Hardness of hinted ISIS from the space-time hardness of lattice problems

Martin R. Albrecht[1,2], Russell W. F. Lai[3], and Eamonn W. Postlethwaite[1]

[1] King's College London, London, UK
[2] SandboxAQ, Palo Alto, CA, USA
[3] Aalto University, Espoo, Finland

Version: 7th June 2026

Abstract We initiate the study of basing the hardness of hinted ISIS problems (i.e. with trapdoor information, or 'hints') on the previously conjectured space-time hardness of lattice problems without hints. We present two main results.

1. If there exists an efficient algorithm for hinted ISIS that outputs solutions a constant factor longer than the hints, then there exists a single-exponential time and polynomial memory zero-centred spherical Gaussian sampler solving hinted SIS with norm a constant factor shorter than the hints.
2. Assume the existence of a chain of algorithms for hinted ISIS each taking as input Gaussian hints whose norms decrease by a constant factor at each step in the chain, then there exists a single-exponential time and polynomial memory algorithm for SIS with norm a quasilinear factor from optimal.

The existence of such hinted ISIS solvers implies single-exponential time and polynomial memory algorithms for worst-case lattice problems, contradicting a conjecture by Lombardi and Vaikuntanathan (CRYPTO'20) and all known algorithms. This suggests that hinted ISIS is hard.

Apart from advancing our understanding of hinted lattice problems, an immediate consequence is that signing the same message twice in GPV-style [Gentry–Peikert–Vaikuntanathan, STOC'08] schemes (without salting or derandomisation) likely does not compromise unforgeability. Also, cryptanalytic attempts on the One-More-ISIS problem [Agrawal–Kirshanova–Stehlé-Yadav, CCS'22] likely will need to overcome the conjectured space-time hardness of lattices.

Contents

1 Introduction . . . 3
1.1 Contributions . . . 4
2 Technical overview . . . 8
2.1 Hinted variants of SIS and ISIS . . . 8
2.2 Technical tools . . . 8
2.3 Polynomial memory 'sieving' . . . 9
2.4 Open problems . . . 10
3 Preliminaries . . . 11
3.1 Probability and entropy . . . 11
3.2 Problems and algorithms . . . 12
3.3 Classical tail bounds . . . 13
3.4 Graph theory . . . 14
3.5 (Smooth) Rényi divergence . . . 14
3.6 Lattices . . . 15
3.7 Gaussians . . . 16
3.8 Leftover hash lemma . . . 19
3.9 Upper bounding $\mu(\Lambda_q^{\perp}(m,n))$ . . . 20
3.10 Computational problems . . . 20
3.11 Finding short or close vectors . . . 21
3.12 Finding short vectors in polynomial memory . . . 23
4 Lower bounding the probability of sampling near-neighbours . . . 26
4.1 An upper bound on the number of not close points in a ball . . . 26
4.2 Randomisation via Motzkin–Straus . . . 26
5 Gaussian vectors generate the lattice . . . 28
6 Rejection sampling . . . 32
7 Hinted SIS problems . . . 38
7.1 Definitions of hinted problems . . . 38
7.2 'Sieving with centres' . . . 39
8 Polynomial memory sieve . . . 40
8.1 Sampling a shorter generating set . . . 40
8.2 Sampling a clean shorter generating set . . . 42
8.3 Chaining adversaries . . . 43
9 Bounds on λ_1 for $\Lambda_q^{\perp}(m,n)$ . . . 45
10 Conclusion . . . 48
A Reduction from One-More-ISIS . . . 56
A.1 Multi instance repetition gadget . . . 56
A.2 Reduction . . . 56

1 Introduction

Computing short vectors in Euclidean lattices is a well studied and assumed hard computational problem, both in the worst case and in the average case over random q-ary lattices. On the one hand, the shortest vector problem (SVP) is known to be NP-hard up to sublinear approximation factors under *randomised* reductions [Ajt98, Mic01, Kho05, HR12, Mic12]. On the other hand, the worst-case hardness of these problems, e.g. the shortest independent vector problem (SIVP), with polynomial approximation factors implies the average-case hardness of the short integer solution (SIS) [Ajt96, MR07] and learning with errors (LWE) [Reg05, Pei09, BLP+13a] problems, cornerstones of lattice-based cryptography.

Space-time hardness of lattice problems. A long line of works has studied both the provable, worst-case and heuristic, average-case complexity of finding short vectors in lattices. In both regimes – provable and heuristic – two classes of algorithms are well established; for lattices of rank n, *enumeration* style algorithms [Kan83, FP85, GNR10, MW15, ABF+20, ABLR21] run in time $n^{\Theta(n)}$ and require only $\mathsf{poly}(n)$ memory, whereas *sieve* style algorithms [AKS01, NV08, MV10a, Laa15b, ADS15, BDGL16, Duc18, ADH+19, ALS21] run in time $2^{\Theta(n)}$ but also require $2^{\Theta(n)}$ memory. The high memory requirement for sieving is inherited from the high memory requirement of a near-neighbour search (NNS) subroutine, which takes in a database of exponentially many vectors and outputs a database of shorter vectors. Feeding the output of this NNS routine to itself a polynomial number of times realises the sieve [BDGL16, AGPS20]. Quantum computers seem not to alter this picture dramatically [LMvdP13, ANS18, AGPS20, CL21, BBTV24]; neither does relaxing SVP to an approximate variant within polynomial factors.

Some works [BLS16, ACKS21] explore trade-offs between sieve style algorithms in time and memory $(\tau, \mu) = \left(2^{\Theta(n)}, 2^{\Theta(n)}\right)$ and enumeration style algorithms in time and memory $(\tau, \mu) = \left(n^{\Theta(n)}, \mathsf{poly}(n)\right)$, and establish somewhat smooth curves interpolating between the two. Despite significant effort, there seems to be a barrier to finding short vectors in time $n^{o(n)}$ with $\mathsf{poly}(n)$ memory, in both the worst and average case. This is discussed in e.g. [MV10b, HPS11, ADRS15]. Indeed, [LV20, Cor. 2.2] – pointing to a 25-year history of studying polynomial-space algorithms for GapSVP – conjectured the non-existence of an $n^{o(n)}$ time $\mathsf{poly}(n)$ memory algorithm for GapSVP, a decision version of SVP, to within any fixed $\mathsf{poly}(n)$ approximation factor. Since GapSVP immediately reduces to SIVP via a transference theorem [Ban93, Thm. 2.2] with a factor n loss in the approximation factor, this conjecture translates to SIVP.

Here, we focus on the SIS problem, which is an average-case analogue of approximate SVP over q-ary lattices, and its inhomogeneous variant the ISIS problem. Let $n \in \mathbb{N}^+$, let $m \in \mathbb{Z}$ with $m \geq n$, $q \in \mathbb{Z}_{\geq 2}$ and $\beta \in \mathbb{R}^+$ be functions of n. A SIS instance is a uniform $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. A solution is a non zero vector $\mathbf{u} \in \mathbb{Z}^m$ such that $\mathbf{A} \cdot \mathbf{u} = \mathbf{0}$ and $\|\mathbf{u}\| \leq \beta$, that is, a short non zero vector in the rank m lattice $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{u} \in \mathbb{Z}^m \colon \mathbf{A} \cdot \mathbf{u} = \mathbf{0}\}$. For ISIS a solution is $\mathbf{u} \in \mathbb{Z}^m$ such that $\mathbf{A} \cdot \mathbf{u} = \mathbf{t}$ and $\|\mathbf{u}\| \leq \beta$ for uniform $\mathbf{t} \in \mathbb{Z}_q^n$, that is a short vector in the coset of $\Lambda_q^\perp(\mathbf{A})$ where $\{\mathbf{u} \in \mathbb{Z}^m \colon \mathbf{A} \cdot \mathbf{u} = \mathbf{t}\}$.

