* (see Alg. 2). This algorithm can easily be explained as follows: in dimension two, it is easy to make $\mathbf{b}_2^*$ arbitrarily short by lengthening $\mathbf{b}_1$ (adding a suitable multiple of $\mathbf{b}_2$), since $\|\mathbf{b}_1\| \times \|\mathbf{b}_2^*\| = \mathrm{vol}(L)$ is invariant. Unbalanced reduction works by iterating this process on two-dimensional projected lattices, similarly to the classical size-reduction process. However, one would like to make sure that the resulting first basis vector $\mathbf{c}_1$ does not become too large, which is quantified by the following result:

Theorem 4.1 (Unbalanced reduction) *Given an n-dimensional projected block $B = B'_{[i,i+n-1]}$ of an integer lattice $L \subseteq \mathbb{Z}^m$ and a target $\sigma \in \mathbb{Q}^+$, Algorithm 2 outputs in polynomial time an $n \times n$ unimodular matrix U such that $C = UB$ satisfies $\|\mathbf{c}_1\| \leq n\sigma\delta_\sigma(B)$ and $\|\mathbf{c}_i^*\| \leq \sigma$ for $i \geq 2$, and:*

$$\delta_\nu(B) \;\leq\; \delta_\nu(C) \;\leq\; \frac{\|\mathbf{c}_1\|}{\sigma\delta_\sigma(B)} \times \delta_\nu(B) \;\; \textit{for all } \nu \leq \sigma \tag{1}$$

$$\textit{where } \delta_\sigma(B) \underset{def}{=} \prod_{j=1}^{n} \max\left(1,\; \left\|\mathbf{b}_j^*\right\|/\sigma\right). \tag{2}$$

We call $\delta_\sigma(B)$ the *cubicity-defect* of B relatively to σ: it basically measures by which amount the hypercube of side σ should be scaled up to cover the parallelepiped spanned by $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$. The proofs of Th. 4.1 and Alg. 2 can be found in App. C.2. Th. 4.1 shows that Alg. 2 solves co-cyclic lattice reduction for $q \geq n\delta_\sigma(B)$. However, this may not be suitable for our applications, since this lower bound depends on B and might be unbounded. To address this issue, we now show that LLL can bound $\delta_\sigma(B)$ depending only on n for appropriate σ:

Algorithm 2 Unbalanced Reduction

Input: an $n \times m$ basis B of an integer lattice $L \subseteq \mathbb{Z}^m$ and a target length $\sigma \in \mathbb{Q}^+$. More generally, B can be any n-dimensional projected block $B = B'_{[i,i+n-1]}$ of some basis B' of $L \subseteq \mathbb{Z}^m$.

Output: an $n \times n$ unimodular matrix U such that $C = UB$ satisfies $\|c_i^*\| \leq \sigma$ for $i \geq 2$ and $\|\mathbf{c}_1\| \leq n\sigma\delta_\sigma(B)$.

```
$C \leftarrow B$, $U \leftarrow I_n$ and compute the Gram-Schmidt matrices $\mu$ and $C^*$
If $\|\mathbf{c}_i^*\| \leq \sigma$ for all $i$, return $U$
for $i = k-1$ downto 1 where $k$ is the largest index such that $\|\mathbf{c}_k^*\| > \sigma$ do
    if $\|\mathbf{c}_i^*\| \leq \sigma$ then
        $\alpha \leftarrow \lfloor -\mu_{i+1,i} \rceil$
    else
        $\alpha \leftarrow \left\lceil -\mu_{i+1,i} + \frac{\|\mathbf{c}_{i+1}^*\|}{\|\mathbf{c}_i^*\|}\sqrt{(\|\mathbf{c}_i^*\|/\sigma)^2 - 1}\right\rfloor$
    end if
    $(\mathbf{c}_i, \mathbf{c}_{i+1}) \leftarrow (\mathbf{c}_{i+1} + \alpha \cdot \mathbf{c}_i,\ \mathbf{c}_i)$, $(\mathbf{u}_i, \mathbf{u}_{i+1}) \leftarrow (\mathbf{u}_{i+1} + \alpha \cdot \mathbf{u}_i,\ \mathbf{u}_i)$ and update the Gram-Schmidt matrices $\mu$ and $C^*$.
end for
return $U$
```

Theorem 4.2 (LLL's cubicity-defect) *Let L be a full-rank lattice in $\mathbb{R}^n$ and $\sigma \geq ((1+\varepsilon_{_{LLL}})\sqrt{4/3})^r \cdot \mathrm{bl}(L)$ for some $r \geq 0$. If B is an LLL-reduced basis of L with factor $\varepsilon_{_{LLL}}$, then $\delta_\sigma(B) \leq ((1+\varepsilon_{_{LLL}})\sqrt{4/3})^{\frac{(n-2r)^2}{8}+\frac{(n-2r)}{4}}$.*

By combining Th. 4.1 and 4.2, we obtain:

Theorem 4.3 (Co-cyclic Reduction) *Given an $n \times m$ basis of a lattice $L \subseteq \mathbb{Z}^m$, $\varepsilon > 0$ and a rational $\sigma \geq ((1+\varepsilon_{_{LLL}})\sqrt{4/3})^r \cdot \mathrm{bl}(L)$ for some $r \geq 0$, and an integer $q \geq n((1+\varepsilon_{_{LLL}})\sqrt{4/3})^{\frac{(n-2r)^2}{8}+\frac{(n-2r)}{4}}$, Alg. 3 computes a basis $\bar{B}$ of an overlattice $\bar{L} \supseteq L$ in time polynomial in the basis size, σ and $1/\varepsilon$, such that $\left\|\bar{B}^*\right\| \leq \sigma$ and $(q\bar{\mathbf{b}}_1, \bar{\mathbf{b}}_2, \ldots, \bar{\mathbf{b}}_n)$ is a basis of L. In particular, $\bar{L}/L \simeq \mathbb{Z}_q$.*

For instance, Th. 4.3 with $r = n$ implies that given a lattice L and any cyclic group G of sufficiently large order (*i.e.* $2^{\Omega(n^2)}$), one can efficiently obtain a basis $\bar{B}$ of some overlattice $\bar{L}$ of L such that $\bar{L}/L \simeq G$ and $\|\bar{B}^*\| \leq \mathrm{bl}(L)$: by comparison, an LLL-reduced basis only approximates $\mathrm{bl}(L)$ to some exponential factor in the worst case.

Algorithm 3 Co-cyclic Reduction

Input: a basis of a full-rank integer lattice $L \subseteq \mathbb{Z}^n$, a factor $\varepsilon > 0$, and a rational $\sigma \geq ((1+\varepsilon_{\text{LLL}})\sqrt{4/3})^r \cdot \text{bl}(L)$ for some $r \geq 0$, and an integer $q \geq n((1+\varepsilon_{\text{LLL}})\sqrt{4/3})^{\frac{(n-2r)^2}{8}+\frac{(n-2r)}{4}}$

Output: a basis $\bar{B}$ of an overlattice $\bar{L}$ such that $\|\bar{B}^*\| \leq \sigma$ and $\bar{L}/L \simeq \mathbb{Z}_q$.

```
Apply Alg. 2 on an LLL-reduced basis with factor $\varepsilon_{\text{LLL}}$ output by the LLL algorithm to find a basis $C$ of $L$.
return $\bar{B} = (\frac{\mathbf{c}_1}{q}, \mathbf{c}_2, \ldots, \mathbf{c}_n)$
```

4.3 Arbitrary Groups

Using unbalanced reduction, we prove that for an arbitrary sufficiently large finite abelian group G of rank $\leq n$, given any basis B of the lattice $L \subseteq \mathbb{Z}^n$, one can compute a basis $\bar{B}$ of some overlattice $\bar{L}$ of L s.t. $\bar{L}/L \simeq G$ effectively and $\|\bar{B}^*\|$ is essentially lower than $\|B^*\|/\#G^{1/n}$. In particular, $\text{bl}(\bar{L})$ is essentially $\#G^{1/n}$ smaller than $\text{bl}(L)$. Although this is slightly weaker than the result we obtained (in the previous subsection) for cyclic groups G, it is sufficient for our worst-case to average-case reductions.

Theorem 4.4 (Structural Lattice Reduction) *Given an $n \times m$ basis B of a lattice $L \subseteq \mathbb{Z}^n$, and $k \leq n$ integers $q_1 \geq \cdots \geq q_k$ defining the group $G = \prod_{i=1}^k \mathbb{Z}_{q_i}$ s.t. $n^k(\|B^*\|/\sigma)^n \leq \#G$ or:*
$\#G \geq \frac{n!}{(n-k)!}\delta_\sigma(B)$ and for all $i \leq k$, $\|B^*\|/\sigma \leq q_i/(n+1-i)$
Alg. 4 outputs a basis $\bar{B}$ of an overlattice $\bar{L} \supseteq L$ such that $\left\|\bar{B}^*\right\| \leq \sigma$ and $(q_1\bar{\mathbf{b}}_1, \ldots, q_n\bar{\mathbf{b}}_n)$ is a basis of L where $q_i = 1$ for $i > k$. In particular, $\bar{L}/L \simeq G$.

For instance, the condition $n^k(\|B^*\|/\sigma)^n \leq \#G$ in Th. 4.4 means that σ (and therefore $\|\bar{B}^*\|$) can be chosen as low as $n^{k/n}\|B^*\|/(\#G)^{1/n}$. The proof of Th. 4.4 can be found in App. C.3.

Algorithm 4 Structural Lattice Reduction

Input: σ, an $n \times m$ basis B of an integer lattice L, and $(q_1, \ldots, q_k)$ s.t. $G = \mathbb{Z}_{q_1} \times \cdots \mathbb{Z}_{q_k}$ satisfies the conditions of Th. 4.4

Output: an $n \times m$ basis $\bar{B}$ of an overlattice $\bar{L}$ of L such that $\|\bar{B}^*\| \leq \sigma$ and $\bar{L}/L \simeq G$.

```
$C \leftarrow B$
for $i = 1$ to $k$ do
    if $\left\|C^*_{[i,n]}\right\| \leq \sigma$ return $\bar{B} = (\frac{\mathbf{c}_1}{q_1}, \ldots, \frac{\mathbf{c}_k}{q_k}, \mathbf{c}_{k+1}, \ldots, \mathbf{c}_n)$
    Compute the smallest $\ell \geq \sigma$ such that $\ell \cdot \delta_\ell(C_{[i,n]}) = q_i\sigma/(n-i+1)$.
    $V \leftarrow$ UnbalancedReduction$(C_{[i,n]}, \sigma)$ using Alg. 2.
    Apply $V$ on $(\mathbf{c}_i, \ldots, \mathbf{c}_n)$
end for
return $\bar{B} = (\frac{\mathbf{c}_1}{q_1}, \ldots, \frac{\mathbf{c}_k}{q_k}, \mathbf{c}_{k+1}, \ldots, \mathbf{c}_n)$
```

Intuitively, Alg. 4 simply applies unbalanced reduction iteratively, cycle by cycle of G.

4.4 Application

Structural reduction finds a short overlattice basis, which can typically be used to sample short (overlattice) vectors, and which provides the following effective isomorphisms:

Proposition 4.5 *Let L and $\bar{L}$ be two full-rank lattices such that $\bar{L} \supseteq L$ and $\bar{L}/L \simeq G$ where $G = \mathbb{Z}_{q_1} \times \cdots \times \mathbb{Z}_{q_k}$. Given bases B and $\bar{B}$ of resp. L and $\bar{L}$, one can compute in polynomial time a morphism φ s.t. the sequence $0 \to L \xrightarrow{\text{id}} \bar{L} \xrightarrow{\varphi} G \to 0$ is exact, and a "dual" morphism $\varphi^\times : L^\times \to \hat{G}$ s.t.*

$$[\varphi^\times(\mathbf{u})](\varphi(\mathbf{v})) = \langle \mathbf{u}, \mathbf{v} \rangle \mod 1 \text{ for all } \mathbf{u} \in L^\times \text{ and all } \mathbf{v} \in \bar{L} \quad (3)$$

Furthermore, preimages of $\varphi^\times$ can be computed in polynomial time.

Proof. (Sketch) Let $(e_1, \ldots, e_k)$ be the canonical generators of $G = \prod_{i=1}^k \mathbb{Z}_{q_i}$. Find any basis C of L and $\bar{C}$ of $\bar{L}$ such that $C = (q_1\bar{\mathbf{c}}_1, \ldots, q_k\bar{\mathbf{c}}_k, \bar{\mathbf{c}}_{k+1}, \ldots \bar{\mathbf{c}}_n)$, then let φ be the morphism mapping C to $(e_1, \ldots, e_k, 0, \ldots, 0)$ and $\varphi^\times$ be the mapping from $C^\times$ to $(\hat{e}_1, \ldots, \hat{e}_k, \hat{0}, \ldots, \hat{0})$. □

This proposition still holds if G is an explicit finite abelian group.

5 Hardness of Group-SIS

Our hardness result for GSIS requires that the finite abelian group G is *explicit* (see Sect. 2).

5.1 Overview

We first sketch how to adapt the SIS reduction to GSIS using structural lattice reduction.

The main idea behind the SIS reduction can be traced back to 1935, when Mordell [22] published an arithmetical proof of Minkowski's theorem. To prove the existence of short non-zero vectors in an arbitrary full-rank lattice $L \subseteq \mathbb{R}^n$, Mordell implicitly presented an algorithm to find short vectors from (exponentially many) long vectors, as follows. Let $q \geq 1$ be an integer and $\mathbf{w}_1, \ldots, \mathbf{w}_m \in L$ be distinct vectors of norm $\leq R$, where $m > q^n$: for large R, m can be essentially chosen as large as the volume of the R-radius ball divided by the volume of L. Letting $\mathbf{v}_i = q^{-1}\mathbf{w}_i$, we have $\mathbf{v}_i \in q^{-1}L$. Since $m > q^n = [(q^{-1}L) : L]$, there must be $i \neq j$ such that $\mathbf{v}_i \equiv \mathbf{v}_j \bmod L$, *i.e.* $\mathbf{v}_i - \mathbf{v}_j = q^{-1}(\mathbf{w}_i - \mathbf{w}_j) \in L$ whose (nonzero) norm is $\leq 2R/q$, which is suitably short for appropriate choices of q and R.

This algorithm is not efficient since m is exponential in q, but it can be made polynomial by reducing m to $\mathrm{poly}(n)$, using a SIS oracle for (q, m, n). Indeed, let L be an arbitrary full-rank integer lattice in $\mathbb{Z}^n$. The lattice $\bar{L} = q^{-1}L$ is an overgroup of L such that $\bar{L}/L \simeq \mathbb{Z}_q^n = G$: namely, there is an exact sequence of groups $0 \to L \xrightarrow{\mathrm{id}} \bar{L} \xrightarrow{\varphi} G \to 0$, where φ is efficiently computable, *e.g.* for any fixed basis $(\bar{\mathbf{b}}_1, \ldots, \bar{\mathbf{b}}_n)$ of $\bar{L}$, one may take $\varphi(\sum_{i=1}^n x_i \bar{\mathbf{b}}_i) = (x_1 \bmod q, \ldots, x_n \bmod q) \in G$.

Furthermore, if the basis $\bar{B}$ is short enough compared to the minima of L, it is possible to sample short vectors $\mathbf{v}_1, \ldots, \mathbf{v}_m \in \bar{L}$ with Gaussian distribution of parameter as small as $\eta_\varepsilon(L)$. Fourier analysis guarantees that for such parameter of the Gaussian distribution, each projection $g_i = \varphi(\mathbf{v}_i)$ is uniformly distributed over G. This allows us to call an SIS oracle on $(g_1, \ldots, g_m)$, which outputs a short $\mathbf{x} \in \mathbb{Z}^m$ such that $\sum_{i=1}^m x_i g_i = 0$, *i.e.* $\sum_{i=1}^m x_i \varphi(\mathbf{v}_i) = 0$ which implies that $\mathbf{v} = \sum_{i=1}^m x_i \mathbf{v}_i \in L$. This $\mathbf{v}$ can be proved to be non-zero with overwhelming probability, and it is short because the $\mathbf{v}_i$'s and $\mathbf{x}$ are short, which concludes the reduction from worst-case SIVP to SIS.