Write $\mathsf{SIS}_{n,m,q,\beta}$ to represent the SIS problem and (τ, μ)-$\mathsf{SIS}_{n,m,q,\beta}$ for the set of algorithms that solve $\mathsf{SIS}_{n,m,q,\beta}$ in time τ, memory μ and with non negligible success probability. The standard SIS assumption states that $(\mathsf{poly}(n), \mathsf{poly}(n))$-$\mathsf{SIS}_{n,m,q,\beta}$ is empty. We are interested in $(2^{O(m)}, \mathsf{poly}(m))$-$\mathsf{SIS}_{n,m,q,\beta}$. For technical reasons, we express (τ, μ) as functions of m, while noting that any algorithm solving (τ, μ)-$\mathsf{SIS}_{n,m,q,\beta}$ can be used to solve (τ, μ)-$\mathsf{SIS}_{n,m',q,\beta}$ when $m \leq m'$.

It has been established [Ajt96, MR07] that solving (τ, μ)-$\mathsf{SIS}_{n,m,q,\beta}$ solves SIVP in time $\mathsf{poly}(\tau)$ and space $\mathsf{poly}(\mu)$ within certain polynomial approximation factors: [MR07] obtained such a reduction for $m, \beta \in \mathsf{poly}(n)$, $q \geq 8n\sqrt{m}\beta$ and approximation factor $\gamma = 8\beta\sqrt{n} \cdot \omega(\sqrt{\log n})$. If we further restrict to $m \in o(n \log n)$ then any algorithm in $\left(2^{O(m)}, \mathsf{poly}(m)\right)$-$\mathsf{SIS}_{n,m,q,\beta}$ implies a $n^{o(n)}$

time and $\mathsf{poly}(n)$ memory algorithm for SIVP to within a $\mathsf{poly}(n)$ approximation factor, violating the conjecture in [LV20, Cor. 2.2]. Thus, the latter implies the following conjecture on SIS:

Conjecture 1 (Space-time hardness of SIS.) *If $m \in o(n \log n)$, $\beta \in \mathsf{poly}(n)$ and $q \geq 8n\sqrt{m}\beta$ then $\left(2^{O(m)}, \mathsf{poly}(m)\right)$-$\mathsf{SIS}_{n,m,q,\beta} = \emptyset$.*

Remark 1. When $q \in \mathsf{poly}(n)$ this implies $m \in o(n \log q)$. There are no known algorithms even when $q \in 2^{\mathsf{poly}(n)}$ and $m = o(n \log q)$ instead of $m = o(n \log n)$. This is despite significant cryptanalytic effort: the cost of the high memory requirements of lattice sieving algorithms are an active area of study to establish security levels for post-quantum schemes, e.g. [Ber16, BBC+20, AS17, NIS23, Jaq24, Sch24]. For $m = \Omega(n \log q)$ the conjecture would not hold (§3.12).

Hinted assumptions. There is a tension between the expressiveness and assumed hardness of a computational problem, where more structured hardness assumptions give rise to more advanced or efficient schemes, but also plausibly provide additional avenues to invalidate the assumption. A classical example of this tension is pairing-based cryptography in comparison to cryptography that relies solely on the Diffie–Hellman assumption.

A traditional way to add structure to a problem is to include additional 'hints' and conjecture that the presence of these hints does not weaken the problem. These hints typically encode some algebraic structure that enables new functionalities. Since it has become increasingly apparent that the standard SIS and LWE assumptions are not expressive enough for some advanced applications or to fulfil stringent efficiency requirements, many recent works have explored such hinted assumptions in order to achieve more advanced functionalities: functional commitments [ACL+22, WW23], polynomial commitments [FMN24], blind signatures [AKSY22], anonymous credentials [BLNS23, DKLW25], attribute-based encryption [Wee22, Wee24] and indistinguishability obfuscation [CLW25, HJL25], to name but a few.

Unlike for SIS and LWE problems, however, little is known about the hardness of these hinted lattice problems, in complexity-theoretic and cryptanalytic terms. In particular, it is unclear how to obtain reductions from worst-case problems such as SIVP to these hinted problems.

1.1 Contributions

We study how to use the space-time hardness of SIS (Conjecture 1) to prove conditional hardness for hinted variants of SIS and ISIS.

Hinted SIS with non-trivial space-time $\leq$ Hinted ISIS. Our first main result is of the following form.

Theorem 1 (Informal, cf. Theorem 6 and Lemma 29). *Let $1 < \gamma_\uparrow, \gamma_\downarrow \leq \mathsf{polylog}(m)$ be growth and shrink factors. Let $\mathcal{A}^{\mathsf{H\text{-}ISIS}}$ be a $\mathsf{poly}(m)$ time algorithm which takes as input $(\mathbf{A}, \mathbf{t}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ and $k \leq \mathsf{poly}(m)$ hints, which are vectors in $\Lambda_q^\perp(\mathbf{A})$ of norm at least $\beta/\gamma_\uparrow$, and outputs a solution to the ISIS instance $(\mathbf{A}, \mathbf{t})$ of norm at most β. Then there exists a $2^{O(m \log(\gamma_\uparrow \cdot \gamma_\downarrow))}$ time and $\mathsf{poly}(m)$ memory algorithm $\mathcal{A}^{\mathsf{H\text{-}SIGS}}$ which, on input $\mathbf{A}$ and the same hints, finds a set of solutions $\{\mathbf{x}_1, \ldots, \mathbf{x}_\ell\}$ to the SIS instance $\mathbf{A}$, each of norm at most $\beta/(\gamma_\uparrow \cdot \gamma_\downarrow)$ such that $\{\mathbf{x}_1, \ldots, \mathbf{x}_\ell\}$ follow a zero-centred spherical discrete Gaussian distribution and generate the lattice $\Lambda_q^\perp(\mathbf{A})$.*

Taking $\gamma_\uparrow \cdot \gamma_\downarrow$ to be a constant yields a $2^{O(m)}$ time $\mathsf{poly}(m)$ memory algorithm and $\beta/(\gamma_\uparrow \cdot \gamma_\downarrow)$ is shorter than the hints by a factor of $\gamma_\downarrow$.