With this formalization, we can replace the SIS oracle by a GSIS oracle, as while as we are able to sample short vectors $\mathbf{v}_1, \ldots, \mathbf{v}_m \in \bar{L}$ with Gaussian distribution, where $\bar{L}/L \simeq G$. And this is exactly what structural lattice reduction ensures.

5.2 Reducing Worst-case ApproxSIVP to GSIS

Our main result formalizes the previous sketch and states that for appropriate choices of (G, m, β), if one can solve $\mathrm{GSIS}(G, m, \beta)$ on average, then one can approximate SIVP in the worst case, *i.e.* one can efficiently find short vectors in every n-dimensional lattice:

Theorem 5.1 *Let $n \in \mathbb{N}$ and $\varepsilon = \mathrm{negl}(n)$. Given as input a basis B of a full-rank integer lattice $L \subseteq \mathbb{Z}^n$ and $\sigma \geq \sqrt{2}\,\mathrm{bl}(L)$, and an explicit finite abelian group G of rank $k \leq n$ such that $\#G \geq n^k(\|B^*\|/\sigma)^n$, Alg. 5 outputs (in random polynomial time) n linearly independent vectors of L with norm $\leq \sigma\eta_\varepsilon(\mathbb{Z}^n)\sqrt{n\pi}\beta$, using polynomially many calls to an oracle solving $\mathrm{GSIS}(G, m, \beta)$ with probability $\geq 1/\mathrm{poly}(n)$.*

Algorithm 5 Reducing ApproxSIVP to GSIS

```
Input: a basis $B$ of a full-rank integer lattice $L \in \mathbb{Z}^n$, a parameter $\sigma \geq \sqrt{2}\,\mathrm{bl}(L)$, an explicit finite abelian group $G$ satisfying the
    condition of Th. 5.1, and an oracle $\mathcal{O}$ solving $\mathrm{GSIS}(G, m, \beta)$ with probability $\geq 1/\mathrm{poly}(n)$, a negligible $\varepsilon > 0$
Output: A set $S$ of $n$ linearly independent vectors of $L$ of norm $\leq \sigma\eta_\varepsilon(\mathbb{Z}^n)\sqrt{n/2\pi}\beta$.
$S \leftarrow \emptyset$.
Call structural lattice reduction (Alg. 4) on $(B, G, \sigma)$ to get $\bar{B}$ s.t. $\|\bar{B}^*\| \leq \sigma$ and $\varphi : \bar{L} \to G$ (Prop. 4.5) where $\bar{L} = L(\bar{B})$.
repeat
    Sample $\mathbf{v}_1, \cdots, \mathbf{v}_m \in \bar{L}$ with distribution $D_{\bar{L}, \sigma\eta_\varepsilon(\mathbb{Z}^n), \mathbf{0}}$ using $\bar{B}$.
    $g_i = \varphi(\mathbf{v}_i)$ for $1 \leq i \leq m$, forming a sequence $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$.
    Call the GSIS-oracle $\mathcal{O}$ on $\mathbf{g}$, which returns $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{Z}^m$ s.t. $\sum_{i=1}^m x_i g_i = 0$.
    $\mathbf{v} \leftarrow \sum_{i=1}^m x_i \mathbf{v}_i \in L$
    if $\|\mathbf{v}\| \leq \sigma\eta_\varepsilon(\mathbb{Z}^n)\sqrt{n\pi}\beta$ and $\mathbf{v} \notin \mathrm{span}(S)$ then $S \leftarrow S \cup \{\mathbf{v}\}$
until $\dim(S) = n$
Return $S$
```

In particular, by letting $\sigma = \frac{\|B\|^*}{2\eta_\varepsilon(\mathbb{Z}^n)\sqrt{n/\pi}\beta}$, we can obtain an incremental version of the reduction, where the output basis is twice as short as the input. This generalizes [21, Th. 5.9] and [14, Th. 9.2] with a GSIS oracle instead of SIS. Iterating Th. 5.1 until we reach $\sigma = \sqrt{2}\,\mathrm{bl}(L)$ allows to connect the average-case hardness of GSIS to the worst-case of ApproxSIVP.

Corollary 5.2 *Let $n \in \mathbb{N}$ and $\varepsilon = \operatorname{negl}(n)$. Let $(G_n)_{n\in\mathbb{N}}$ be a sequence of explicit finite abelian groups s.t. $\#G_n \leq (\beta_n/\sqrt{m_n})^{m_n}$ for $m_n \in \mathbb{N}$ and G_n has rank k_n. If $\#G_n \geq n^{k_n}\left(\eta_\varepsilon(\mathbb{Z}^n)\sqrt{2n/\pi}\beta_n\right)^{\max(n,k_n)}$, then using polynomially many calls to an oracle solving* $\mathrm{GSIS}(G_n, m_n, \beta_n)$ *with probability $\geq 1/\operatorname{poly}(n)$, one can solve the worst-case n-dimensional* $\mathrm{ApproxSIVP}_{\eta_\varepsilon(\mathbb{Z}^n)\sqrt{n/\pi}\beta_n}$ *in (randomized) polynomial time.*

Consider the set of all full-rank integer lattices $\subseteq \mathbb{Z}^m$ of volume $\geq \omega_n = n^m\left(\eta_\varepsilon(\mathbb{Z}^n)\sqrt{2n/\pi}\beta_n\right)^m$. This set can be partitioned as $\cup_G \mathcal{L}_{G,m}$ where G runs over all finite abelian groups of order $\geq \omega_n$ and rank $\leq m$. Each such G satisfies the conditions of Cor. 5.2, and therefore GSIS over G is as hard as worst-case lattice problems: for any of the partition cells $\mathcal{L}_{G,m}$, finding short vectors in a random lattice from this cell is as hard as finding short vectors in any n-dim lattice.

6 Hardness of Search-Group-LWE

Like GSIS, our hardness result for GLWE requires that the finite abelian group G is *explicit*. The main result of this section states that for appropriate choices of (G, m, α), if one can solve Search-$\mathrm{GLWE}_{G,m,\leq\alpha}$ on average with probability $\geq 1/\operatorname{poly}(n)$, then one can quantumly approximate SIVP in the worst case, *i.e.* one can (quantum)-efficiently find short vectors in every n-dimensional lattice, which generalizes Regev's quantum Search-LWE reduction [31]. To do this, we only need to modify the classical part of Regev's proof, not the quantum part. More precisely, we only need to prove that a GLWE-oracle allows us to approximate bounded distance decoding (BDD) for dual lattices in the worst-case for some factor β: given a basis $B^\times$ of a dual lattice $L^\times$, and a target $\mathbf{t} \in \operatorname{span}(L^\times)$ within distance $\leq \beta\lambda_1(L^\times)$ to $L^\times$, find the lattice point $\mathbf{u} \in L^\times$ closest to $\mathbf{t}$.

Let us first explain the main difference with LWE. In previous proofs, the LWE-oracle is used to transform any β-BDD on $L^\times$ into an β/q-BDD over the same lattice $L^\times$. One iterates this process k times until the distance β/q^k becomes smaller than $2^{-O(n)}\lambda_1(L^\times)$, at which point Babai's nearest plane algorithm [4] solves the BDD instance in polynomial time. To allow arbitrary structures G, we reinterpret this as reducing β-BDD on $L^\times$ to β-BDD over $\bar{L}^\times$, where $\bar{L} = L/q$. Thus, instead of reducing the distance, we modify the lattice to increase $\lambda_1(L^\times)$ until the BDD instance can be solved by Babai's algorithm. This approach allows arbitrary overlattices $\bar{L}$, just like in our GSIS reduction.

More precisely, consider a BDD instance over $L^\times$: we have a target $\mathbf{t} \in \operatorname{span}(L^\times)$ close to some secret $\mathbf{u} \in L^\times$. Let $\hat{s} = \varphi^\times(\mathbf{u}) \in \hat{G}$. Remember that structural lattice reduction gives an exact sequence of groups $0 \to L \xrightarrow{\mathrm{id}} \bar{L} \xrightarrow{\varphi} G \to 0$, where φ is efficiently computable. Let $\varphi^\times$ be as in Prop. 4.5. Like in the GSIS reduction, we sample short vectors $\mathbf{v}_1, \ldots, \mathbf{v}_m \in \bar{L}$ with Gaussian distribution, in such a way that each projection $g_i = \varphi(\mathbf{v}_i)$ is uniformly distributed over G. Then: $\hat{s}(g_i) = [\varphi^\times(\mathbf{u})](\varphi(\mathbf{v}_i)) \equiv \langle \mathbf{u}, \mathbf{v}_i\rangle \pmod 1$. Since $\mathbf{t}$ is close to $\mathbf{u}$, each $\langle \mathbf{t}, \mathbf{v}_i\rangle$ is therefore close to $\hat{s}(g_i)$ mod 1, namely $\langle \mathbf{t}, \mathbf{v}_i\rangle - \hat{s}(g_i) \equiv \langle \mathbf{t} - \mathbf{u}, \mathbf{v}_i\rangle \pmod 1$. By adding a suitable noise, it is possible to simulate the distribution of the GLWE noisy approximation of $\hat{s}(g_i)$ (using Lemma 2.4). Then one can recover the character $\hat{s}$ by calling a Search-GLWE oracle: this allows to compute $\mathbf{u}' \in L^\times$ s.t. $\varphi^\times(\mathbf{u}') = \varphi^\times(\mathbf{u})$. One can compute $\mathbf{t} - \mathbf{u}'$, which is a target equally close to $\mathbf{u} - \mathbf{u}' \in \bar{L}^\times$, as $\mathbf{t}$ was close to $\mathbf{u} \in L^\times$. Hence, we have transformed a BDD-instance over $L^\times$ into a BDD-instance over $\bar{L}^\times$ with exactly the same error $\mathbf{t} - \mathbf{u}$. By iterating this process, one is eventually able to solve the BDD instance efficiently. Formally, we have:

Theorem 6.1 *Let $n \in \mathbb{N}$, $\varepsilon = \operatorname{negl}(n)$, a BDD factor $\beta \leq \sqrt{\pi/2}\,(2n\eta_\varepsilon(\mathbb{Z}^n))^{-1}$ and $\theta = \beta\sqrt{2/\pi}\,(2n\eta_\varepsilon(\mathbb{Z}^n))$. Let $\alpha \in \left]\theta\sqrt{\pi/2}, \sqrt{\pi/2}\right]$ and an explicit finite abelian group G of rank $k \leq n$. Given as input a basis B of an n-dimensional lattice L and $\mathbf{t} \in \operatorname{span}(L^\times)$ such that the BDD instance $(B^\times, \mathbf{t})$ admits a unique solution $\mathbf{t} - \mathbf{w} \in L^\times$ with $\|\mathbf{w}\| \leq \beta\lambda_1(L^\times)$, and a Search-*$\mathrm{GLWE}_{G,m,\leq\alpha}$ *oracle satisfying*

$$\#G \geq n^k \left(\frac{\|B^*\|}{\operatorname{bl}(L)} \frac{\theta\sqrt{\pi/2}}{\sqrt{2}\alpha}\right)^n, \tag{4}$$

Alg. 6 finds in time polynomial in n and $\log(1/\varepsilon)$ a basis $\bar{B}$ of some overlattice $\bar{L}$ such that $\bar{L}/L \simeq G$, and a target $\bar{\mathbf{t}} \in \operatorname{span}(\bar{L}^\times)$ such that the BDD instance $(\bar{B}^\times, \bar{\mathbf{t}})$ have the same error $\mathbf{w}$, namely $\bar{\mathbf{t}} - \mathbf{w} \in \bar{L}^\times$. If B is LLL-reduced with factor ε_{LLL}, we can replace $\frac{\|B^*\|}{\operatorname{bl}(L)}$ by $((1+\varepsilon_{LLL})\sqrt{4/3})^{\frac{n-1}{2}}$ in (4).

The proof of Th. 6.1 is essentially summarized by Alg. 6, which makes a few simplifying assumptions.

Algorithm 6 Reducing BDD to GLWE

```
Input: A dimension $n$ and a negligible probability $\varepsilon = \mathrm{negl}(n)$, a basis $B$ of a $n$-dimensional integer lattice $L$, an average-oracle $\mathcal{O}$ for Search-GLWE$_{G,m,\leq\alpha}$ satisfying the conditions of Th. 6.1, a BDD factor $\beta$, a target $\mathbf{t}$ and an upper-bound $d_0 \leq \beta\lambda_1(L^\times)$ on the error norm.
Output: a basis $\bar{B}$ of length $\|\bar{B}^*\| \leq \|B^*\|/2$ of some ($G$-)overlattice $\bar{L}$ such that and a target $\bar{\mathbf{t}} \in \mathrm{span}(\bar{L}^\times)$ such that the BDD instance $(\bar{B}^\times, \bar{\mathbf{t}})$ has the same error $\mathbf{w}$ of norm $\leq d_0$ than $(B^\times, \mathbf{t})$
1: $\sigma_0 \leftarrow \frac{\alpha}{\sqrt{2}d_0\eta_\varepsilon(\mathbb{Z}^n)} \geq \frac{\alpha\sqrt{2}}{\theta\sqrt{\pi/2}}\,\mathrm{bl}(L) \geq \sqrt{2}\,\mathrm{bl}(L)$.
2: Call structural lattice reduction (Alg.4) on $(B, G, \sigma_0)$ to get $(\bar{B}, \bar{L})$ and $\varphi : \bar{L} \to G$, $\varphi^\times : L^\times \to \hat{G}$ (Prop. 4.5)
3: repeat
4:     Sample $m$ random points $(\mathbf{v}_1, \cdots, \mathbf{v}_m) \in \bar{L}$ with distribution $\mathcal{D}_{\bar{L},\sigma_0\eta_\varepsilon(\mathbb{Z}^n)}$ using $\bar{B}$.
5:     Let $a_i = \varphi(\mathbf{v}_i)$ and $b_i \leftarrow \mathcal{D}_{\mathbb{T},\frac{\alpha}{\sqrt{2}},\langle\mathbf{t},\mathbf{v}_i\rangle}$, to form $(a_i, b_i)_{i\in[1,m]} \in (G \times \mathbb{T})^m$.
6:     Call the Search-GLWE$_{G,m,\leq\alpha}$ oracle on $(a_i, b_i)_{i\in[1,m]}$ to find $\hat{s} \in \hat{G}$.
7: until Search-GLWE$_{G,m,\leq\alpha}$ finds a solution
8: $\bar{\mathbf{t}} \leftarrow \mathbf{t} - \mathbf{u}$ where $\mathbf{u} \in \varphi^{\times-1}(\hat{s})$ (take any preimage modulo $\bar{L}^\times$)
9: return $\bar{B}, \bar{t}$
```

In Step. 6 of Alg. 6, the Search-GLWE$_{G,m,\alpha}$ oracle is called directly on the $(a_i, b_i)_{i\in[1,m]}$, whereas, strictly speaking, we should actually randomize these inputs to make sure that the solution $\mathbf{s}$ follows the right distribution: in the classical LWE reduction, one also uses the self-reducibility of LWE. To make sure that the input has the right distribution, the key step is Step. 5. Note that $\langle\mathbf{t}, \mathbf{v}_i\rangle = \langle\mathbf{u}, \mathbf{v}_i\rangle + \langle\mathbf{t} - \mathbf{u}, \mathbf{v}_i\rangle \bmod 1$, where the first term is equal to $\hat{s}(a_i) = \langle\varphi^\times(\mathbf{u}), \varphi(\mathbf{v}_i)\rangle$. Since $b_i \leftarrow \mathcal{D}_{\mathbb{T},\sqrt{\alpha}/2,\hat{s}(a_i)+\langle\mathbf{t}-\mathbf{u},\mathbf{v}_i\rangle}$ where $\mathbf{v}_i \leftarrow \mathcal{D}_{L,\sigma_0\eta_\varepsilon(\mathbb{Z}^n)}$, Lemma 2.4 proves that b_i has the requested distribution $\mathcal{D}_{\mathbb{T},\alpha',\hat{s}(a_i)}$ for some $\alpha' \leq \alpha$

By iterating Alg. 6 and Th. 6.1 a polynomial number of times, as the length of the input basis geometrically decreases, then $\lambda_1(L^\times)$ geometrically increases. Eventually, the BDD instance becomes easy, and the error $\mathbf{w}$ can be retrieved using for instance Babai nearest plane algorithm. Thus we deduce the following result on the hardness of Search-LWE.

Corollary 6.2 *Let $n \in \mathbb{N}$, $\varepsilon = \mathrm{negl}(n)$ and two real sequences $\beta_n \leq \sqrt{\pi/2}\,(2n\eta_\varepsilon(\mathbb{Z}^n))^{-1}$, and $\alpha_n \in \left]\theta_n\sqrt{\pi/2}, \sqrt{\pi/2}\right]$ where $\theta_n = \beta_n\sqrt{2/\pi}\,(2n\eta_\varepsilon(\mathbb{Z}^n))$. Let $(G_n)_{n\in\mathbb{N}}$ be a sequence of explicit finite abelian groups of rank k_n. If $\#G_n \geq n^{k_n}\left((1+\varepsilon_{LLL})\sqrt{4/3})^{\frac{n-1}{2}}\frac{\theta_n\sqrt{\pi/2}}{\sqrt{2}\alpha_n}\right)^{\max(n,k_n)}$, then using polynomially many calls to an oracle solving Search-*GLWE$_{G_n,\leq\alpha_n}$ *with probability* $1/\mathrm{poly}(n)$, *one can solve worst-case n-dimensional* BDD$_{\beta_n}$ *in (randomized) polynomial time and* ApproxSIVP$_{\sqrt{2}n/\beta_n}$ *in quantum polynomial time.*

7 Hardness of Decisional-Group-LWE

In this section, we give two types of reductions for the hardness of decisional-GLWE.

First, it is well-known that in certain cases, such as when q is a small prime number (see [31, 19]), there is an elementary reduction from search-LWE to decisional LWE, which allows to extend hardness results on search-LWE to decisional-LWE. These special search-to-decision reductions can easily be adapted to GLWE, but they only work for very special choices of G (such as $G = \prod_{i=1}^n q_i$ where the q_i's are small prime numbers), which is insufficient to prove the hardness of decisional-GLWE for arbitrary (sufficiently large) finite abelian groups G. More precisely, we have the following adaptation of [19], when the target group order is smooth:

Theorem 7.1 (Search-to-Decision) *Let $n \in \mathbb{N}$ and $(G_n)_{n\in\mathbb{N}} = \langle e_1\rangle \oplus \cdots \oplus \langle e_{k_n}\rangle$ be a sequence of abelian finite groups, where each e_i has order $q_i(n)$. Let $q_i(n)$ have prime factorization $q_i = p_{i,1}^{\mu_{1,i}} \ldots p_{i,t_i}^{\mu_{i,t_i}}$ for pairwise distinct polynomially bounded prime p_{i,j_i} with $\mu_{i,j_i} \geq 1$, where $i \in [1, k_n], j_i \in [1, t_i]$. Let $0 < \alpha_n \leq 1/\omega(\sqrt{\log n})$ be a real sequence and ℓ be the number of prime factors $p_{i,j_i} < \omega\sqrt{\log n}/\alpha_n$. There is a probabilistic polynomial-time reduction from Search-*GLWE$_{G_n,\alpha_n}$ *to Decisional-*GLWE$_{G_n,\alpha'_n}$ *for any $\alpha'_n \geq \alpha_n$ such that $\alpha'_n \geq \omega(\sqrt{\log n})/p_{i,j_i}^{\mu_{i,j_i}}$ for $i \in [1, k_n], j_i \in [1, t_i]$ and $(\alpha'_n)^\ell \geq \alpha_n \cdot \omega(\sqrt{\log n})^{1+\ell}$.*

Our second type of reductions can transfer the following hardness results on Decisional-LWE to Decisional-GLWE:

Theorem 7.2 ([31, 27]) *Let $n \in \mathbb{N}$ and $q_n \geq 1$ be a sequence of integers, and let $\alpha_n \in (0,1)$ be a real sequence such that $\alpha_n q_n \geq 2\sqrt{n}$. There exists a quantum reduction from worst-case n-dimensional* $\mathrm{GapSVP}_{\tilde{O}(n/\alpha_n)}$ *to Decisional-*$\mathrm{GLWE}_{\mathbb{Z}^n_{q_n},\alpha_n}$*. If $q_n \geq 2^{n/2}$ is smooth then there is a classical reduction between them.*

Theorem 7.3 ([9]) *Let $n \in \mathbb{N}$ and $q_n \geq 1$ be a sequence of integers, and let $\alpha_n \in (0,1)$ be a real sequence such that $\alpha_n \geq 2\sqrt{n}/2^{n/2}$. There exists a classical reduction from worst-case $\sqrt{n}$-dimensional* $\mathrm{GapSVP}_{\tilde{O}(\sqrt{n}/\alpha_n)}$ *to Decisional-*$\mathrm{GLWE}_{\mathbb{Z}^n_{q_n},\beta_n}$*, where $\beta_n^2 = 10n\alpha_n^2 + \frac{n}{q_n^2} \cdot \omega(\log n)$*

To do so, we reduce Decisional-LWE to Decisional-GLWE using what we call group switching. This technique transforms GLWE samples over one group G to samples over another group G', generalizing the modulus-dimension switching technique in [9], which is actually the special case $G = \mathbb{Z}_q^n$ and $G' = \mathbb{Z}_{q'}^{n'}$. We believe that the group switching technique proposed below is useful to better understand the core idea of the modulus-dimension switching technique.

Before presenting group switching, we note that the modulus-dimension switching technique from [9] implicitly uses a special case of structural lattice reduction. More precisely, Brakerski *et al.* [9] defined a special lattice Λ (see Th 3.1 of [9]) to transform LWE samples over $G = \mathbb{Z}_q^n$ to LWE samples over $G' = \mathbb{Z}_{q'}^{n'}$, but the meaning of Λ may look a bit mysterious. The lattice Λ is defined as $\Lambda = \frac{1}{q'}\mathbb{Z}^{n'} \cdot H + \mathbb{Z}^n$ where H is some $n' \times n$ integer matrix: this matrix is actually denoted by G in [9], but this would collide with our notation G for finite abelian groups. And [9] provided a good basis of Λ in special cases. We note that the exact definition of Λ is not important: whatever is the quotient $\Lambda/\mathbb{Z}^n$, it turns out to be isomorphic to the group $G' = \mathbb{Z}_{q'}^{n'}$, as shown by the transformation mapping $\frac{1}{q'}\mathbf{x} \cdot H + \mathbf{y} \in \Lambda$ to $\mathbf{x} \mod q' \in G'$. Thus, finding a good basis of Λ is actually a special case of structural lattice reduction for the lattice $\mathbb{Z}^n$ and the group G'. Therefore, it is natural to use structural lattice reduction directly (instead of an ad-hoc process) to obtain a more general statement than the modulus-dimension switching technique of [9].

Since we have two finite abelian groups G and G' and two overlattices $\bar{L}$ and $\bar{L}'$ of $\mathbb{Z}^n$, we will have two morphisms φ from $\bar{L}$ to G and φ' from $\bar{L}'$ to G' with $\ker(\varphi) = \ker(\varphi') = \mathbb{Z}^n$. Both morphisms are associated to their dual morphism as in Prop. 4.5, *i.e.* $\varphi^\times$ from $\mathbb{Z}^n$ to $\hat{G}$ and $\varphi'^\times$ from $\mathbb{Z}^n$ to $\hat{G}'$, satisfying $[\varphi'^\times(\mathbf{u})](\varphi'(\mathbf{v})) = \langle \mathbf{u}, \mathbf{v}\rangle \mod 1$ for all $\mathbf{u} \in \mathbb{Z}^n$ and all $\mathbf{v} \in \bar{L}'$ (resp. without the primes).

We say that a distribution S over $\mathbb{Z}^n$ is K-bounded if $\Pr_{\mathbf{s}\leftarrow S}[\|\mathbf{s}\| > K] \leq \mathrm{negl}(n)$. By extension, given a (public) morphism f from $\mathbb{Z}^n$ to $\hat{G}$, we say that a distribution $\mathcal{S}$ over $\hat{G}$ is K-bounded (for f) if it is the image of a K-bounded distribution[1] by f. In the following, we will choose $\varphi^\times = f$ and φ its dual morphism accordingly. Thus, any secret $\hat{s} \leftarrow \mathcal{S}$ has with overwhelming probability a preimage $\mathbf{s} \in \mathbb{Z}^n$ of norm $\leq K$. Note that the small $\mathbf{s} \in \mathbb{Z}^n$ may be hard to compute from $\hat{s}$, however in this case, what matters is its existence. During group switching, the new secret in $\hat{G}'$ will be $\varphi'^\times(\mathbf{s})$, and the new K-bounded distribution $\mathcal{S}' = \varphi'^\times(S)$.

Lemma 7.4 (Group Switching) *Let G and G' be two finite abelian groups of rank $\leq n$ s.t. G is fully-explicit and G' is explicit. Let $\bar{L}$ be an overlattice of $\mathbb{Z}^n$ such that $\bar{L}/\mathbb{Z}^n \simeq G$. Let $\bar{B}'$ be a basis of an overlattice $\bar{L}'$ of $\mathbb{Z}^n$ such that $\bar{L}'/\mathbb{Z}^n \simeq G'$. Let φ, φ' and $\varphi'^\times$ be defined as in Prop. 4.5. Let $r \geq \max\left(\sqrt{2}\eta_\varepsilon(\bar{L}), \|\bar{B}'^*\| \cdot \eta_\varepsilon(\mathbb{Z}^n)\right)$, where ε is some negligible function. Then, there is an efficient randomized algorithm which, given as input a sample from $G \times \mathbb{T}$, outputs a sample from $G' \times \mathbb{T}$, with the following properties:*
- If the input sample has uniform distribution in $G \times \mathbb{T}$, then the output sample has uniform distribution in $G' \times \mathbb{T}$ (except with negligible distance).
- If the input is distributed according to $A_{G,\alpha}(\hat{s})$ for some $\hat{s} = \varphi^\times(\mathbf{s})$ s.t. $\mathbf{s} \in \mathbb{Z}^n$ and $\|\mathbf{s}\| \leq K$, then the output distribution is statistically close to $A_{G',\beta}(\hat{s}')$, where $\hat{s}' = \varphi'^\times(\mathbf{s}) \in \hat{G}'$ and $\beta^2 = \alpha^2 + r^2(\|\mathbf{s}\|^2 + K^2) \leq \alpha^2 + 2(rK)^2$.

By combining the Group Switching Lemma 7.4 with the structural reduction Theorem 4.4, one obtains the following reduction between Decisional-GLWE with group G to Decisional-GLWE with group G':

[1] Ideally, f should be collision resistant among samples from S. In the classical LWE ($G = \mathbb{Z}_q^n$), f would map $\mathbf{s} \in \mathbb{Z}^n$ to the secret character $\hat{s} : \mathbf{y} \to {}^1\!/_q\langle \mathbf{s}, \mathbf{y}\rangle \mod 1$ in $\hat{G}$.

Corollary 7.5 (GLWE to GLWE) *Let $n \in \mathbb{N}$ and $0 < \sigma_n < 1$ be a real sequence. Let $(G_n)_{n\in\mathbb{N}}$ and $(G'_n)_{n\in\mathbb{N}}$ be two sequences of finite abelian groups with respective rank $k_n \leq n$ and $k'_n \leq n$ s.t. $\#G_n \geq n^{k_n}(\sqrt{2}/\sigma_n)^n$ (or if $G_n = \mathbb{Z}_{q_n}^n$ where $q_n \geq \sqrt{2}/\sigma_n$) and $\#G'_n \geq n^{k'_n}(1/\sigma_n)^n$. Assume that G_n is fully-explicit and G'_n is explicit. Let S be an arbitrary K_n-bounded distribution over $\mathbb{Z}^n$ and $\mathcal{S} = \varphi^\times(S)$ its image by some morphism $\varphi^\times : \mathbb{Z}^n \to \hat{G}_n$, $\alpha_n, \beta_n > 0$ be two real sequences and $\varepsilon = \mathrm{negl}(n)$ satisfying*

$$\beta_n^2 \geq \alpha_n^2 + 2(\sigma_n K_n \cdot \eta_\varepsilon(\mathbb{Z}^n))^2.$$

Then there is an efficient reduction from Decisional-$\mathrm{GLWE}_{G_n,\leq\alpha_n}(\mathcal{S})$ *to Decisional-*$\mathrm{GLWE}_{G'_n,\leq\beta_n}(\mathcal{S}')$*, where $\mathcal{S}' = \varphi'^\times(S)$ for some morphism $\varphi'^\times : \mathbb{Z}^n \to \hat{G}'_n$*

Proof. Given the canonical basis of $\mathbb{Z}^n$ and the group G_n, structural reduction finds an overlattice $\bar{L}$ together with a basis $\bar{C}$ s.t. $\left\|\bar{C}^*\right\| \leq \sigma_n/\sqrt{2}$. Therefore $\sqrt{2}\eta_\varepsilon(\bar{L}) \leq \sigma_n\eta_\varepsilon(\mathbb{Z}^n)$. And structural reduction on G'_n and σ_n gives a short basis $\bar{B}'$ of length $\leq \sigma_n$ and defines $\bar{L}'$. The rest of the proof follows immediately from Lemma 7.4. □

Using the normal form [3] of LWE, namely, if $\mathcal{S}$ is the image of the $\alpha_n q_n\sqrt{n}$-bounded distribution $\mathcal{D}_{\mathbb{Z}^n,\alpha_n q_n}$, through the canonical embedding which maps $\mathbf{s} \in \mathbb{Z}^n$ to the character $\hat{s} = \mathbf{y} \to {}^1/_{q_n}\langle \mathbf{s}, \mathbf{y}\rangle \bmod 1$, we obtain the quantum/classical hardness of Decisional-GLWE problem for any sufficiently large finite Abelian group, together with Theorems 7.2 and 7.3:

Corollary 7.6 (Quantum Hardness of GLWE) *Let $n \in \mathbb{N}$ and $q_n \geq 1$ be a sequence of integers and $(G'_n)_{n\in\mathbb{N}}$ be a sequence of any finite abelian explicit groups such that $\#G'_n \geq n^{k_n}(q_n/\sqrt{2})^n$ where $k_n = \mathrm{rank}\,(G'_n) \leq n$. Let $\alpha_n, \beta_n \in (0,1)$ be two real sequences such that $\alpha_n q_n \geq 2\sqrt{n}$ and $\beta_n = \alpha_n\sqrt{n}\cdot\omega(\sqrt{\log n})$. Then there exists a quantum reduction from worst-case n-dimensional* $\mathrm{GapSVP}_{\tilde{O}(n/\alpha_n)}$ *to Decisional-*$\mathrm{GLWE}_{G'_n,\beta_n}$.