SIS with non-trivial space-time $\leq$ chained Hinted ISIS. Suppose $\mathcal{A}^{\mathsf{H\text{-}ISIS}}$ behaves well when given as new hints the shorter generating set produced by $\mathcal{A}^{\mathsf{H\text{-}SIGS}}$ using $\mathcal{A}^{\mathsf{H\text{-}ISIS}}$ itself. We may formalise this, in the spirit of e.g. lncGDD [MR07], as $\mathcal{A}^{\mathsf{H\text{-}ISIS}}$ being an algorithm that – within limits – takes hints of some norm and outputs solutions of norm at most $\gamma_\uparrow$ longer. Alternatively, we can formalise this as a family (a 'chain') of $\mathcal{A}_i^{\mathsf{H\text{-}ISIS}}$ that on input hints of norm at least $\beta_i/\gamma_\uparrow$, for a sequence of β_i, output solutions of norm at most β_i. While the former is the natural generalisation of sieving algorithms to our setting, we choose the latter, more explicit, formulation. Either choice is a stronger assumption than assuming $\mathcal{A}^{\mathsf{H\text{-}ISIS}}$ for one particular absolute β; see below.

Theorem 2 (Informal, cf. Theorem 7). *Let $\gamma_\uparrow > 1$. Let $s_1 \geq s_2 \geq \ldots \geq s_z > 0$ be a sequence of Gaussian parameters, where $z \leq \mathsf{poly}(m)$. For $i \in [z]$, suppose there exists a sequence of $2^{O(m \log \gamma_\uparrow)}$ time $\mathsf{poly}(m)$ memory algorithms $\mathcal{B}_i$ which, on input $k \leq \mathsf{poly}(m)$ vectors sampled from $D_{\Lambda_q^\perp(\mathbf{A}),s_i}$, samples a Gaussian generating set of $\Lambda_q^\perp(\mathbf{A})$ whose vectors closely follow $D_{\Lambda_q^\perp(\mathbf{A}),s_{i+1}}$. Then there exists a $2^{O(m \log \gamma_\uparrow)}$ time and $\mathsf{poly}(m)$ memory algorithm $\mathcal{C}$ which, on input the Gaussian hints with parameter s_1, finds a generating set of $\Lambda_q^\perp(\mathbf{A})$ whose vectors are of norm at most β_z.*

Then, given a SIS instance $\mathbf{A}$, by considering the Hermite normal form basis of $\Lambda_q^\perp(\mathbf{A})$ we may sample (long) hints for $\mathcal{A}_1^{\mathsf{H\text{-}ISIS}}$ to start the chain. Composing the above, we obtain the following result.

Theorem 3 (Informal, cf. §§ 8 and 10). *Let $z \leq \mathsf{poly}(m)$, $\gamma_\uparrow \leq \mathsf{polylog}(m)$, $s_1 > q \cdot \omega(\sqrt{\log m})$ and $\beta_z \geq \Omega(\sqrt{\log m \cdot m^3}) \cdot q^{n/m}$. For $i \in [z-1]$, let β_i and s_i be such that $\beta_i \leq \gamma_\uparrow \cdot \sqrt{m} \cdot s_i$ and $s_i/s_{i+1} \approx \gamma_\downarrow = C > 1$. Let $\{\mathcal{A}_i\}_{i\in[z]}$ be $\mathsf{poly}(m)$ time algorithms such that with probability $\varepsilon_i \geq 1/\mathsf{poly}(m)$ they return a solution of norm bound β_i for an ISIS instance $(\mathbf{A}, \mathsf{U}(\mathbb{Z}_q^n))$ given $k \leq \mathsf{poly}(m)$ hints from $D_{\Lambda_q^\perp(\mathbf{A}),s_i}$. Then there exists an algorithm solving SIS on $\mathbf{A}$ with norm bound β_z running in time $\tau \leq 2^{O(m \log \gamma_\uparrow)}$ using memory $\mu \in \mathsf{poly}(m)$ and succeeding with probability $\varepsilon \geq 1/\mathsf{poly}(m)$.*

Remark 2. Theorem 3 implies that the algorithm constructed in Theorem 1 asymptotically outperforms known algorithms (§3.12) for some choice of parameters. There must be a minimal index $j \in [z]$ such that we can produce the required hints from $D_{\Lambda_q^\perp(\mathbf{A}),s_j}$ but known algorithms cannot solve $\mathsf{H\text{-}ISIS}$ for norm bound β_j. Otherwise, we can instantiate Theorem 3 to invalidate Conjecture 1. Thus, there exist $\mathsf{H\text{-}ISIS}$ parameters (and not just chains of such parameters) where the absence of $\mathsf{H\text{-}ISIS}$ adversaries follows from the absence of $\mathsf{SIS}_{n,m,q,\beta_j}$ adversaries in time $2^{O(m \log \gamma_\uparrow)}$ and polynomial memory. In other words, there must exist some β_j such that making $\gamma_\downarrow = s_j/s_{j+1}$ progress costs $2^{\omega(m \log \gamma_\uparrow)}$ time for all known polynomial memory algorithms. We note, however, that our reduction does not establish this β_j. In particular, while all cryptanalytic evidence suggests that the problem does not get easier as β_j shrinks, our reductions do not formally establish this.

Techniques. Several fine-grained improvements over existing techniques are required to prove the aforementioned statements.

§4 relies on the Motzkin–Straus theorem from graph theory [MS65] to translate an argument about the maximum number of points that fit into a ball of radius β without being within distance β/γ, into an argument about the probability of two such points being within distance β/γ when sampled from *any* distribution.

§5 shows that sufficiently many samples from a sequence of Gaussian distributions over a lattice Λ with arbitrary centres generate the lattice Λ with high probability. This generalises and tightens an existing result [HR14] on zero centred Gaussians.

§6 provides a rejection-sampling lemma which turns samples of non zero centred Gaussians to samples of zero centred Gaussians. We note that the results of §6 suggest that we do not require §5. However, the presence of §5 allows for a more modular proof and may be of independent interest.

§7 defines the hinted variants of SIS and ISIS, called H-SIS and H-ISIS respectively.[4] We also give a variant of the existing reduction from SIS to ISIS for the hinted setting which takes a form compatible with an approach that we dub 'sieving with centres'. In particular, by not taking some difference but outputting a pair, we can reason about the distribution of the output.

§8 formalises our main results highlighted above. To obtain the first result, i.e. using a hinted ISIS adversary to find a short generating set in non-trivial space-time, we use that the distinct outputs of the many runs of the SIS to ISIS reduction are distributed according to Gaussian distributions with varying centres. We then call rejection sampling on exponentially many candidates to output samples from a zero-centred Gaussian distribution. We then instantiate our chain.

§9 proves probabilistic bounds on the first minimum λ_1 of $\Lambda_q^\perp(\mathbf{A})$ for uniform $\mathbf{A}$. In particular, let $m \geq cn$ for some $c > 1$, $m \in o(n \log q)$ and $q \geq n^\gamma$ for $\gamma > c/2(c-1)$. Except with exponentially small probability in n we establish that $\lambda_m(\Lambda_q^\perp(\mathbf{A})) \in O(\sqrt{\log m}) \cdot \lambda_1(\Lambda_q^\perp(\mathbf{A}))$. This allows us to describe the quality (shortness) of the output of our chain.