The lower bound on $\#G'_n$ is better than for $\#G_n$ in Cor. 6.2 and 5.2, because group switching relies on structural reduction over $\mathbb{Z}^n$, rather than over an arbitrary lattice: the canonical basis of $\mathbb{Z}^n$ is orthonormal, which simplifies the bound of Sect. 4.

Corollary 7.7 (Classical Hardness of GLWE) *Let $n \in \mathbb{N}$ and $q_n \geq 1$ be a sequence of integers and $(G'_n)_{n\in\mathbb{N}}$ be a sequence of any finite abelian explicit groups such that $\#G'_n \geq n^{k_n}(q_n/\sqrt{2})^n$ where $k_n = \mathrm{rank}\,(G'_n) \leq n$. Let $\alpha_n, \beta_n \in (0,1)$ be two real sequences such that $\alpha_n \geq 2\sqrt{n}/2^{n/2}$ and $\beta_n^2 = n^2\alpha_n^2 \cdot \omega(\log n) + \frac{n^2}{q_n^2}\cdot\omega(\log^2 n)$. There exists a classical reduction from worst-case $\sqrt{n}$-dimensional* $\mathrm{GapSVP}_{\tilde{O}(\sqrt{n}/\alpha_n)}$ *to Decisional-*$\mathrm{GLWE}_{G'_n,\beta_n}$.

References

[1] M. Ajtai. Generating hard instances of lattice problems. In *STOC*, pages 99–108, 1996.

[2] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proc. of 29th STOC*, pages 284–293. ACM, 1997. Available at [11] as TR96-065.

[3] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proc. of Crypto '09*, LNCS 5677, pages 595–618. IACR, Springer-Verlag, 2009.

[4] L. Babai. On Lovász lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986.

[5] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993.

[6] G. Baumslag, N. Fazio, A. Nicolosi, V. Shpilrain, and W. E. Skeith. Generalized learning problems and applications to non-commutative cryptography. In *Proc. ProvSec '11*, volume 6980 of *Lecture Notes in Computer Science*, pages 324–339. Springer, 2011.

[7] A. Becker, N. Gama, and A. Joux. Solving shortest and closest vector problems: The decomposition approach. Cryptology ePrint Archive, Report 2013/685, 2013. `http://eprint.iacr.org/`.

[8] D. Boneh and R. Venkatesan. Hardness of computing the most significant bits of secret keys in Diffie-Hellman and related schemes. In *Proc. of Crypto '96*, LNCS. IACR, Springer-Verlag, 1996.

[9] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D.Stehlé. Classical hardness of learning with errors. In *Proc. of 45th STOC*, pages 575–584. ACM, 2013.

[10] Y. Chen and P. Q. Nguyen. BKZ 2.0: Better lattice security estimates. In *Proc. of Asiacrypt*, pages 1–20, 2011.

[11] ECCC. `http://www.eccc.uni-trier.de/eccc/`. The Electronic Colloquium on Computational Complexity.

[12] N. Gama and P. Q. Nguyen. Predicting Lattice Reduction. In *Proc. of Eurocrypt*, 2008.

[13] C. Gentry and S. Halevi. Implementing Gentry's fully-homomorphic encryption scheme. In *Advances in Cryptology - Proc. EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 129–148. Springer, 2011.

[14] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.

[15] D. Goldstein and A. Mayer. On the equidistribution of Hecke points. *Forum Math.*, 15(2):165–189, 2003.

[16] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Ann.*, 261:513–534, 1982.

[17] A. Lubotzky. The expected number of random elements to generate a finite group. *J. Algebra*, 257(2):452–459, 2002.

[18] D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM J. Comput.*, 34(1):118–169, 2004.

[19] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proc. EUROCRYPT '12*, LNCS. IACR, Springer-Verlag, 2012.

[20] D. Micciancio and C. Peikert. Hardness of SIS and LWE with small parameters. In *Proc. CRYPTO '13*, volume 8042 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2013.

[21] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. In *SIAM J. Comput.*, 2007.

[22] L. J. Mordell. On some arithmetical results in the geometry of numbers. *Compositio Mathematica*, 1:248–253, 1935.

[23] P. Q. Nguyen and I. E. Shparlinski. The insecurity of the digital signature algorithm with partially known nonces. *J. Cryptology*, 15(3):151–176, 2002.

[24] P. Q. Nguyen and I. E. Shparlinski. Counting co-cyclic lattices. Preprint, 2013.

[25] I. Pak. On probability of generating a finite group. Preprint, 1999.

[26] A. Paz and C.-P. Schnorr. Approximating integer lattices by lattices with cyclic factor groups. In *Proc. of ICALP*, pages 386–393, 1987.

[27] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC '09*, pages 333–342. ACM, 2009.

[28] C. Peikert. An efficient and parallel gaussian sampler for lattices. In *Proc. of Crypto '10*, LNCS 6223, pages 80–97. Spinger-Verlag, 2010.

[29] M. Pohst. A modification of the LLL reduction algorithm. *J. Symbolic Comput.*, 4(1):123–127, 1987.

[30] O. Regev. Lattices in computer science #12: Average-case hardness. Regev's webpage, 2004.

[31] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.

A Missing Proofs of Sect. 2

In order to ease the proofs of indistinguishability between two distributions, we introduce an approximate equality which will also be used in App. E: for $\varepsilon \in]0, \frac{1}{2}]$ and $n, m \in \mathbb{N}$, the notation $a =_{\varepsilon,n,m} b$ means $b \cdot \frac{(1-\varepsilon)^n}{(1+\varepsilon)^m} \leq a \leq b \cdot \frac{(1+\varepsilon)^n}{(1-\varepsilon)^m}$, which implies $b =_{\varepsilon,m,n} a$ and also $a = b(1 + (m+n)\varepsilon + O(\varepsilon^2))$. If $a =_{\varepsilon,n,m} b$ and $c =_{\varepsilon,n',m'} d$, we then have $ac =_{\varepsilon,n+n',m+m'} bd$.

With this notation, the main property of the smoothing parameter is that for all lattice L, $\mathbf{c} \in \operatorname{span}(L)$, $\sigma \geq \eta_\varepsilon(L)$, we have that $\rho_{\mathbb{R}^n,\sigma}(\mathbf{c} + L) =_{\varepsilon,1,0} 1/\operatorname{vol}(L)$. Thus for all overlattice $\bar{L} \supseteq L$ and all $\mathbf{c} \in \bar{L}$, $\rho_{\bar{L}/L,\sigma}(\mathbf{c}) =_{\varepsilon,1,1} \frac{\operatorname{vol}(\bar{L})}{\operatorname{vol}(L)}$

A.1 Proof of Prop. 2.2

We know that for $\varepsilon = 1/10$, independent Gaussian samples $(\mathbf{y}_i \leftarrow_\varepsilon \mathcal{D}_{L,s_i})$ such that $\sqrt{2}\eta_\varepsilon(L) \leq s_i$ have probability $\leq 9/10$ to be in any fixed hyperplane (see [30, Lemma 14]). This can be adapted to any $\varepsilon > 0$, so that one can extract a full-rank family F of n vectors of norm $\|F\| \leq \sqrt{n/2\pi} \cdot \max s_i$.

Then the generalized LLL algorithm for linearly dependent vectors [29] $F \cup B$ returns a basis C of length $\|C^*\| \leq \sqrt{n/2\pi} \cdot \max s_i$ in polynomial time.

A.2 Proof of Lemma 2.4 on Discrete Convolution

We now prove the dot product convolution Lemma.

The proof relies on the following equality: For $\alpha, \sigma \in \mathbb{R}_{ge0}$, for $\gamma = \left(\frac{1}{\sigma^2} + \frac{u^2}{\alpha^2}\right)^{-1/2}$ and $\Gamma = \sqrt{\alpha^2 + \sigma^2 u^2}$

$$\begin{aligned}&\frac{1}{\sigma^2}x^2 + \frac{1}{\alpha^2}(t - ux)^2 \\ =&\frac{1}{\gamma^2}\left(x - \frac{\gamma^2 tu}{\alpha^2}\right)^2 + \frac{1}{\Gamma^2}t^2\end{aligned}$$

Let $C = \mathbf{z} + L$ be some coset of a n-dimensional lattice L, $\mathbf{u} \in \mathbb{R}^n$, $\alpha, \sigma \in \mathbb{R}_{ge0}$ and $\varepsilon \in (0, 1/2)$. Let $(\mathbf{e}_1, \ldots, \mathbf{e}_n)$ be an orthonormal basis of $\mathbb{R}^n$ such that $\mathbf{u} = u \cdot \mathbf{e}_n$. A vector $\mathbf{v} \in \mathbb{R}^n$ will be expressed as $\sum_{i=1}^n v_i \mathbf{e}_i$.

$$\begin{aligned}\sum_{\mathbf{v}\in C} \mathcal{D}_{\mathbb{R}^n,\sigma}(\mathbf{v})\mathcal{D}_{\mathbb{R},\alpha}(t - \langle \mathbf{u}, \mathbf{v}\rangle) &= \sum_{\mathbf{v}\in C} \mathcal{D}_{\mathbb{R},\sigma}(v_1)\ldots\mathcal{D}_{\mathbb{R},\sigma}(v_n)\mathcal{D}_{\mathbb{R},\alpha}(t - uv_n) \\ &= \sum_{\mathbf{v}\in C} \mathcal{D}_{\mathbb{R},\sigma}(v_1)\ldots\mathcal{D}_{\mathbb{R},\sigma}(v_{n-1})\frac{1}{\sigma\alpha}\exp\left(-\pi\left(\frac{1}{\sigma^2}v_n^2 + \frac{1}{\alpha^2}(t - uv_n)^2\right)\right) \\ &= \frac{1}{\sigma^n\alpha}\sum_{\mathbf{v}\in C}\exp\left(-\pi\left(\frac{1}{\sigma}v_1^2 + \cdots + \frac{1}{\sigma}v_{n-1}^2 + \frac{1}{\gamma^2}\left(v_n - \frac{\gamma^2 u}{\alpha^2}t\right)^2 + \frac{1}{\Gamma^2}t^2\right)\right)\end{aligned}$$

Let f be the affine function which maps $\sum_{i=1}^n v_i\mathbf{e}_i$ to $\frac{v_1}{\sigma}\mathbf{e}_1 + \cdots + \frac{v_{n-1}}{\sigma}\mathbf{e}_{n-1} + \frac{v_n - \gamma^2 ut/\alpha^2}{\gamma}\mathbf{e}_n$. Then,

$$\begin{aligned}\sum_{\mathbf{v}\in C} \mathcal{D}_{\mathbb{R}^n,\sigma}(\mathbf{v})\mathcal{D}_{\mathbb{R},\alpha}(t - \langle \mathbf{u}, \mathbf{v}\rangle) &= \frac{1}{\sigma^n\alpha}\sum_{\mathbf{v}\in C}\exp\left(-\pi\left(\|f(\mathbf{v})\|^2 + \frac{1}{\Gamma^2}t^2\right)\right) \\ &= \frac{1}{\sigma^n\alpha}\sum_{\mathbf{v}'\in f(C)}\exp\left(-\pi\left(\|\mathbf{v}\|^2 + \frac{1}{\Gamma^2}t^2\right)\right) \\ &= \frac{\Gamma}{\sigma^n\alpha}\mathcal{D}_{\mathbb{R}^n,1}(f(C))\mathcal{D}_{\mathbb{R},\Gamma}(t)\end{aligned}$$

Note that the largest eigenvalue of the linear part of f is $1/\gamma$, thus since $C = \mathbf{z} + L$, $f(C) = \mathbf{z}' + L'$ where $\eta_\varepsilon(L') \leq \eta_\varepsilon(L)/\gamma \leq 1$. Therefore, $\mathcal{D}_{\mathbb{R}^n,1}(f(C)) =_{\varepsilon,1,1} 1/\operatorname{vol}(L') = \sigma^{n-1}\gamma/\operatorname{vol}(L)$. We finally obtain:

$$\sum_{\mathbf{v}\in C} \mathcal{D}_{\mathbb{R}^n,\sigma}(\mathbf{v})\mathcal{D}_{\mathbb{R},\alpha}(t - \langle \mathbf{u}, \mathbf{v}\rangle) =_{\varepsilon,1,1} \mathcal{D}_{\mathbb{R},\Gamma}(t)$$

Now, we can prove the lemma. First, let $\mathbb{K} = \mathbb{R}$ and $b \leftarrow \mathcal{D}_{\mathbb{K},\alpha,c+\langle\mathbf{u},\mathbf{v}\rangle}$ where $\mathbf{v} \leftarrow \mathcal{D}_{C,\sigma}$. Then the density of b is $\sum_{v\in C}\mathcal{D}_{\mathbb{K},\alpha}(b - c - \langle\mathbf{u},\mathbf{v}\rangle)\mathcal{D}_{C,\sigma}(\mathbf{v}) =_{\varepsilon,1,1} \mathcal{D}_{\mathbb{K},\Gamma}(b - c) = \mathcal{D}_{\mathbb{K},\Gamma,c}(b)$.

Second, let $\mathbb{K} = \frac{1}{N}\mathbb{Z}$, but assume that $\alpha \geq \eta_\varepsilon(\mathbb{K})$. the density of $b \in \mathbb{K}$ is

$$\sum_{\mathbf{v}\in C} \frac{\mathcal{D}_{\mathbb{R},\alpha}(b - c - \langle \mathbf{u}, \mathbf{v}\rangle)}{\mathcal{D}_{\mathbb{R},\alpha}(\mathbb{K} - c - \langle \mathbf{u}, \mathbf{v}\rangle)} \mathcal{D}_{C,\sigma}(\mathbf{v})$$

Since the denominator is $=_{\varepsilon,1,0} 1/\operatorname{vol}(\mathbb{K})$, the whole expression is nearly equal to $N\mathcal{D}_{\mathbb{R},\Gamma,c}(b)$. Since $\Gamma \geq \alpha \geq \eta_\varepsilon(\mathbb{K})$, then the N can be viewed as $1/\mathcal{D}_{\mathbb{R},\Gamma,c}(\mathbb{K})$, and finally, the density of b is $\mathcal{D}_{\mathbb{K},\Gamma,c}(b)$.

B Addendum on Sect. 3

B.1 Proof of Th. 3.1

Let $L \in \mathcal{L}_{G,m}$. Then G has rank $\leq m$ because $L \subseteq \mathbb{Z}^m$. And G is isomorphic to some product $\mathbb{Z}_{q_1} \times \cdots \times \mathbb{Z}_{q_k}$ of cyclic groups, where $1 \leq k \leq m$, $q_i \geq 1$ and q_{i+1} divides q_i for all i. By [26], there exist primitive vectors $\mathbf{z}_1, \ldots, \mathbf{z}_k \in \mathbb{Z}^m$ s.t. $L = \{\mathbf{y} \in \mathbb{Z}^m,\ \langle \mathbf{y}, \mathbf{z}_i\rangle \equiv 0 \pmod{q_i}, i \in [1,k]\}$. This shows that there exists $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$ generating G such that $L = \mathcal{L}_\mathbf{g}$, where we recall that $\mathcal{L}_\mathbf{g} = \{(x_1, \ldots, x_m) \in \mathbb{Z}^m \text{ s.t. } \sum_{i=1}^m x_i g_i = 0 \text{ in } G\}$.

Reciprocally, let $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$ generate G. Consider the morphism ψ which maps $(x_1, \ldots, x_m) \in \mathbb{Z}^m$ to $\sum_{i=1}^m x_i g_i \in G$. By definition, the image of ψ is G and its kernel is $\mathcal{L}_\mathbf{g}$, therefore $\mathbb{Z}^m/\mathcal{L}_\mathbf{g} \simeq G$.