Applications. Our results are inspired by the problem of double signing in GPV signatures [GPV08]. A GPV signature consists of a short $\mathbf{u} \in \mathbb{Z}_q^m$ such that $\mathbf{t} = \mathbf{A} \cdot \mathbf{u}$, where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is a uniformly sampled verification key and $\mathbf{t} = H(\mathsf{msg}) \in \mathbb{Z}_q^n$ is a uniform target (in the Random Oracle Model) derived from the message being signed. A forgery in this framework thus amounts to the ISIS problem: find a short preimage $\mathbf{u}$ under uniform $\mathbf{A}$ of uniform target $\mathbf{t}$. The security proof argues that, *provided a single signature is required per message*, a signing adversary $\mathcal{A}$ allows one to compute two short preimages of $\mathbf{t}$, one sampled in the reduction and one returned by $\mathcal{A}$. This solves SIS for $\mathbf{A}$ via subtraction. The proof makes critical use of entropy in the ISIS solutions to show that solving ISIS (solve for $\mathbf{t}$) solves SIS (solve for $\mathbf{0}$).

The security proof cannot provide different signatures for the same message, and does not apply when this is required. Also, the signing key is a particular collection of SIS solutions: a short full rank $\mathbf{Z} \in \mathbb{Z}^{m \times m}$ satisfying $\mathbf{0} = \mathbf{A} \cdot \mathbf{Z}$. Thus, signing the same $\mathbf{t}$ twice and producing likely distinct $\mathbf{u}_i$ such that $\mathbf{t} = \mathbf{A} \cdot \mathbf{u}_i$ may reveal an equivalent $\mathbf{Z}'$ via $\mathbf{u}_i - \mathbf{u}_j$, potentially allowing forgeries and not merely invalidating the security proof.[5] This leads to strongly worded warnings [LP21, OPF+22].

Three approaches exist to prevent requiring multiple signatures per message: (a) derandomised signatures such that the same message always produces the same signature; (b) stateful signing algorithms that refuse to sign the same message twice; and (c) salted messages such that $\mathbf{t} = H(\mathsf{msg}, r)$ for some uniform salt r. The most popular solution for signatures is to salt the message,[6] but essentially the same problem exists for constructions of other primitives based on the GPV framework, such as identity-based encryption (IBE) and aggregate signature schemes (ASS). Unlike in the signature setting, salting is no longer viable. It is thus natural to ask if revealing two preimages of the same image in GPV-style constructions enables attacks.

A tentative negative answer was given in [AKSY22] which introduced the One-More-ISIS (omISIS) assumption. A key ingredient of this tentative answer is that the required preimages are allowed to have norms only marginally longer than the provided hints, which are short preimages of zero with respect to $\mathbf{A}$. As a corollary, we can then argue heuristically that the above mentioned forgery against traditional GPV signature schemes is also hard. Indeed, we may rephrase the cryptanalytic

[4] Compared to k-SIS [BF11, LPSS14], H-ISIS allows the hints $\{\mathbf{u}_i\}_{1 \leq i \leq k}$ to span $\Lambda_q^\perp(\mathbf{A})$ and thus can only plausibly be hard for β small relative to the norms of the hints, see Proposition 5.

[5] Note $\mathbf{u}_i - \mathbf{u}_j$ is close to a zero-centred Gaussian distribution on $\Lambda_q^\perp(\mathbf{A})$, motivating our assumption that $\mathcal{A}$ accepts hints from such a distribution.

[6] This approach is taken in FALCON [PFH+22, Sec. 2.2.2].

task of finding such a forgery given short preimages of zero as the H-ISIS problem above and we show in Appendix A that if H-ISIS is easy then omISIS is too.

The most immediate application of this work then is that the existence of a family of algorithms that forge signatures in a GPV-style signature scheme with norm $\gamma_\uparrow \cdot \|\mathbf{u}\|$ given repeated signatures of norm at most $\|\mathbf{u}\|$ implies algorithms for solving hard lattice problems in $\mathsf{poly}(n)$ memory and in time $2^{O(n)}$ or $n^{o(n)}$ when $\gamma_\uparrow \in O(1)$ or $\gamma_\uparrow \in \mathsf{polylog}\,(n)$ respectively. The conjectured impossibility of such algorithms can thus be used to rule out such forgeries.

We expect that the thus established conjecture – that salting is plausibly not necessary in GPV-style signatures – will have applications in the construction of more advanced authentication-style primitives such as (algebraic) lattice-based IBE, ABE and ASS.

2 Technical overview

We summarise the technical steps obtaining our main results.

2.1 Hinted variants of SIS and ISIS

The H-SIS, H-ISIS and H-SIGS problems. We define hinted variants of SIS and ISIS, called H-SIS and H-ISIS, in §7. Both have parameters $\mathsf{params} = (k, m, q, \beta, \mathsf{Dist})$ parametrised by n, with non-negative integer k, rank m where $n \leq m$, modulus q, norm bound $\beta > 0$, and distribution Dist mapping $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ to a distribution over $\Lambda_q^\perp(\mathbf{A})^k$. We define H-ISIS:

> H-ISIS: Given $\mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n \times m})$, target $\mathbf{t} \leftarrow \mathsf{U}(\mathbb{Z}_q^n)$ and hints $\{\mathbf{u}_i\}_{1 \leq i \leq k}$ from $\mathsf{Dist}(\mathbf{A})$, find $\mathbf{u}^* \in \mathbb{Z}^m$ s.t. $\mathbf{t} = \mathbf{A} \cdot \mathbf{u}^*$ and $\|\mathbf{u}^*\| \leq \beta$.

In H-SIS, $\mathbf{t} = \mathbf{0}$ and $\mathbf{u}^*$ must be non-zero. We also define the hinted short integer generating set (H-SIGS) problem. A solution to H-SIGS is a short generating set of $\Lambda_q^\perp(\mathbf{A})$.

For H-ISIS, we focus on the regime where $\|\mathsf{Dist}(\mathbf{A})\| \geq \beta/\gamma_\uparrow$ with high probability for some $1 < \gamma_\uparrow \leq \mathsf{polylog}(m)$, so the problem is non-trivial even when $m \ll k \leq \mathsf{poly}(m)$. Here, $\|\mathsf{Dist}(\mathbf{A})\|$ means the norm of the vectors following the distribution $\mathsf{Dist}(\mathbf{A})$. For H-SIS, solving is trivial for any $k \geq 1$ since $\|\mathbf{u}_1\| \leq \beta$ suffices: output $\mathbf{u}^* = \mathbf{u}_1$, or a short linear combination of $\{\mathbf{u}_i\}_{1 \leq i \leq k}$. However, such methods yield few distinct solutions. A non-trivial requirement is to return H-SIS solutions from an entropic distribution beyond such combinations.

2.2 Technical tools

To convert H-ISIS adversaries into $(2^{O(m)}, \mathsf{poly}(m))$ algorithms finding shorter H-SIS solutions, we introduce some technical tools in §§ 5 and 6.

Arbitrarily centred Gaussian vectors generate the lattice. To chain reductions from H-SIGS to H-ISIS, we show that short vectors from a H-SIGS adversary, each following a discrete Gaussian, form a generating set of $\Lambda_q^\perp(\mathbf{A})$. In §5, we generalise a result in [HR14] showing that $\tilde{O}(m^2)$ zero-centred Gaussian vectors in rank m lattice Λ span it with high probability under mild conditions, to the case of arbitrarily centred Gaussians. We also tighten the reduction to $\tilde{O}(m)$ vectors.

Rejection sampling. Our core reduction outputs vectors distributed close to $D_{\Lambda_q^\perp(\mathbf{A}),s,\mathbf{c}_i}$ with norm approximately β/γ where $\|\mathbf{c}_i\| \approx \beta/2\gamma$. In §6, we prove rejection sampling can 'clean up' this distribution to zero-centred $D_{\Lambda_q^\perp(\mathbf{A}),s}$ in single-exponential time and polynomial memory when $\|\mathbf{c}_i\|/s \in O(\sqrt{m})$, which holds here. Our results extend [DFPS22]. There, the smooth Rényi divergence $R_\infty^\varepsilon(D_{\mathbb{Z}^m,s} \| D_{\mathbb{Z}^m,s,\mathbf{c}})$ is established. We generalise this to arbitrary Λ straightforwardly. They also give a theorem bounding rejection sampling's running time and statistical distance — when run until outputting not $\perp$ — as a function of $R_\infty^\varepsilon(P\|Q)$ for target P and source Q. However, [DFPS22] does not establish $R_\infty^\varepsilon(D_{\mathbb{Z}^m,s} \| D_{\mathbb{Z}^m,s,C})$ when centres $\mathbf{c}_i$ vary with C.[7] Instead, in our application we bound the number of rejection sampling trials since we aim for a single-exponential time algorithm with $\mathsf{negl}(n)$ failure probability, not an expected single-exponential one with no failure. This allows us to bound the output's statistical distance via a simple hybrid argument.

[7] If $R_\infty^\varepsilon(P\|Q_i) \leq M$ we cannot conclude e.g. $R_\infty^\varepsilon(P\|(Q_0 + Q_1)/2) \leq M$.

Sampling near-neighbours We build an algorithm $\mathcal{B}$ that runs $\mathcal{A}$ many times to find t distinct solutions, aiming to locate a pair satisfying a predicate (a short difference). To ensure $\mathsf{poly}(m)$ memory, $\mathcal{B}$ cannot store all t candidates since $t = 2^{\Omega(m)}$ may be needed. We provide in §4 a gadget converting known results on the number of 'not-close' points in a ball of radius β into a statement about finding two close points when sampled from an *arbitrary* distribution. The key ingredient is the Motzkin–Straus theorem [MS65].

2.3 Polynomial memory 'sieving'

Equipped with these tools, we build polynomial memory sieving-style algorithms from H-ISIS adversaries. Our construction has two steps below.

Step 1. Reduction from H-SIGS to H-ISIS. We first turn a H-ISIS adversary into a H-SIS adversary where the latter samples solutions $\gamma_\downarrow$ times shorter than the hints.

A first idea is to run the H-SIS adversary built from the H-ISIS adversary many times, producing solutions of norm at most β. Assuming the hints have norm at least $\beta/\gamma_\uparrow$, solutions are at most $\gamma_\uparrow$ times longer than the hints. One can show that, indeed, many runs yield $2^{O(m\log(\gamma_\uparrow\cdot\gamma_\downarrow))}$ distinct correct solutions. Then, by a counting argument, some pair has difference norm at least $\gamma_\uparrow\cdot\gamma_\downarrow$ times shorter than the solutions, thus at least $\gamma_\downarrow$ times shorter than the hints.

A first issue with this idea is that the output distribution of the algorithm is hard to analyse, making it difficult to compose recursively. To address this, we open up the above reduction from H-SIS to H-ISIS. Recall that the reduction samples $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$, sets $\mathbf{t} = \mathbf{A}\cdot\mathbf{v}$, sends $(\mathbf{A},\mathbf{t})$ to the H-ISIS adversary, and receives $\mathbf{v}'$ such that $\mathbf{A}\cdot(\mathbf{v}-\mathbf{v}') = \mathbf{0}$. Instead of returning $\mathbf{v}-\mathbf{v}'$, we return the pair $(\mathbf{v}',\mathbf{v})$. Then, we consider pairs of $\mathbf{v}'_i$ to get a pair $(\mathbf{v}'_\ell,\mathbf{v}'_r)$ with short difference $\mathbf{v}'_\ell-\mathbf{v}'_r$. Critically, $\mathbf{w} = (\mathbf{v}_\ell-\mathbf{v}'_\ell)-(\mathbf{v}_r-\mathbf{v}'_r)$ is distributed negligibly close to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s,\mathbf{v}'_r-\mathbf{v}'_\ell}$. Another effect of this change, which we call 'sieving with centres', is that $\mathbf{v}'_r = \mathbf{v}'_\ell$ is allowed, since with high probability $\mathbf{v}_\ell \neq \mathbf{v}_r$, ensuring $\mathbf{w}\neq\mathbf{0}$. Overall, we construct an algorithm repeatedly sampling $(\mathbf{v}'_i,\mathbf{v}_i)$ for $i\in\{\ell,r\}$ and accepting a pair when $\|\mathbf{v}'_\ell-\mathbf{v}'_r\| \le \beta/(2\gamma_\uparrow\cdot\gamma_\downarrow)$ which succeeds with probability at least $2^{-O(m\cdot\log(\gamma_\uparrow\cdot\gamma_\downarrow))}$ by Motzkin–Straus.

Given we can now reason about the output distribution, we establish the smooth Rényi divergence to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$, the ideal distribution for GPV-style signatures. This allows us to show that for some parameter choices – satisfied by our reduction – we can use rejection sampling with time $2^{O(m)}$ and $\mathsf{poly}(m)$ memory to output vectors close to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$. Thus, given an algorithm $\mathcal{A}^{\mathsf{H\text{-}ISIS}}$ consuming zero-centred Gaussian hints, we construct $\mathcal{A}^{\mathsf{H\text{-}SIS}}$ outputting shorter zero-centred Gaussian hints.

We then turn this H-SIS adversary into a H-SIGS adversary. The idea is to run it many times, each sampling from $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$. Correctness follows from the lemma that Gaussian samples generate the lattice with high probability.

Step 2. Chaining H-SIGS adversaries. Given H-SIGS adversaries achieving a $\gamma_\downarrow$ factor improvement for $1<\gamma_\downarrow\le\mathsf{polylog}(m)$, our second step chains them to achieve larger improvements. Let $s_i/s_{i+1}=\gamma_\downarrow$ and suppose $\mathcal{A}_1^{\mathsf{H\text{-}SIGS}},\dots,\mathcal{A}_z^{\mathsf{H\text{-}SIGS}}$ are H-SIGS adversaries for $\mathsf{Dist}_i(\mathbf{A}) = D_{\Lambda_q^\perp(\mathbf{A}),s_i}$ that output generating sets following $D_{\Lambda_q^\perp(\mathbf{A}),s_{i+1}}$ (as above) but with probability $\varepsilon_i\ge 1/\mathsf{poly}(n)$ on $\Lambda_q^\perp(\mathbf{A})$. We show we can amplify success to overwhelming, allowing composition into an algorithm $\mathcal{B} := (\mathcal{A}_1^{\mathsf{H\text{-}SIGS}},\dots,\mathcal{A}_z^{\mathsf{H\text{-}SIGS}})$ that accepts hints of norm $\beta_1/\gamma_\uparrow$ but outputs solutions of norm $\beta_z/\gamma_\downarrow$.