B.2 Proof of Lemma 3.2

Let $\mathbf{g} = (g_1, \ldots, g_m) \in G^m$ be such that the g_i's generate G. Let $\mathbf{h} = (h_1, \ldots, h_m) \in G^m$.

Assume that $\mathcal{L}_\mathbf{g} = \mathcal{L}_\mathbf{h}$. Define a map $\psi : G \to G$ as follows: for any $g \in G$, there exists a decomposition $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{Z}^m$ s.t. $g = \sum_{i=1}^m x_i g_i$, and we let $\psi(g) = \sum_{i=1}^m x_i h_i$. This map is well-defined because if there are two decompositions of g, *i.e.* $g = \sum_{i=1}^m x_i g_i = \sum_{i=1}^m y_i g_i$, then $\mathbf{x} - \mathbf{y} \in \mathcal{L}_\mathbf{g} = \mathcal{L}_\mathbf{h}$, thus $\sum_{i=1}^m x_i h_i = \sum_{i=1}^m y_i h_i$ and $\psi(g)$ has the same value. It can be checked that ψ is a morphism. Since $\mathbb{Z}^m/\mathcal{L}_\mathbf{g} \simeq \mathbb{Z}^m/\mathcal{L}_\mathbf{h}$, we know that the h_i's generate G, and therefore ψ is an automorphism of G.

Reciprocally, let ψ be an automorphism of G such that $h_i = \psi(g_i)$ for all $1 \leq i \leq m$. Then $\psi(\sum_{i=1}^m x_i g_i) = \sum_{i=1}^m x_i h_i$ for any $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{Z}^m$. It follows that $\sum_{i=1}^m x_i g_i = 0$ if and only if $\sum_{i=1}^m x_i h_i = 0$. Hence, $\mathcal{L}_\mathbf{g} = \mathcal{L}_\mathbf{h}$.

Finally, unicity follows from $\psi(\sum_{i=1}^m x_i g_i) = \sum_{i=1}^m x_i h_i$.

B.3 Proof of Lemma 3.4

We prove the decisional version here, and the search version is analogous. Given the Decisional-GLWE$_{G,m',\alpha}(\mathcal{S})$ oracle (m' to be specified later), we now construct an efficient algorithm B to distinguish m samples $(a_i, b_i)_{i\in[1,m]}$ either from $A_{G,\beta}(\hat{s})$ or uniform in $G \times \mathbb{T}$ for some unknown β ($\leq \alpha$) and secret $\hat{s}$ sampled from $\mathcal{S}$. Let Z be the set of integer multiples of $\frac{1}{m^2}\alpha^2$ between 0 and α^2. For each $z \in Z$, B does the following. B picks m uniform samples $(\tilde{a}_i, \tilde{b}_i)_{i\in[1,m]}$ from $G \times \mathbb{T}$ and receives m samples $(a_i, b_i)_{i\in[1,m]}$ from his challenger. B estimates the acceptance probability of the Decisional-GLWE$_{G,m',\alpha}(\mathcal{S})$ oracle on the following two inputs: The first input is of the form $(\tilde{a}_i, \tilde{b}'_i)_{i\in[1,m]}$, where $\tilde{b}'_i \leftarrow \mathcal{D}_{\mathbb{T},z,\tilde{b}_i}$ and the second input is of the form $(a_i, b'_i)_{i\in[1,m]}$, where $b'_i \leftarrow \mathcal{D}_{\mathbb{T},z,b_i}$. If in any of these polynomial attempts a non-negligible difference is observed between two acceptance probabilities, output "non-uniform"; otherwise, output "uniform".

Note that $(\tilde{a}_i, \tilde{b}'_i)_{i\in[1,m]}$ is uniformly random in $G \times \mathbb{T}$. If $(a_i, b_i)_{i\in[1,m]}$ is uniformly random in $G \times \mathbb{T}$, then the two acceptance probabilities are exactly the same. If $(a_i, b_i)_{i\in[1,m]}$ is distributed as $A_{G,\beta}(\hat{s})$, then by classical convolution, $(a_i, b'_i)_{i\in[1,m]}$ is distributed as $A_{G,\sqrt{\beta^2+z}}(\hat{s})$. Consider the smallest $z \in Z$ such that $z \geq \alpha^2 - \beta^2$. Clearly, $z \leq \alpha^2 - \beta^2 + \frac{1}{m^2}\alpha^2$. Then

$$\alpha \leq \sqrt{\beta^2 + z} \leq \sqrt{\alpha^2 + \frac{1}{m^2}\alpha^2} \leq (1 + \frac{1}{m^2})\alpha.$$

Therefore, the statistical distance between $\mathcal{D}_{\mathbb{T},\alpha,\hat{s}(a_i)}$ and $\mathcal{D}_{\mathbb{T},\sqrt{\beta^2+z},\hat{s}(a_i)}$ is at most $O(\frac{1}{m^2})$ for $i \in [1,m]$. The statistical distance of m samples from $A_{G,\sqrt{\beta^2+z}}(\hat{s})$ and m samples from $A_{G,\alpha}(\hat{s})$ is at most $O(\frac{1}{m})$. Hence, for our choice of z, and by Chernoff bound, non-negligible $(O(\frac{1}{m}))$ difference will be observed with probability $\geq \frac{1}{3}$, if $(a_i, b_i)_{i\in[1,m]}$ is distributed as $A_{G,\beta}(\hat{s})$. Notice that we can set $m' = m^3$ in the Decisional-GLWE oracle.

B.4 Discretization of GLWE

In this subsection, we discuss discrete versions of the GLWE problem, where the LWE samples (a, b) are not taken in $G \times \mathbb{T}$, but b is instead chosen from a discrete subset of $\mathbb{T}$.

The first option is to use the rounded Gaussian distributions, which is suitable for a floating point representation. By convention, the distance between two numbers $x, y \in \mathbb{T}$ is $\min_{z\in\mathbb{Z}}(|x - y + z|)$. Let $h_1, \ldots, h_p$ be p real numbers such that $0 \leq h_1 < \cdots < h_p < 1$. We denote by H the values $h_1, \ldots, h_p$ mod 1, which is a finite subset of $\mathbb{T}$. We define the Rounded Discrete GLWE distribution denoted by $Ar_{G,H,\alpha}(\hat{s})$ the distribution of tha pair (a, b) over $G \times H$ where a is uniformly random in G and b is sampled according to $\mathcal{D}_{\mathbb{T},\alpha,\hat{s}(a)}$ and rounded to the nearest value over H. For the decisional variant, the uniform distribution of b over $\mathbb{T}$ is replaced by the distribution over H where b is sampled uniformly at random over $\mathbb{T}$ and rounded to its nearest value in H. With this definition, it is clear that starting from continuous GLWE (or uniform) samples $(a, b) \in G \times \mathbb{T}$, it suffices to take $a' = a$ and round b to its nearest value $b' \in H$ to obtain a discrete and rounded sample (a', b'). We denote by (Search) Decisional-GLWE$_r(G, H, \alpha)(\mathcal{S})$ the corresponding problems. If an oracle solves Decisional-GLWE$_r(G, H, \alpha)(\mathcal{S})$ (resp. Search), it automatically solves the underlying continuous Decisional-GLWE$_{G,\alpha}(\mathcal{S})$ (resp. Search) instance (provided that the solution remains unique). Reciprocally, one can turn a discrete rounded GLWE sample into a continuous one by adding some Gaussian noise larger than the maximal distance between two consecutive values in H:

Lemma B.1 *Let $h_1, \ldots, h_p$ be p real numbers such that $0 \leq h_1 < \cdots < h_p < 1$ and H their representatives in $\mathbb{T}$. By convention, we set $h_{p+1} = 1 + h_1$. For all parameter β such that $\beta \geq \sqrt{2}\max_{i\in[1,p]}(h_{i+1} - h_i)$. Then there is a reduction from Decisional-*GLWE$_r(G, H, \alpha)(\mathcal{S})$ *(resp. Search) to Decisional-*GLWE$_{G,\sqrt{\alpha^2+\beta^2}}(\mathcal{S})$ *(resp. Search) for any distribution $\mathcal{S}$ over $\hat{G}$.*

The second option is to discretize GLWE over a finite subgroup $\mathbb{K} = \frac{1}{N}\mathbb{Z}/\mathbb{Z}$ of the torus using the discrete Gaussian distribution. For $\beta > 0$ and some positive integer N, we denote $\bar{A}_{G,\beta,N}(\hat{s})$ the distribution over $G \times \mathbb{K}$ which chooses $a \leftarrow G$ uniformly at random, sets $b \leftarrow \mathcal{D}_{\mathbb{K},\beta,\hat{s}(a)}$ and outputs (a, b). We call (Search) Decisional-DGLWE$_{G,\alpha,N}(\mathcal{S})$ this discretization.

Again, we show that the discrete version is at least as hard as the continuous version for some suitable parameters:

Lemma B.2 *Let G be any finite abelian group and $N > 0$ an integer. Let $0 < \alpha, \beta < 1$ be reals such that $\beta \geq \eta_\varepsilon(\frac{1}{N}\mathbb{Z})$ for some negligible function ε. Then there is a reduction from Decisional-*GLWE$_{G,\alpha}(\mathcal{S})$ *(resp. Search) to Decisional-*DGLWE$_{G,\sqrt{\alpha^2+\beta^2},N}(\mathcal{S})$ *(resp. Search) for any distribution $\mathcal{S}$ over $\hat{G}$.*

Proof. The reduction does the following: given a sample $(a, b) \in G \times \mathbb{T}$, it sets $a' = a$ and samples $b' \leftarrow D_{\mathbb{K},\beta,b}$. If the distribution of (a, b) is $A_{G,\alpha}(\hat{s})$, then $b \leftarrow D_{\mathbb{T},\alpha,\hat{s}(a)}$. Since $\beta \geq \eta_\varepsilon(\frac{1}{N}\mathbb{Z})$, the distribution of b' is statistically close to $\mathcal{D}_{\mathbb{K},\sqrt{\alpha^2+\beta^2},\hat{s}(a)}$ by simple convolution. If (a, b) is uniformly random over $G \times \mathbb{T}$, then b is uniformly random over $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ and independent of a. We obtain that b' is uniformly random over $\mathbb{K}$. □

C Missing Proofs of Sect. 4

C.1 Proof of Th. 4.2

Let B be an LLL-reduced basis with factor ε_{LLL} of an n-dimensional lattice $L \subseteq \mathbb{R}^n$. Let $\alpha = (1 + \varepsilon_{\text{LLL}})\sqrt{4/3}$.

Let $x_1, \dots, x_n$ denote the $\|\mathbf{b}_i^*\|$'s ordered by decreasing value. Let $k = \min \{i \in [1,n] \text{ st. } x_i/x_{i+1} > (1+\varepsilon_{\text{LLL}})\sqrt{4/3}\}$, where $x_{n+1} = 0^+$.

Th. 4.2 follows from the following inequalities:

$$\mathrm{bl}(L) \geq \left(\prod_{i=1}^{k} x_i\right)^{1/k}. \tag{5}$$

$$\delta_{\mathrm{bl}(L)}(B) \leq \alpha^{\frac{n^2}{8}+\frac{n^2}{4}} \tag{6}$$

$$\delta_{\alpha^r \,\mathrm{bl}(L)}(B) \leq \alpha^{\frac{(n-2r)^2}{8}+\frac{(n-2r)^2}{4}} \tag{7}$$

Proof of (5) The proof follows from the following two facts: First, we have: $\mathrm{bl}(L) \geq \left(\prod_{j=n+1-i}^{n} \|\mathbf{b}_j^*\|\right)^{1/i}$ for all $i \in [1,n]$. Indeed, consider the projection π_{n+1-i} over $(\mathbf{b}_1, \dots, \mathbf{b}_{n-i})^{\perp}$. Then $\mathrm{bl}(L) \geq \mathrm{bl}(\pi_{n+1-i}(L))$ because projections cannot increase Gram-Schmidt norms, and $\mathrm{bl}(\pi_{n+1-i}(L)) \geq \mathrm{vol}(\pi_{n+1-i}(L)) = \left(\prod_{j=n+1-i}^{n} \|\mathbf{b}_j^*\|\right)^{1/i}$ because $B_{[n+1-i,n]}$ is a basis of $\pi_{n+1-i}(L)$. Second, let $A = \{i \in [1,n] \text{ s.t } \|\mathbf{b}_i^*\| \geq x_k\}$. By definition of k, $i \notin A \implies \|\mathbf{b}_i^*\| < x_k/((1+\varepsilon_{\text{LLL}})\sqrt{4/3})$. Therefore, Lovász' condition implies that for all $i \in A$, $i+1 \in A$. Thus, A is necessarily the right-most integer interval with k elements, *i.e.* $[n+1-k, n]$ and $\prod_{i=1}^{k} x_i = \prod_{i=n+1-k}^{n} \|\mathbf{b}_i^*\|$. □

Proof of (6) Let $\sigma_0 = \left(\prod_{i=1}^{k} x_i\right)^{1/k}$ and $j = \max\{i \in [1,k],\ x_i \geq \sigma_0\}$. Note that $\delta_{\sigma_0}(B) = \prod_{l=1}^{j} \frac{x_l}{\sigma_0} = \prod_{l=j+1}^{k} \frac{\sigma_0}{x_l}$. If $j \leq k/2$, then $x_l/\sigma_0 \leq \alpha^{j+1-l}$ for all $l \leq j$, therefore $\delta_{\sigma_0}(B) \leq \alpha^{(1+\ldots+j)} \leq \alpha^{\frac{k}{4}(\frac{k}{2}+1)}$. If $j > k/2$, $\sigma/x_l \leq \alpha^{l-j}$ for all $l \geq j$, therefore $\delta_{\sigma}(B) \leq \alpha^{(1+\ldots+(k-j))} \leq \alpha^{\frac{k}{4}(\frac{k}{2}+1)}$. In all cases, $\delta_{\sigma_0}(B) \leq \alpha^{\frac{k}{4}(\frac{k}{2}+1)} \leq \alpha^{\frac{n}{4}(\frac{n}{2}+1)}$. Finally, the cubicity-defect decreases with σ: since $\mathrm{bl}(L) \geq \sigma_0$, $\delta_{\mathrm{bl}(L)}(B) \leq \delta_{\sigma_0}(B)$. □

Proof of (7) Assume by contradiction that $\delta_{\alpha^r \,\mathrm{bl}(L)} > \alpha^{(\frac{n}{2}-r)+\cdots+2+1}$. Let $j = \max\{i \text{ s.t. } x_i \geq \alpha^r \,\mathrm{bl}(L)\}$, since $\delta_{\alpha^r \,\mathrm{bl}(L)} \leq \frac{x_j}{\alpha^r \,\mathrm{bl}(L)} \cdots \frac{x_1}{\alpha^r \,\mathrm{bl}(L)} \leq \alpha\alpha^2 \ldots \alpha^j$ then $j > \frac{n}{2} - r$. Thus

$$\begin{aligned}
&\delta_{\mathrm{bl}(L)}(B) \\
&\geq \prod_{i=1}^{j} \frac{x_i}{\mathrm{bl}(L)} \prod_{i=j+1}^{j+r} \frac{x_i}{\mathrm{bl}(L)} \\
&\geq \delta_{\alpha^r \,\mathrm{bl}(L)} \alpha^{rj} \alpha^{r-1} ... \alpha^1 \\
&> \alpha^{(\frac{n}{2}-r+r)+\cdots+(1+r)+(0+r)} \alpha^{(r-1)+\cdots+2+1} \\
&> \alpha^{\frac{n}{2}+\cdots+1}
\end{aligned}$$

This contradicts (6), thus $\delta_{\alpha^r \,\mathrm{bl}(L)} \leq \alpha^{\frac{(n-2r)^2}{8}+\frac{(n-2r)}{4}}$. □

C.2 Proof of Theorem 4.1 and Alg. 2

Let $a_i = \max(1, \|\mathbf{b}_i^*\|/\sigma)$ for $i \in [1,n]$. For each i from $k-1$ downto 1, we use the suffix "old" and "new" to denote respectively the values of the variables at the beginning and at the end of the "for" loop (line 3 of Alg. 2). Furthermore, we call x_i the value $\|\mathbf{c}_i^{*\text{new}}\|$ during iteration i. Note that x_i is also $\|\mathbf{c}_i^{*\text{old}}\|$ during the next iteration (of index $i-1$ since i goes backwards).