Finally, to jump-start the chain, we can provide the first adversary hints long enough to sample efficiently given $\mathbf{A}$. We thus turn the chain of H-SIGS adversaries into a SIGS adversary, which is also a SIS adversary. Assuming all H-SIGS adversaries behave well, chaining stops when outputs are too

short to (a) provide sufficient entropy to assume $\mathbf{v}_\ell \neq \mathbf{v}_r$, or (b) form a generating set. Avoiding the first failure condition requires $\beta \geq \tilde{\Omega}(\sqrt{m})$ and avoiding the second requires $\beta \geq \Omega(m) \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A}))$, cf. §5. In §9 we establish that $\lambda_m(\Lambda_q^\perp(\mathbf{A})) \in O(\sqrt{\log m}) \cdot \lambda_1(\Lambda_q^\perp(\mathbf{A}))$ for parameters of interest. For example, with $m = 5n$ and $q \approx n^2$, we obtain $\beta \geq \Omega(\sqrt{\log n} \cdot n^{19/10})$.

Remark 3. We could drop the condition that the hints must span a generating set since we assume our H-ISIS adversaries expect Gaussian hints without any stipulation that these generate $\Lambda_q^\perp(\mathbf{A})$. However, it is plausible that such adversaries would only work for parameters where this condition is also met. Thus, we conservatively assume this here.

2.4 Open problems

Recursion. Our reduction assumes a family of algorithms $\mathcal{A}_i$ for a given $\mathbf{A}$ rather than a single algorithm. Relaxing this is an important open problem.

Cosets. Our results relate to omISIS [AKSY22]. While our reduction does not formally establish the hardness of omISIS even assuming single-exponential time polynomial memory algorithms do not exist, it suggests all known cryptanalytic approaches rely on solving a hard problem. Our reduction is phrased in terms of finding short vectors in a lattice given short vectors in it. It can plausibly be generalised to finding short vectors in a coset given short vectors in some cosets of the lattice. All known cryptanalytic approaches for omISIS fit this framework. We leave this generalisation for future work.

Number rings. We state our results for lattices over $\mathbb{Z}$ but cryptographic constructions often use modules over rings of integers of number fields. Generalising to this setting would cover these applications.

3 Preliminaries

We give an overview of notation used throughout in Table 1. The natural numbers are $\mathbb{N} = \{0, 1, 2, \dots\}$ and $\mathbb{N}^+ = \mathbb{N} \setminus \{0\}$. For $k \in \mathbb{N}^+$ let $[k] = \{1, \dots, k\}$. Let $n \in \mathbb{N}$ denote the security parameter. Algorithms implicitly take 1^n as input. Logarithms written log have base two, the natural logarithm is written ln. Column vectors are written $\mathbf{x} \in \mathbb{R}^n$ and matrices $\mathbf{S} \in \mathbb{R}^{n \times m}$. We write $\hat{\mathbf{S}} = (\hat{\mathbf{s}}_1 \cdots \hat{\mathbf{s}}_m)$ for the Gram–Schmidt orthogonalisation of $\mathbf{S}$. We write $\|\mathbf{x}\|$ for the Euclidean norm and $\|\mathbf{S}\|_{\max}$ and $\|\mathbf{S}\|_{\min}$ for the maximum and minimum Euclidean norm over columns respectively. We often suppress the max subscript and write $\|\mathbf{S}\| = \|\mathbf{S}\|_{\max}$. We may treat matrices as sets of their column vectors where convenient. We use the same conventions for the infinity norm $\|\cdot\|_\infty$. Let $X \subseteq [0, \infty)$ be unbounded and $f, g \colon X \to [0, \infty)$. We write $f(x) \in \alpha(g(x))$ with $\alpha \in \{O, o, \Omega, \omega, \Theta\}$ and $f \sim g$. We may write $f(x) = \alpha(g(x))$, remembering this is a 'one way' equality. We write $f(x) \in 2^{\alpha(g(x))}$ to denote that $\log f(x) \in \alpha(g(x))$, so for example $2^{x^2} \in 2^{\Omega(x)}$ and $2^{-x/2} \in 2^{-\Theta(x)}$. Note that if $f(x) \in \omega(g(x))$ then we cannot conclude $\log f(x) \in \omega(\log g(x))$ but only that $\log f(x) \in \Omega(\log g(x))$. Let $g(n) = 2^{f(n)}$ then we call g exponential if $f \in \Theta(n)$ and superexponential if $f \in \omega(n)$. We also define strongly and weakly superexponential which will be useful for our results; superexponential g is weakly superexponential if $f \in o(n \log n)$ and otherwise strongly superexponential. The gamma function is defined on $\mathbb{C}$ except the non positive integers and is the analytic continuation of $\Gamma(z) = \int_0^\infty t^{z-1} e^{-t}\, dt,\ \Re(z) > 0$.

3.1 Probability and entropy

We implicitly consider discrete random variables throughout. Let $(\Omega, \mathcal{F}, P)$ be a probability space and $(E, \mathcal{E})$ be another measurable space. A random variable X is a measurable function $X \colon \Omega \to E$ with a countable codomain. The probability distribution associated to X is the pushforward of P by X, that is $P_X \colon \mathcal{E} \to [0, 1]$, $A \mapsto P(X^{-1}(A))$. The support of X is the support of P_X as a measure; in the above setup this is exactly the codomain of X. We write $\mathrm{Supp}(X)$ for the support of X and may write simply $x \in X$ when quantifying over elements of it. For random variable X its expectation is $\mathbb{E}[X]$ and its min-entropy is

$$H_\infty(X) = -\log\left(\max_{x \in X} \Pr[X = x]\right) \geq 0.$$

For random variables X and Y their statistical distance is $\Delta(X, Y)$ and the conditional min-entropy of X given Y is

$$H_\infty(X \mid Y) = -\log\left(\sum_{y \in Y} \Pr[Y = y] \cdot \max_{x \in X} \Pr[X = x \mid Y = y]\right) \geq 0.$$

For a finite set S we write $\mathsf{U}(S)$ for the uniform random variable such that $\Pr[\mathsf{U}(S) = s] = 1/|S|$ for all $s \in S$ and write $x \leftarrow S$ to denote a sample from this distribution. We write $x \leftarrow D$ to denote a sample from some distribution D. We may sometimes write the arguments of expectations, statistical distances and (conditional) min entropies as distributions or elements sampled from them when it is clear from context what is meant.

If X is uniformly distributed on $N \in \mathbb{N}$ elements then the probability that $(x_1, \dots, x_t)$ sampled i.i.d. from X are pairwise distinct is $\binom{N}{t} \cdot t!/N^t$ for $t \in \{0\} \cup [N]$. We consider X with lower bounded min entropy and derive a lower bound on the probability of samples being pairwise distinct.

Lemma 1. *Let $H_\infty(X) \geq \alpha$, $N = \lceil 2^\alpha \rceil$ and $t \in \mathbb{N}$. The probability that $(x_1, \dots, x_t)$ sampled i.i.d. from X are pairwise distinct is at least $\binom{N-1}{t} \cdot t!/(N-1)^t$ for $t \in \{0\} \cup [N-1]$ and at least 0 for $t \geq N$.*

Proof. There are at least N elements in the support of X and our lower bound gives probability one to $t = 0$ and $t = 1$, which is correct regardless of X, and the trivial probability zero to $t \geq N$. Henceforth we consider $t \in [N-1] \setminus \{1\}$. The probability that $(x_i)_{i=1}^t$ are pairwise distinct is

$$\prod_{i=1}^{t} (1 - \Pr[X \in \{x_1, \ldots, x_{i-1}\}]) \geq \prod_{i=1}^{t} \left(1 - (i-1) \cdot 2^{-\alpha}\right) \geq \prod_{i=1}^{t} \frac{N-i}{N-1}.$$

Simple manipulations give the lemma. □

The lower bound of Lemma 1 can be unwieldy. We give the following corollary which simplifies matters for some functions $N, t \colon \mathbb{N} \to \mathbb{N}$ such that $t \in o(\sqrt{N})$.