We show by induction over i that the following invariant holds at the end of each iteration (line 3 of Alg. 2):

$$a_i x_{i+1} \leq x_i \leq a_i x_{i+1} + \sigma a_i \tag{8}$$

At the first iteration ($i = k-1$), it is clear that $x_k = \|\mathbf{c}_k^{*\text{old}}\| = \sigma a_k$. At the beginning of iteration i (line 3), there are two cases:

1. if $\left\|\mathbf{c}_i^{*\text{old}}\right\| \leq \sigma$ and $\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| > \sigma$ (line 9), we size-reduce and swap the two vectors, so that $\|\mathbf{c}_i^{*\text{new}}\|$ satisfies:
$$\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| \leq \|\mathbf{c}_i^{*\text{new}}\| \leq \left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| + \sigma.$$
Since $a_i = 1$, $x_{i+1} = \left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|$ and $x_i = \|\mathbf{c}_i^{*\text{new}}\|$, the invariant (8) holds.

2. If $\left\|\mathbf{c}_i^{*\text{old}}\right\| > \sigma$ and $\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| > \sigma$ (line 7), we transform the block so that the norm of the first vector satisfies
$$R \leq \|\mathbf{c}_i^{*\text{new}}\| \leq R + \left\|\mathbf{c}_i^{*\text{old}}\right\|. \qquad (9)$$
$$\text{where } R = \left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| \left\|\mathbf{c}_i^{*\text{old}}\right\| / \sigma$$
This condition can always be fulfilled with a primitive vector of the form $\mathbf{c}_i^{\text{new}} = \mathbf{c}_{i+1}^{\text{old}} + \alpha \mathbf{c}_i^{\text{old}}$ for some $\alpha \in \mathbb{Z}$. Since the volume is invariant, the new $\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\|$ is upper-bounded by σ. And by construction, Equation (9) is equivalent to the invariant (8) since $\left\|\mathbf{c}_i^{*\text{old}}\right\| = a_i\sigma$, $\|\mathbf{c}_i^{*\text{new}}\| = x_i$ and $\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| = x_{i+1}$.

By expanding, this invariant implies that

$$\begin{aligned} x_1 &\leq \sigma \sum_{i=1}^{k} a_1 \dots a_i \\ &= \sigma \sum_{i=1}^{k} \delta_\sigma(B_{[1,i]}) \\ &\leq n\sigma\delta_\sigma(B) \end{aligned}$$

Note that the transformation matrix of the unbalanced reduction algorithm is

$$\left[\begin{array}{cccc|ccc} \alpha_1 & \cdots & \alpha_{k-1} & 1 & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 & \vdots & & \vdots \\ 0 & \ddots & \ddots & \vdots & \vdots & & \vdots \\ 0 & 0 & 1 & 0 & 0 & \cdots & 0 \\ \hline 0 & \cdots & \cdots & 0 & 1 & 0 & 0 \\ \vdots & & & \vdots & 0 & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & 0 & 0 & 1 \end{array}\right]$$

where α_i is either $\lfloor -\mu_{i+1,i} \rceil$ or $\left\lceil -\mu_{i+1,i} + \frac{x_{i+1}}{\sigma}\sqrt{1 - \frac{1}{a_i^2}} \right\rceil$. Since each x_{i+1} is bounded by $\delta_\sigma(B_{[i+1,n]})$, all coefficients have a size polynomial in the input basis and the overall complexity is therefore polynomial.

Now let us show (1). It suffices to prove that the following invariant holds at the beginning of each iteration:

$$\forall\ \nu \leq \sigma, \qquad \delta_\nu\left(C_{[i,n]}^{\text{old}}\right) \leq \delta_\nu\left(C_{[i,n]}^{\text{new}}\right) \leq \frac{\|\mathbf{c}_i^{*\text{new}}\|}{\sigma\delta_\sigma\left(C_{[i,n]}^{\text{old}}\right)} \times \delta_\nu\left(C_{[i,n]}^{\text{old}}\right) \qquad (10)$$

Since all $\left\|\mathbf{c}_j^{*\text{old}}\right\| \leq \sigma$ for $j = i+2, \dots, n$, $\delta_\sigma\left(C_{[i,n]}^{*}\right) = \delta_\sigma\left(C_{[i,i+1]}^{*}\right)$, where (*) is either old or new. Hence, showing (10) amounts to show

$$\forall\ \nu \leq \sigma, \qquad \delta_\nu\left(C_{[i,i+1]}^{\text{old}}\right) \leq \delta_\nu\left(C_{[i,i+1]}^{\text{new}}\right) \leq \frac{\|\mathbf{c}_i^{*\text{new}}\|}{\sigma\delta_\sigma\left(C_{[i,i+1]}^{\text{old}}\right)} \times \delta_\nu\left(C_{[i,i+1]}^{\text{old}}\right) \qquad (11)$$

Two cases can occur in the for loop of Alg. 2:

- First case: $\left\|\mathbf{c}_i^{*\text{old}}\right\| \leq \sigma$ and $\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| \geq \sigma$ (swap case). Since $\|\mathbf{c}_i^{*\text{new}}\|$ is projected on a space of higher dimension than $\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|$, we have $\|\mathbf{c}_i^{*\text{new}}\| \geq \left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| \geq \nu$ and since the projected volume $\operatorname{vol}\left(C_{[i,i+1]}\right)$ remains unchanged, we have $\frac{\|\mathbf{c}_i^{*\text{new}}\|}{\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|} = \frac{\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|}{\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\|}$ and hence $\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\| \leq \left\|\mathbf{c}_i^{*\text{old}}\right\|$.

By definition,

$$\begin{aligned}
\delta_\nu\left(C^{\text{new}}_{[i,i+1]}\right) &= \frac{\left\|\mathbf{c}_i^{*\text{new}}\right\|}{\nu} \times \max\left(1, \frac{\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\|}{\nu}\right) \\
\delta_\nu\left(C^{\text{old}}_{[i,i+1]}\right) &= \frac{\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|}{\nu} \times \max\left(1, \frac{\left\|\mathbf{c}_i^{*\text{old}}\right\|}{\nu}\right)
\end{aligned}$$

We obtain $\frac{\delta_\nu\left(C^{\text{new}}_{[i,i+1]}\right)}{\delta_\nu\left(C^{\text{old}}_{[i,i+1]}\right)} = \frac{\|\mathbf{c}_i^{*\text{new}}\|}{\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|} \times \underbrace{\frac{\max\left(1, \frac{\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\|}{\nu}\right)}{\max\left(1, \frac{\left\|\mathbf{c}_i^{*\text{old}}\right\|}{\nu}\right)}}_{\geq 1}$. Now $\sigma\delta_\sigma\left(C^{\text{old}}_{[i,i+1]}\right) = \left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|$ proves the right side of (11).

For the left side, it suffices to notice that $\frac{\delta_\nu\left(C^{\text{new}}_{[i,i+1]}\right)}{\delta_\nu\left(C^{\text{old}}_{[i,i+1]}\right)} = \frac{\max\left(\|\mathbf{c}_i^{*\text{new}}\|, \frac{\text{vol}\left(C_{[i,i+1]}\right)}{\nu}\right)}{\max\left(\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\|, \frac{\text{vol}\left(C_{[i,i+1]}\right)}{\nu}\right)} \leq 1$.

- Second case: $\left\|\mathbf{c}_i^{*\text{old}}\right\| \geq \sigma$ and $\left\|\mathbf{c}_{i+1}^{*\text{old}}\right\| \geq \sigma$. Thus we have $\|\mathbf{c}_i^{*\text{new}}\| \geq \sigma$ and, $\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\| \leq \sigma$. Therefore, the two equality hold: $\delta_\sigma\left(C^{\text{old}}_{[i,i+1]}\right) = \frac{\text{vol}\left(C_{[i,i+1]}\right)}{\sigma^2}$, and a fortiori, $\delta_\nu\left(C^{\text{old}}_{[i,i+1]}\right) = \frac{\text{vol}\left(C_{[i,i+1]}\right)}{\nu^2}$. Since by definition, $\delta_\nu\left(C^{\text{new}}_{[i,i+1]}\right) = \frac{\|\mathbf{c}_i^{*\text{new}}\|}{\nu} \times \max\left(1, \frac{\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\|}{\nu}\right)$, the left side of (11) easily follows. Furthermore,

$$\begin{aligned}
\frac{\delta_\nu\left(C^{\text{new}}_{[i,i+1]}\right)}{\delta_\nu\left(C^{\text{old}}_{[i,i+1]}\right)} &= \frac{\nu\left\|\mathbf{c}_i^{*\text{new}}\right\|}{\sigma^2\delta_\sigma\left(C^{\text{old}}_{[i,i+1]}\right)} \times \max\left(1, \frac{\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\|}{\nu}\right) \\
&= \frac{\nu\left\|\mathbf{c}_i^{*\text{new}}\right\|}{\sigma\delta_\sigma\left(C^{\text{old}}_{[i,i+1]}\right)} \times \underbrace{\max\left(\frac{\nu}{\sigma}, \frac{\left\|\mathbf{c}_{i+1}^{*\text{new}}\right\|}{\sigma}\right)}_{\leq 1}
\end{aligned}$$

This proves the right side of Inequality (11).

C.3 Proof of Th. 4.4

We will first prove the theorem using the first condition, which is tighter than the second one. The invariant of the main for loop is that at the beginning of i_{th} iteration, the current basis $C_{[i,n]}$ satisfies:

$$\left\|\mathbf{c}_j^*\right\| \leq q_j\sigma \text{ for all } j < i \qquad \text{and} \qquad \delta_\sigma(C_{[i,n]}) \leq \prod_{j=i}^{k} \frac{q_j}{n+1-j} \tag{12}$$

With this invariant, it is clear that the returned $\bar{B}$ at line 3 or 8 satisfies the upper-bound $\left\|\bar{B}^*\right\| \leq \sigma$.

Let us show (12) by induction on i. Clearly, the condition holds for $i = 1$.

At step 4, ℓ exists and is easy to compute, since the function $\nu \to \log(\delta_\nu(C_{[i,n]}))$ is a piecewise affine positive decreasing continuous function which is zero when $\nu = \left\|C^*_{[i,n]}\right\| \leq \|B^*\|$. With this ℓ, unbalanced reduction always produces a new basis such that $\|\mathbf{c}_i^{*\text{new}}\| \leq q_i \cdot \sigma$. Then, there are two cases: either $\ell = \sigma$, and in this case, $\delta_\sigma(C^{\text{new}}_{[i+1,n]}) = 1$. Or we have the equality $\ell\delta_\ell(C^{\text{old}}_{[i,n]}) = q_i\sigma/(n+1-i)$. By replacing σ and ν by ℓ and σ respectively in (1), we obtain:

$$\delta_\sigma\left(C^{\text{new}}_{[i,n]}\right) \leq \frac{\|\mathbf{c}_i^{*\text{new}}\| \times \delta_\sigma\left(C^{\text{old}}_{[i,n]}\right)}{\ell \times \delta_\ell\left(C^{\text{old}}_{[i,n]}\right)} \leq \|\mathbf{c}_i^{*\text{new}}\| \times \frac{\prod_{j=i}^{k} \frac{q_j}{n+1-j}}{\frac{q_i\sigma}{n-i+1}}$$

Since $\delta_\sigma\left(C^{\text{new}}_{[i,n]}\right) = \delta_\sigma\left(C^{\text{new}}_{[i+1,n]}\right) * \frac{\|b_i^{*\text{new}}\|}{\sigma}$, we have:

$$\delta_\sigma\left(C^{\text{new}}_{[i+1,n]}\right) \leq \prod_{j=i+1}^{k} \frac{q_j}{n+1-j}$$

To prove the theorem with the second condition, it suffices to notice that at the first iteration, the length ℓ is smaller than $(n\,\|B^*\|^n/q_1)^{1/(n-1)}$, because for this value, $\ell\delta_\ell(B) \leq \ell(\|B^*\|/\ell)^n = \frac{q_1}{n}$. Therefore, for the next iteration, $\left\|B^*_{[2,n]}\right\| \leq \ell$ and $(\ell/\sigma)^{n-1} \leq \frac{n\sigma}{q_1}\|B^*\|^n/\sigma^n \leq q_2 \cdot \cdots \cdot q_n/n^{k-1}$. The proof goes on by induction.

D Missing Proofs of Sect. 5

D.1 Proof of Th. 5.1

Calls to GSIS. Since $\sigma\eta_\varepsilon(\mathbb{Z}^n) \geq \text{bl}(L)\eta_\varepsilon(\mathbb{Z}^n) \geq \eta_\varepsilon(L)$, $\mathcal{D}_{\bar{L}/L,\sigma}$ is statistically close to the uniform distribution over $\bar{L}/L$. Therefore the $\mathbf{v}_i$'s are uniform $\mod L$ by Lemma 2.3. Thus, the elements $g_i = \varphi(\mathbf{v}_i)$ (defined at Line 5) have uniform distribution over G, which allows to make calls to the GSIS oracle at Line 6.

Correctness. It is easy to see that $\mathbf{v} = \sum_{i=1}^m x_i\mathbf{v}_i$ (defined in Line 7) is indeed a short vector of L, since $\varphi(\mathbf{v}) = \sum_{i=1}^m x_i g_i = 0$ and

$$\begin{aligned} E[\|\mathbf{v}\|] &\leq \|\mathbf{x}\| \times E[\|\mathbf{v}_i\|] \\ &\leq \beta \times \sqrt{n/2\pi}\eta_\varepsilon(\mathbb{Z}^n)\sigma \end{aligned}$$

Termination. It remains to prove that the algorithm indeed outputs n linearly independent vectors (and in particular that the output vectors are non-zero). This part of the proof is similar to the proof of [21]. The distribution of the output vectors $\mathbf{v}$'s depends on the $\mathbf{v}_i$'s and on the answer $\mathbf{x}$ of the GSIS-oracle, which only depends on $\mathbf{g} = (g_1, \ldots, g_m)$. The distribution of $(\mathbf{v}_i), (g_i), \mathbf{x}$ during the algorithm can be equivalently simulated as follows: First choose $(g_1, \ldots, g_m)$ uniformly in G, and call the GSIS oracle which returns a non-zero solution $\mathbf{x}$ with non-negligible probability. Now, for each g_i, sample the preimages $\mathbf{v}_i$, which necessarily have the conditional distribution of $\mathbf{v}_i \leftarrow D_{\bar{L},\sigma\eta_\varepsilon(\mathbb{Z}^n)}$ where $\varphi(\mathbf{v}_i) = g_i$, *i.e.* the distribution $D_{\varphi^{-1}(g_i),\sigma\eta_\varepsilon(\mathbb{Z}^n)}$ where $\varphi^{-1}(g_i)$ is a coset of L. From Proposition 2.2, since $\sigma\eta_\varepsilon(\mathbb{Z}^n) \geq \sqrt{2}\eta_\varepsilon(L)$, one can form a full rank family from $O(n)$ of such samples, which proves that the algorithm terminates.

D.2 Proof of Cor. 5.2

We consider two cases, depending on the rank k_n of G_n.