Corollary 1. *Let $N, t \in \mathbb{N}$ such that $t \leq N^\gamma$ for $\gamma \in [0, 1/2)$. We have*

$$\binom{N}{t} \frac{t!}{N^t} \geq 1 - N^{2\gamma - 1}.$$

Proof. A careful unpacking of [Das16, (3)]. □

3.2 Problems and algorithms

Let $\mathcal{P} = \{P_n\}_n$ be a problem ensemble with instance distributions $\mathcal{D} = \{D_n\}_n$ and implicit criteria for solutions. We call the pair $\mathsf{Prob} = (\mathcal{P}, \mathcal{D})$ an (average case) problem and an algorithm that attempts to solve Prob a Prob adversary. We often drop the dependence on n for problem ensembles and distributions. We also denote by ε_p the success probability of $\mathcal{A}$ on a particular problem instance $p \in P$ and may refer to p sampled according to Prob rather than according to D. Typically we consider problems which are parametrised by some parameters $\mathsf{params} = \mathsf{params}(n)$. We denote such a problem by $\mathsf{Prob}_{\mathsf{params}}$. We sometimes consider Prob such that an adversary has access to stateful oracles $\mathsf{O}_1, \ldots, \mathsf{O}_k$ which may depend on params and the problem instance. We write $\mathsf{Prob}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}_{\mathsf{params}} = (\mathcal{P}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}, \mathcal{D})$ in this case. Whether the solution output by the adversary is correct may depend on the state of the oracles.

Definition 1 ((τ, μ, ε) adversary). *Let $\tau, \mu, \varepsilon, \mathsf{params}$ be parametrised by n. A randomised algorithm $\mathcal{A}$ is a (τ, μ, ε) adversary against $\mathsf{Prob}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}_{\mathsf{params}} = (\mathcal{P}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}, \mathcal{D})$ if for all n it holds that*

$$\Pr\left[\mathsf{Exp\text{-}Prob}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}_{\mathsf{params},\mathcal{A}}(1^n) = 1\right] \geq \varepsilon(n),$$

with probability taken over $\mathsf{Exp\text{-}Prob}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}_{\mathsf{params},\mathcal{A}}$ defined in Fig. 1 (i.e. over $\mathsf{O}_1, \ldots, \mathsf{O}_k$, D_n and the random tape of $\mathcal{A}$), and $\mathcal{A}$ runs in time at most $\tau(n)$ using at most $\mu(n)$ memory cells for all $p \in P_n$.

$\mathsf{Exp\text{-}Prob}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}_{\mathsf{params},\mathcal{A}}(1^n)$

1 : $p \leftarrow D_n$
2 : $x \leftarrow \mathcal{A}^{\mathsf{O}_1,\ldots,\mathsf{O}_k}(p)$
3 : **return** $[\![x$ is correct relative to p and the states of $\mathsf{O}_1, \ldots, \mathsf{O}_k]\!]$

Figure 1. Generic experiment for any parametrised average-case problem $\mathsf{Prob}_{\mathsf{params}}$.

3.3 Classical tail bounds

We use specific forms of common probabilistic tools. First a tail bound for non-negative real random variables depending only on expectation.

Proposition 1 (Inverse Markov inequality, [RS92, Loè77]). *Let $a, c \geq 0$ and X be a random variable. If $\Pr[0 \leq X \leq c] = 1$ then $\Pr[X \geq a] \geq (\mathbb{E}[X] - a)/c$.*

Proposition 1 is frequently used as follows in cryptography, similarly to [PS00, Lem. 7]. Let $\mathcal{A}$ be a (τ, μ, ε) adversary against $\mathsf{Prob} = (\mathcal{P}, \mathcal{D})$ and X be the random variable that samples $p \in P$ according to D and returns ε_p. Note $\mathbb{E}[X] = \varepsilon$ and we may therefore set $c = 1$ and $a = \varepsilon/2$ so that $\Pr[X \geq \varepsilon/2] \geq \varepsilon/2$. That is, the success probability of $\mathcal{A}$ on a random problem instance is at least $\varepsilon/2$ with probability at least $\varepsilon/2$.

For $p \in [0, 1]$ let $\mathrm{Ber}(p)$ denote a Bernoulli random variable with parameter p, so that $X \sim \mathrm{Ber}(p)$ has $\Pr[X = 1] = p$ and $\Pr[X = 0] = 1 - p$. Also let $\mathrm{Bin}(N, p)$ denote a binomial random variable of N trials each with success probability p. We give Chernoff's bound.

Lemma 2 (Chernoff's bound, implicit in [MU05, Thm. 4.5, 2.]). *Suppose $X_1, \ldots, X_n \sim \mathrm{Ber}(p)$ are independent and $X = \sum_i X_i$. Then for $\delta \in (0, 1)$ we have*

$$\Pr[X \leq (1 - \delta) \cdot np] \leq \exp(-np \cdot \delta^2/2).$$

We give a utility lemma and corollary for the binomial distribution.

Lemma 3. *Let $X \sim \mathrm{Bin}(N, p)$ for $p \in (0, 1)$. Let $k \in \mathbb{N}^+$ and $c \in (0, 1)$. Let $y = y(k, c) = 2\ln(1/(1 - c))/k$. If $N \geq (k/p) \cdot 2/(2 + y - \sqrt{y^2 + 4y})$ then $\Pr[X \geq k] \geq c$.*

Proof. For $\delta \in (0, 1)$ let $f(N, p, \delta) = \exp(-Np \cdot \delta^2/2)$ then Lemma 2 implies

$$\Pr[X \geq (1 - \delta) \cdot Np] \geq 1 - f(N, p, \delta).$$

We may choose $(N, \delta) \in \mathbb{N}^+ \times (0, 1)$. We solve $(1 - \delta) \cdot Np = k$ and $1 - f(N, p, \delta) = c$ for $\delta \in (0, 1)$, i.e. the appropriate root of $\delta^2 + y\delta - y = 0$. Therefore $\delta = (\sqrt{y^2 + 4y} - y)/2$. That $(1 - \delta) \cdot Np = k$ then implies $N = (k/p) \cdot 2/(2 + y - \sqrt{y^2 + 4y})$ as required. □

We apply the above to $c = 1 - 2^{-k}$. Note that since $k \geq 1$ we have $c \geq 1/2$.