If $k_n \leq n$ and $\#G_n \geq n^{k_n}\left(\eta_\varepsilon(\mathbb{Z}^n)\sqrt{2n/\pi}\beta_n\right)^n$, then it is a direct consequence of Th. 5.1.

Now, assume that $k_n > n$ and $\#G_n \geq n^{k_n}\left(\eta_\varepsilon(\mathbb{Z}^n)\sqrt{2n/\pi}\beta_n\right)^{k_n}$. Consider the decomposition of G_n into elementary divisors: $G_n \simeq \prod_{i=1}^{k_n}\mathbb{Z}_{q_i}$ where each q_{i+1} divides q_i. Then:

$$\left(\prod_{i=1}^{n} q_i\right)^{1/n} \geq \left(\prod_{i=1}^{k_n} q_i\right)^{1/k_n}.$$

Letting $H_n = \prod_{i=1}^n \mathbb{Z}_{q_i}$, we get that $\#H_n \geq \#G_n^{n/k_n} \geq n^n\left(\eta_\varepsilon(\mathbb{Z}^n)\sqrt{2n/\pi}\beta_n\right)^n$ and H_n has rank n. Therefore solving GSIS(H_n, m_n, β_n) with probability $\geq 1/\text{poly}(n)$ can be used to solve worst-case n-dim ApproxSIVP$_{\eta_\varepsilon(\mathbb{Z}^n)\sqrt{n/\pi}\beta_n}$ But since $G_n \simeq H_n \times J_n$ for some finite abelian group J_n, we know that solving GSIS(H_n, m_n, β_n) with probability $\geq 1/\text{poly}(n)$ can be reduced to solving GSIS(G_n, m_n, β_n) with probability $\geq 1/\text{poly}(n)$.

E Missing proofs of Section 6

E.1 Proof of Theorem 6.1

Let $\mathbf{t}, B$ be the BDD-β instance on the dual $L(B)^{\times}$, and call $d_0 \leq \beta\lambda_1(L^{\times})$ an upper-bound on the error norm. Like in Theorem 6.1, we suppose that $\beta \leq \sqrt{\pi/2}\,(2n\eta_{\varepsilon}(\mathbb{Z}^n))^{-1}$, and call $\theta = \beta\sqrt{2/\pi}\,(2n\eta_{\varepsilon}(\mathbb{Z}^n)) < 1$.

In Alg. 6, the parameter $\alpha \in [\theta\sqrt{\pi/2}, \sqrt{\pi/2})$ is a valid noise parameter for GLWE oracles.

The parameter $\sigma_0 = \alpha/(\sqrt{2}d_0\eta_{\varepsilon}(\mathbb{Z}^n))$ is larger than $2n\beta/\sqrt{2}d_0$. Note that by Banaszczyk theorem [5], $\mathrm{bl}(L) \cdot \lambda_1(L^{\times}) \leq n$, so $\sigma_0 \geq \sqrt{2}\,\mathrm{bl}(L)$. Since $\#G$ is larger than $n^k(\|B^*\|/\sigma_0)^n$, one can indeed apply structural reduction to obtain a basis $\bar{B}$ of $\bar{L}$ such that $\|\bar{B}^*\| \leq \sigma_0$ (line 3 of Alg. 6).

There exists a (unique) vector $\mathbf{u} \in L^{\times}$ such that $\mathbf{t} = \mathbf{u} + \mathbf{w}$ with $\|\mathbf{w}\| \leq d_0$. We now prove that the instance $(a_i, b_i)_{i\in[1,m]}$ generated lines 6,7 is indistinguishable from a random GLWE$(G, m, \leq \alpha)$ instance of solution $\hat{s} = \varphi^{\times}(\mathbf{u}) \in \hat{G}$. Namely the a_i's must be uniform in G, and for each $i \in [1, m]$, b_i must have distribution $\mathcal{D}_{\mathbb{T},\alpha,\hat{s}(a_i)}$.

The uniformity of the a_i's in G comes from the same reason as in Section 5, since they are isomorphic (by φ) to the $\mathbf{v}_i \mod L$, and the $\mathbf{v}_i$'s are drawn from a Gaussian distribution of parameter $\sigma_0\eta_{\varepsilon}(\mathbb{Z}^n) \geq \eta_{\varepsilon}(L)$. To show that the b_i's have the correct distribution, the idea is that $\hat{s}(a_i) = [\varphi^{\times}(\mathbf{u})](\varphi'(\mathbf{v}_i)) = \langle \mathbf{v}_i, \mathbf{u}\rangle \mod 1$. Suppose that a_i is fixed. Then the conditional distribution of $\mathbf{v}_i$ is $\mathcal{D}_{\varphi^{-1}(a_i),\sigma_0\eta_{\varepsilon}(\mathbb{Z}^n)}$ where $\varphi^{-1}(a_i)$ is a coset of L. Since the distribution of b is $\mathcal{D}_{\mathbb{T},\alpha/\sqrt{2},\langle \mathbf{t},\mathbf{v}_i\rangle}$ and $\langle \mathbf{t}, \mathbf{v}_i\rangle = \hat{s}(a_i) + \langle \mathbf{w}, \mathbf{v}_i\rangle$ where $\mathbf{v}_i$ has a discrete Gaussian distribution over a coset of L, then by the convolution Lemma 2.4, the distribution of b_i is at distance 4ε from the distribution $\mathcal{D}_{\mathbb{T},\nu,\hat{s}(a_i)}$ where the parameter ν is $= \sqrt{\alpha^2/2 + (\|\mathbf{w}\|\,\sigma_0\eta_{\varepsilon}(\mathbb{Z}^n))^2} \leq \alpha$.

Subsequent Iterations. Since the Search-GLWE oracle cannot distinguish the distribution of $(a_i, b_i)_{i\in[1,m]}$ from random GLWE samples, it will output the solution $\hat{s} = \varphi^{\times}(\mathbf{u})$ after a polynomial number of trials. Unfortunately, $\varphi^{\times}$ is not invertible, we can only recover $\mathbf{u}$ modulo $\ker(\varphi^{\times}) = \bar{L}^{\times}$. Let $\mathbf{u}_0$ be one preimage in $\varphi^{\times-1}(\hat{s})$. The vector $\mathbf{t} - \mathbf{u}_0$ is now at distance $\leq d_0$ of $\bar{L}^{\times}$ instead of $L^{\times}$. Thus we can iterate the whole process by replacing L with $\bar{L}$.

Since $\mathrm{bl}(L)$ has decreased, the authorized interval for α increases, so α remains a valid noise parameter, and the same oracle may be used for all subsequent iterations.

Since the structural reduction always computes bases such that $\|B^*\|$ decreases by a constant factor compared to the previous basis, the while loop can be iterated $O(\log n)$ times, until $\|B^*\|$ becomes smaller than $1/d_0$. At this point, the BDD is very easy to solve exactly, for example using Babai nearest-plane algorithm.

E.2 Proof of Cor. 6.2

Let $n \in \mathbb{N}$, and let $Q_n = (1 + \varepsilon_{_{\mathrm{LLL}}})\sqrt{4/3})^{\frac{n-1}{2}}\,\frac{\theta_n\sqrt{\pi/2}}{\sqrt{2}\alpha_n}$. Like in Cor. 5.2, the case $k_n \leq n$ and $\#G_n \geq n^{k_n}\,(Q_n)^n$, is a direct consequence of (multiple iterations of) Th. 6.1 and Regev's quantum connection between BDD_{β_n} and $\mathrm{ApproxSIVP}_{\sqrt{2}n/\beta_n}$.

Now, assume that $k_n > n$ and $\#G_n \geq n^{k_n}\,(Q_n)^{k_n}$. Again, from the decomposition of G_n into elementary divisors: $G_n \simeq \prod_{i=1}^{k_n} \mathbb{Z}_{q_i}$ where each q_{i+1} divides q_i, we can decompose G_n as $H_n \oplus J_n$ where the subgroup $H_n = \prod_{i=1}^{n} \mathbb{Z}_{q_i}$ has rank n and satisfies $\#H_n \geq n^n\,(Q_n)^n$. Note that any GLWE sample (a, b) on H_n with (unknown) secret $\hat{s}$ can be combined with a randomly generated GLWE sample (a', b') over J_n with a randomly chosen secret $\hat{s} \in \hat{J_n}$ to form a GLWE sample on G. Therefore solving Search-GLWE(G, α_n) with probability $\geq 1/\mathrm{poly}(n)$ can be used to solve Search-GLWE(H_n, α_n) with probability $\geq 1/\mathrm{poly}(n)$ which in turns can be used to solve worst-case n-dim BDD_{β_n}.

F Missing Proof of Sect. 7

F.1 Proof of Theorem 7.1

For simplicity, we denote $G = G_n$, $k = k_n$, $\alpha' = \alpha'_n$ and $\alpha = \alpha_n$. First, observe that one can transform a sample (a, b) from $A_{G,\alpha}(\hat{s})$ to a sample (a', b') from $A_{G,\alpha'}(\hat{s})$ by simply setting $a' = a$ and $b' =$

$\mathcal{D}_{\mathbb{T},\sqrt{\alpha'^2-\alpha^2},b}$.

Let $\hat{s} = \sum_{i=1}^{k} s_i \hat{e}_i$ for $s_i \in \mathbb{Z}_{q_i}$. We now show how to recover each coordinate s_1 modulo powers of any prime p_{1,j_1} for $j_1 \in [1, t_1]$ (similarly for $s_2, ..., s_k$). Let $p = p_{1,j}$ and $\mu = \mu_{1,j_1}$, and for $\tau \in [0, \mu]$ define $A^{\tau}_{G,\alpha'}(\hat{s})$ to be the distribution by sampling (a, b) from $A_{G,\alpha'}(\hat{s})$ and outputting $(a, b + r/p^{\tau} \in \mathbb{T})$, where $r \leftarrow \mathbb{Z}_{q_1}$ is uniformly random. Notice that when $\alpha' \geq \omega(\sqrt{\log n})/p^{\tau} \geq \eta_{\varepsilon}({}^{1}\!/_{q^{\tau}}\mathbb{Z})$ for some negligible ε, $A^{\tau}_{G,\alpha'}(\hat{s})$ is statistically close to uniformly random in $G \times \mathbb{T}$ and this holds at least for $\tau = \mu$ by hypothesis. Therefore, given an oracle solving Decisional-GLWE$_{G,\alpha'}$, there exists some minimal $\tau \in [1, \mu]$ such that the oracle has a non-negligible advantage in distinguishing $A^{\tau-1}_{G,\alpha'}(\hat{s})$ and $A^{\tau}_{G,\alpha'}(\hat{s})$. Note that when $p \geq \omega(\sqrt{\log n})/\alpha \geq \omega(\sqrt{\log n})/\alpha'$ the minimal τ must be 1. We can find that τ efficiently by estimating the behaviour of the oracle. By standard self-reduction and amplification techniques, we can assume that the oracle accepts (respectively, rejects) the distribution $A^{\tau-1}_{G,\alpha'}(\hat{s})$ (resp. $A^{\tau}_{G,\alpha'}(\hat{s})$) with overwhelming probability.

Given access to $A^{\tau-1}_{G,\alpha'}(\hat{s})$ and the oracle, we can test whether $s_1 = 0 \bmod p$ by invoking the oracle on the samples defined as follows: transform each sample $(a, b = \mathcal{D}_{\mathbb{T},\alpha',\hat{s}(a)} + r/p^{\tau-1}) \leftarrow A^{\tau-1}_{G,\alpha'}(\hat{s})$ into

$$a' = a - (r' q_1/p^{\tau}) \cdot e_1, \quad b' = b = \mathcal{D}_{\mathbb{T},\alpha',\hat{s}(a')} + (pr + r' s_1)/p^{\tau} \in \mathbb{T},$$

where $r' \leftarrow \mathbb{Z}_{q_1}$ is uniformly random. Observe that if $s_1 = 0 \bmod p$, the transformed samples are distributed as $A^{\tau-1}_{G,\alpha'}(\hat{s})$, otherwise they are distributed as $A^{\tau}_{G,\alpha'}(\hat{s})$ because r is uniformly random in $\mathbb{Z}_{q_1}$ and $r' s_1$ is uniformly random in $\mathbb{Z}_p$. Hence, the oracle will tell us which is the case.

Using the above test, we can recover $s_1 \bmod p$ by shifting s_1 by each of $0, 1, ..., p-1 \bmod p$ using standard transform that maps from $A_{G,\alpha'}(\hat{s})$ to $A_{G,\alpha'}(\hat{s} + \hat{t})$ for known $\hat{t} \in \hat{G}$. This procedure is efficient since each prime factor is polynomially bounded. Furthermore, we can iteratively recover s_1 mod $p^2, ..., p^{\mu-\tau+1}$ as follows: after recovering $s_1 \bmod p^i$, first shift $A_{G,\alpha}(\hat{s})$ to $A_{G,\alpha}(\hat{s}')$ such that $s'_1 = 0 \bmod p^i$, then apply a similar procedure to recover $s'_1 \bmod p^{i+1}$ by letting $a' = a - (r' q_1/p^{\tau+i}) \cdot e_1$ and $b' = b = \mathcal{D}_{\mathbb{T},\alpha',\hat{s}'(a')} + (pr + r'(s'_1/p^i))/p^{\tau}$. This procedure works as long as $p^{\tau+i}$ divides q_1, so we can recover s_1 mod $p^{\mu-\tau+1}$.

Let τ_{i,j_i} for $i \in [1,k], j_i \in [1, t_i]$ be the minimal value of τ for $p = p_{i,j_i}$ (of which at most ℓ of these value > 1). Therefore, using the above reduction and Chinese remainder theorem we can recover s_i mod T_i for $i \in [1, k]$, where

$$T_i = \prod_{j_i=1}^{t_i} p_{i,j_i}^{u_{i,j_i} - \tau_{i,j_i} + 1} = q_i / \prod_{j_i=1}^{t_i} p_{i,j_i}^{\tau_{i,j_i} - 1}.$$

The product of all the T_i's satisfies

$$\prod_{i=1}^{k} T_i = \prod_{i=1}^{k} (q_i / \prod_{j_i=1}^{t_i} p_{i,j_i}^{\tau_{i,j_i}-1}) = \#G / \prod_{i=1}^{k} \prod_{j_i=1}^{t_i} p_{i,j_i}^{\tau_{i,j_i}-1} \geq \#G \cdot (\frac{\alpha'}{\omega(\sqrt{\log n})})^{\ell} \geq \#G \cdot \alpha \cdot \omega(\sqrt{\log n})$$

because $\alpha' < \omega(\sqrt{\log n})/p_{i,j_i}^{\tau_{i,j_i}-1}$ for all i, j_i by the definition of τ_{i,j_i} and the hypothesis of α'. By applying the shift transform on $A_{G,\alpha}(\hat{s})$, we can assume that $s_i = 0 \bmod T_i$ for $i \in [1, k]$. For any sample (a, b) from $A_{G,\alpha}(\hat{s})$, we have that $b = \sum_{i=1}^{k} (a_i s_i)/q_i + \delta \bmod 1$, where $a_i \in \mathbb{Z}_{q_i}$ is the coordinate of a and $\delta \leftarrow \mathcal{D}_{\mathbb{R},\alpha}$. Note that each $(a_i s_i)/q_i$ is an integer multiple of T_i/q_i. Define

$$g \triangleq \mathrm{lcm}(\frac{q_1}{T_1}, ..., \frac{q_k}{T_k}) \leq \prod_{i=1}^{k} \frac{q_i}{T_i} = \frac{\#G}{\prod_{i=1}^{k} T_i} \leq \frac{1}{\alpha \cdot \omega(\sqrt{\log n})},$$

and compute $g \cdot b = g \cdot (\sum_{i=1}^{k} (a_i s_i)/q_i + \delta + z) = \sum_{i=1}^{k} g \cdot (a_i s_i)/q_i + g\delta + gz$, where $z \in \mathbb{Z}$. Now observe that $\sum_{i=1}^{k} g \cdot (a_i s_i)/q_i$ and gz are integers. Since $\delta \leftarrow \mathcal{D}_{\mathbb{R},\alpha}$ has magnitude $< (\alpha/2) \cdot \omega(\sqrt{\log n})$ except with negligible probability, we have $\lfloor gb \rceil = \sum_{i=1}^{k} g \cdot (a_i s_i)/q_i + gz$. Therefore, we can recover $\hat{s}$ by Gaussian elimination.

F.2 Overview

The following figure represents the main components of group switching and their relationship.

$$\begin{tikzcd}
& & \mathbb{Z}^n \arrow[dl, dashed] \arrow[dr, dashed] & & \\
G & \bar{L} \arrow[l, "\varphi"'] \arrow[rr, bend right, "\text{SampleGauss}"'] & & \bar{L}' \arrow[r, "\varphi'"] & G' \\
a \arrow[r, "\varphi^{-1}"] & \mathbf{u} \arrow[rr] & & \mathbf{v} \arrow[r, "\varphi'"] & a' \\
& (\text{mod } \mathbb{Z}^n) & & (\text{mod } \mathbb{Z}^n) &
\end{tikzcd}$$

$$\begin{tikzcd}
& \bar{L}^{\times} \arrow[d, dashed] & \bar{L}'^{\times} \arrow[dl, dashed] & \\
\hat{G} & \mathbb{Z}^n \arrow[l, "\varphi^{\times}"'] \arrow[rr, "\varphi'^{\times}"] & & \hat{G}' \\
\hat{s} \arrow[r, "\varphi^{\times -1}"] & \mathbf{s} \arrow[rr, "\varphi'^{\times}"] & & \hat{s}' \\
\mathcal{S} \arrow[r] & S \arrow[rr] & & \mathcal{S}' \\
& (K\text{-Bounded}) & &
\end{tikzcd}$$

Dotted arrows represent overlattices. The left side are the primal extensions, the right side contains dual overlattices. On the left of each diagram, we recognize the dual pair of surjective morphisms φ from $\bar{L}$ (mod $\mathbb{Z}^n$) to G, and $\varphi^{\times}$ from $\mathbb{Z}^n$ (mod $\bar{L}^{\times}$) to $\hat{G}$. On the right side of each diagram, the dual morphisms $\varphi', \varphi'^{\times}$ related to the second group G'.

To transform a $A_{G,\alpha}(\hat{s})$ sample (a, b) to a $A_{G',\beta}(\hat{s}')$ sample (a', b'), we compute $\mathbf{u} \in \varphi^{-1}(a)$ a preimage in $\bar{L}$ (mod $\mathbb{Z})^n$, we sample a point $\mathbf{v}$ in $\bar{L}'$ (mod $\mathbb{Z})^n$ close to $\mathbf{u}$, and we apply φ' to get a'.

On the dual side, the secret $\hat{s}$ admits in general a unique preimage $\mathbf{s}$ by $\varphi^{\times}$ which has small norm. Finding this small $\mathbf{s}$ may in some cases be very difficult, as it requires to solve a CVP in $\bar{L}^{\times}$. However, only the existence of the small $\mathbf{s}$ matters. Then, the new (unknown) secret is $\hat{s}' = \varphi'^{\times}(\mathbf{s})$. In general, we know the K-bounded distribution S over $\mathbb{Z}^n$ from which $\mathcal{S} = \varphi^{\times}(S)$ is defined. Thus, the distribution of secrets $\mathcal{S}' = \varphi'^{\times}(S)$ in the new GLWE is entirely determined.

Finally, note that by duality of the morphisms, $\hat{s}(a) = \langle \mathbf{s}, \mathbf{u} \rangle \approx \langle \mathbf{s}, \mathbf{v} \rangle = \hat{s}'(a')$. We just need to add some noise $b' = b$+noise to compensate the Gaussian approximation between $\langle \mathbf{s}, \mathbf{u} \rangle$ and $\langle \mathbf{s}, \mathbf{v} \rangle$ in the middle.

F.3 Proof of Lemma 7.4

The main idea consists in the following: given an element of G, sample randomly an element of G' so that the evaluations on these two elements of the corresponding characters is almost preserved. The approximate equivalence of evaluations of characters comes from the duality of the maps φ' and $\varphi'^{\times}$. Given a sample $(a, b) \in G \times \mathbb{T}$, the procedure is as follows:

```
1: Choose one preimage $\mathbf{u} \in \varphi^{-1}(a)$ and sample $\mathbf{v} \leftarrow \mathcal{D}_{\bar{L}',r,\mathbf{u}}$ using the basis $\bar{B}'$ and Lemma 2.1. The
   preimage can be computed because $G$ is fully-explicit.
2: Let $a' = \varphi'(\mathbf{v})$.
3: Choose $b' \leftarrow \mathcal{D}_{\mathbb{T},rK,b}$.
4: Output $(a', b')$.
```

We now analyze the algorithm. We first show that the distribution of a' is nearly uniform in G'. It suffices to show that $\mathbf{v} \mod \mathbb{Z}^n \in \bar{L}'/\mathbb{Z}^n$ is (nearly) uniformly random. We note that, if $r \geq \eta_{\varepsilon}(\mathbb{Z}^n)$, we have that $\mathbf{v} \mod \mathbb{Z}^n$ is (almost) uniform (see [14]). However, this would require a very large r, which is not suitable for our reduction. Since a is uniform in G, a much smaller r is sufficient to show the uniformity of $\mathbf{v}$ mod $\mathbb{Z}^n$. Indeed, let $\mathcal{A} \in \bar{L}$ be a (finite) set containing exactly one representative of each class of $\bar{L}/\mathbb{Z}^n$. Note that $\varphi^{-1}(a)$ is a uniformly random coset $\mathbf{u} + \mathbb{Z}^n$ where $\mathbf{u} \in \mathcal{A}$. Let $\mathbf{v}_0 \in \bar{L}'$.

$$\begin{aligned}
\Pr[\mathbf{v} = \mathbf{v}_0 \bmod \mathbb{Z}^n] &= \frac{1}{\#G} \sum_{a \in G} \mathcal{D}_{\bar{L}'/\mathbb{Z}^n, r, \varphi^{-1}(a)}(\mathbf{v}_0 + \mathbb{Z}^n) = \frac{1}{\#G} \sum_{\mathbf{u} \in \mathcal{A}} \mathcal{D}_{\bar{L}'/\mathbb{Z}^n, r, \mathbf{u}+\mathbb{Z}^n}(\mathbf{v}_0 + \mathbb{Z}^n) \\
&= \frac{1}{\#G} \sum_{\mathbf{u} \in \mathcal{A}} \sum_{\mathbf{z} \in \mathbb{Z}^n} \mathcal{D}_{\bar{L}', r, \mathbf{u}}(\mathbf{v}_0 + \mathbf{z}) = \frac{1}{\#G} \sum_{(\mathbf{u}-\mathbf{z}) \in \bar{L}} \mathcal{D}_{\bar{L}', r}(\mathbf{v}_0 + (\mathbf{z} - \mathbf{u})) \\
&= \frac{1}{\#G} \sum_{\mathbf{w} \in \bar{L}} \frac{\rho_{\mathbb{R}^n, r}(\mathbf{v}_0 + \mathbf{w})}{\rho_{\mathbb{R}^n, r}(\mathbf{v}_0 + \mathbf{w} + \bar{L}')} =_{2\varepsilon} \frac{1}{\#G} \cdot \frac{1/\operatorname{vol}(\bar{L})}{1/\operatorname{vol}(\bar{L}')} = \frac{1}{\#G} \cdot \frac{\#G}{\#G'} = \frac{1}{\#G'}.
\end{aligned}$$

Clearly, if the input b is uniformly random in $\mathbb{T}$, then b' is also uniform in $\mathbb{T}$. It remains to show that given as input a sample distributed from $A_{G,\alpha}(\hat{s})$, the algorithm outputs a sample distributed from $A_{G',\beta}(\hat{s}')$. Let $\mathbf{f} = \mathbf{v} - \varphi^{-1}(a)$, the distribution of $\mathbf{f}$ is $\mathcal{D}_{\bar{L}' - \varphi^{-1}(a), r}$. We also have, that $\hat{s}'(a') = [\varphi'^{\times}(\mathbf{s})](\varphi'(\mathbf{v})) = \langle \mathbf{s}, \mathbf{v} \rangle = \langle \mathbf{s}, \varphi^{-1}(a) \rangle + \langle \mathbf{s}, \mathbf{f} \rangle = \hat{s}(a) + \langle \mathbf{s}, \mathbf{f} \rangle$. Assume a' is fixed. Since b' is sampled from $\mathcal{D}_{\mathbb{T}, rK, b}$ where $b \leftarrow \mathcal{D}_{\mathbb{T}, \alpha, \hat{s}(a)}$, by classical convolution, the distribution of b' is $\mathcal{D}_{\mathbb{T}, \sqrt{\alpha^2 + (rK)^2}, \hat{s}(a)}$ where $\hat{s}(a) = \hat{s}'(a') - \langle \mathbf{s}, \mathbf{f} \rangle$. Since $\mathbf{f}$ has Gaussian distribution over a coset, by the dot-product convolution lemma 2.4, the distribution of b' is statistically close to $\mathcal{D}_{\mathbb{T}, \sqrt{\alpha^2 + (\|\mathbf{s}\| r)^2 + (rK)^2}, \hat{s}'(a')}$

G Applications: Abstracting Lattice Cryptography

We showed that GSIS and GLWE are hard under worst-case assumptions, provided that the finite Abelian group G is sufficiently large. This suggests to abstract all lattice schemes based on SIS and/or decisional/search-LWE using an arbitrary finite Abelian group G, and check that the security proof carries through, under the assumption that GSIS or GLWE is hard, which holds under the same worst-case lattice assumptions than for SIS and LWE. We believe that such an abstraction leads to a better understanding of the scheme and a clearer presentation: lattice schemes are typically presented using matrices and vectors, which our abstraction avoids. Furthermore, it could be that special choices of G (other than the homocyclic group $\mathbb{Z}_q^n$) may have other benefits.

G.1 GLWE Encryption

As a simple example of abstraction, we generalize the so-called dual version of Regev's [31] LWE-based encryption, which was used to build ID-based encryption from lattices [14].

$\mathsf{Gen}(1^n)$: Takes as input a security parameter n, it chooses a Gaussian parameter $0 < \alpha < 1$, a (sufficiently large) finite Abelian group G and $m \in \mathbb{N}$ group elements $\mathbf{g} = (g_1, ..., g_m) \in G^m$ chosen uniformly at random.

The secret key is a uniformly random vector $\mathbf{x} \in \{0,1\}^m$, whose corresponding public key is $y = \sum_{i=1}^m x_i g_i \in G$.

$\mathsf{Enc}(pk, b)$: Given as input a public key y and a message $b \in \{0,1\}$, it selects $\hat{s} \in \hat{G}$ uniformly at random and $m+1$ Gaussian errors $(e, e_1, ..., e_m) \leftarrow \mathcal{D}_\alpha^{m+1}$. Then the ciphertext is $\mathbf{c} = (c, \{c_i\}_{1 \le i \le m})$, where $c = \hat{s}(y) + e + b/2$ and $c_i = \hat{s}(g_i) + e_i$ for $1 \le i \le m$.

$\mathsf{Dec}(sk, \mathbf{c})$: Given as input a secret key $\mathbf{x}$ and a ciphertext $\mathbf{c}$, it first parses $\mathbf{c} = (c, \{c_i\}_{1 \le i \le m})$ and then computes $b' = c - \sum_{i=1}^m x_i c_i \in \mathbb{T}$. If b' is closer to 0 than to 1/2, it outputs 0; otherwise, it outputs 1.

Lemma G.1 (Correctness) *If $0 < \alpha < 1/(4 \cdot \sqrt{m+1} \cdot \omega(\sqrt{\log n}))$, then the above public key encryption scheme will decrypt correctly with probability $1 - negl(n)$.*

Proof. Hence we have:

$$c - \sum_{i=1}^m x_i c_i = \hat{s}(y) + e + b/2 - \Big(\sum_{i=1}^m x_i (\hat{s}(g_i) + e_i) \Big) = e - \sum_{i=1}^m x_i e_i + b/2.$$

It is sufficient to show $|e - \sum_{i=1}^{m} x_i e_i| < 1/4$. Let $k \leq m$ be the Hamming weight of $\mathbf{x}$, we know that $e - \sum_{i=1}^{m} x_i e_i$ is distributed as $\mathcal{D}_{\sqrt{k+1}\alpha}$. Therefore, it implies that $|e - \sum_{i=1}^{m} x_i e_i| < \sqrt{k+1}\alpha \cdot \omega(\sqrt{\log n})$ with probability $1 - \exp(-\pi \cdot \omega(\log n)) = 1 - \mathrm{negl}(n)$. We obtain that $|e - \sum_{i=1}^{m} x_i e_i| < \sqrt{k+1}\alpha \cdot \omega(\sqrt{\log n}) \leq 1/4$ with probability $1 - \mathrm{negl}(n)$, as desired. □

Lemma G.2 (Security) *If $m \geq \log \#G + \omega(\log n)$ and the $\mathrm{GLWE}_{G,m+1,\alpha}$ assumption holds, then the above scheme is IND-CPA secure.*

Proof. Since $m \geq \log \#G + \omega(\log n)$, $\mathbf{g} \in G^m$ has distribution statistically close to uniform. Since $\mathbf{x} \in \{0,1\}^m$, the leftover hash lemma ensures that $y = \sum_{i=1}^{m} x_i g_i$ is statistically close to uniform over G. Therefore, one can replace $(c, \{c_i\}_{i=1}^{m})$ with a uniformly random vector over $\mathbb{T}^{m+1}$ under the $\mathrm{GLWE}_{G,m+1,\alpha}$ assumption. This proves the IND-CPA security of the above scheme. □

Similarly, one can generalize the GPV signature scheme [14] based on SIS, and by combining it with the previous GLWE encryption, one thus generalizes the ID-based encryption of [14] using GSIS and GLWE.

G.2 Analogy with El Gamal Encryption

The abstract presentation of GLWE encryption allows to make an easy analogy with El Gamal encryption based on DDH. In El Gamal encryption, one works with a cyclic group G' generated by some g' of order q, which we denote additively:

- The secret key is a $x' \in \mathbb{Z}_q$ chosen uniformly at random, and the public key is $y' = x'g'$. Thus, the public key y' is the image of the secret key x' through the DL one-way function over G'. Similarly, in the previous GLWE encryption, the public key y is the image of the secret key $\mathbf{x}$ through the GSIS one-way function over G.

- The El Gamal ciphertext of a message $b' \in G'$ is a pair $(c', d') \in G \times G$ where $c' = b' + e'y'$, $d' = e'g'$ and $e' \in \mathbb{Z}_q$ is a one-time key chosen uniformly at random. The first element c' is a one-time pad encryption of the message b' with the Diffie-Hellman key $e'y' = x'd'$. The second element d' is the image of the one-time key e' through the DL one-way function over G'. Similarly, in the previous GLWE encryption, the first element c is a one-time pad encryption of the message $b/2$ with the common (noisy) key $\hat{s}(y) + e \approx \hat{s}(y)$. The second element $\{c_i\}_{1 \leq i \leq m}$ is the image of the one-time key $(\hat{s}, e_1, ..., e_m)$ through the GLWE one-way function over G.

In other words, El Gamal encryption is based on the Diffie-Hellman key exchange, which pairs two DL one-way functions. On the other hand, GLWE encryption is based on a noisy key exchange which pairs the GSIS one-way function with the GLWE one-way function.