Corollary 2. *Let $X \sim \mathrm{Bin}(N, p)$ for $p \in (0, 1)$ and $k \in \mathbb{N}^+$. If $c = 1 - 2^{-k}$ then $N \geq 4 \cdot k/p$ implies $\Pr[X \geq k] \geq c$.*

Proof. Let $R(k, c) = 2/(2 + y(k, c) - \sqrt{y^2(k, c) + 4y(k, c)})$ where $y(k, c)$ is as defined in Lemma 3. Note that R strictly increases in y. If $c = 1 - 2^{-k}$ then $y(k, c) = 2k\ln(2)/k = 2\ln(2)$. Therefore $R(k, c) < 4$. □

If $k \in \omega(1)$ then the probability of k successes from $N \geq 4 \cdot k/p$ trials tends to one. If $k(n)$ is a parametrised quantity but constant, e.g. $k(n) = 1$, and we require a success probability in $1 - \mathsf{negl}(n)$ then we may select $N = 4n/p$ as $n \geq k$ success immediately implies k successes.

We will make use of the above repetition machinery restricted to 'good' instances.

Lemma 4. *Let $\mathsf{Prob} = (\mathcal{P}, \mathcal{D})$ be a problem. Let $P^{\checkmark} \subseteq P$ and $\alpha, \beta > 0$ be such that $D(P^{\checkmark}) \geq \alpha$ and for $p \in P^{\checkmark}$ we have $\varepsilon_p \geq \beta$. Let $N, t \in \mathbb{N}^+$ and $X_p \sim \mathrm{Bin}(N, \varepsilon_p)$. If $N\beta \geq 4t$ then $\Pr\left[X_p \geq t \wedge p \in P^{\checkmark}\right] \geq \alpha \cdot (1 - 2^{-t}) \geq \alpha/2$.*

Proof. We consider

$$\begin{aligned}\Pr\left[X_p \geq t \wedge p \in P^{\checkmark}\right] &= \Pr\left[X_p \geq t \;\middle|\; p \in P^{\checkmark}\right] \cdot \Pr\left[p \in P^{\checkmark}\right] \\ &\geq \alpha \cdot \Pr\left[X_p \geq t \;\middle|\; p \in P^{\checkmark}\right].\end{aligned}$$

Note $p \in P^{\checkmark}$ implies $\varepsilon_p \geq \beta$ so if $X \sim \mathrm{Bin}(N, \beta)$ then

$$\Pr[X \geq t] \leq \Pr\left[X_p \geq t \;\middle|\; p \in P^{\checkmark}\right]$$

. Setting $p = \beta$ and $k = t$ in Corollary 2 concludes. □

Given $p \in P^{\checkmark}$ this lemma states that if $N, t \colon \mathbb{N} \to \mathbb{N}^+$ and $\beta \colon \mathbb{N} \to [0,1]$ are such that $N\beta$ is eventually at least $4t$ then the probability of having at least t correct outputs is at least one half.

3.4 Graph theory

Definition 2 (Graph and adjacency matrix). *Let V be a non empty set and $E \subseteq \{\{v, w\} \colon v, w \in V\}$. The undirected graph $G = (V, E)$ is said to have vertices V and edges E. We say G is simple if $E \cap \{\{v, v\}\}_{v \in V} = \emptyset$. Letting $V = \{v_i\}_{i \in I}$ its adjacency matrix $\mathbf{M} \in \{0,1\}^{|V| \times |V|}$ has $m_{i,j} = 1$ whenever $\{v_i, v_j\} \in E$.*

Simple graphs have no 'self loops' from a vertex to itself. The adjacency matrix $\mathbf{M}$ is a symmetric binary matrix representing the (non)existence of edges and has a zero diagonal for simple graphs. All graphs we consider are undirected (without repeated edges).

Definition 3 (Graph complement). *Let $G = (V, E)$ be a simple graph then its complement is $\bar{G} = (V, \bar{E})$ with $\bar{E} = \{\{v, w\} \colon v, w \in V \wedge v \neq w\} \setminus E$.*

In particular if $\mathbf{M}$ is the adjacency matrix for a simple $G = (V, E)$ and similarly $\bar{\mathbf{M}}$ for $\bar{G}$ then $\bar{\mathbf{M}} + \mathbf{M} = \mathbf{1}_{|V|} - \mathbf{I}_{|V|}$ where $\mathbf{1}_{|V|}$ denotes the $|V|$-by-$|V|$ all-one matrix. That is, $\bar{G}$ is simple and $\bar{\mathbf{M}}$ differs from $\mathbf{M}$ except on the diagonal where they are both zero.

Definition 4 (Clique and independence number). *Given a graph $G = (V, E)$ we define the clique and independence numbers respectively as*

$$\begin{aligned}\omega(G) &= \max\{|X| \colon X \subseteq V \wedge (\forall (x, y \in X \wedge x \neq y),\, \{x, y\} \in E)\} \\ \alpha(G) &= \max\{|X| \colon X \subseteq V \wedge (\forall (x, y \in X \wedge x \neq y),\, \{x, y\} \notin E)\}\end{aligned}$$

In particular, cliques do not require self loops and independent sets may have them. For simple G we have $\omega(G) = \alpha(\bar{G})$ and $\alpha(G) = \omega(\bar{G})$.

3.5 (Smooth) Rényi divergence

We recall the definition of (exponential) Rényi divergence and the notion of smooth Rényi divergence defined in [DFPS22].

Definition 5. *Let P and Q be discrete probability distributions with supports such that* $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$*. Their (exponential) Rényi divergence of order $\alpha \in (1, \infty)$ is*

$$R_\alpha(P \| Q) \coloneqq \left(\sum_{x \in \mathrm{Supp}(P)} \frac{P(x)^\alpha}{Q(x)^{\alpha - 1}} \right)^{\frac{1}{\alpha - 1}}.$$

Definition 6 (Smooth Rényi divergence, discrete case of [DFPS22, Def. 2.1]). *Let $\varepsilon \geq 0$. Let P, Q be two probability distributions such that $\sum_{x \in \mathrm{Supp}(Q)} P(x) \geq 1 - \varepsilon$. Their ε-smooth Rényi divergence (of infinite order) is*

$$R^{\varepsilon}_{\infty}(P\|Q) \coloneqq \inf\left\{M > 0 : \Pr_{x \leftarrow P}[P(x) \leq M \cdot Q(x)] \geq 1 - \varepsilon\right\}.$$

If $\sum_{x \in \mathrm{Supp}(Q)} P(x) < 1 - \varepsilon$ then $R^{\varepsilon}_{\infty}(P\|Q) \coloneqq +\infty$.

Lemma 5 ([DFPS22, Lemma A.7]). *Let P and Q be discrete probability distributions with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$. For any $\varepsilon > 0$ and order $\alpha \in (1, \infty)$ it holds that*

$$R^{\varepsilon}_{\infty}(P\|Q) \leq \frac{R_{\alpha}(P\|Q)}{\varepsilon^{1/(\alpha-1)}}.$$

3.6 Lattices

We write $\Lambda = \Lambda(\mathbf{B})$ for the Euclidean lattice generated by integer combinations of the columns of $\mathbf{B} \in \mathbb{R}^{m \times d}$. If the columns of $\mathbf{B}$ are linearly independent (so $m \geq d$) we call Λ a rank d lattice and when $m = d$ a full rank lattice. The dual of a full rank lattice Λ is $\Lambda^* = \{\mathbf{y} \in \mathbb{R}^m : \langle \mathbf{y}, \Lambda \rangle \subseteq \mathbb{Z}\}$. If $\mathbf{B}$ is a basis the determinant of $\Lambda(\mathbf{B})$ is $\sqrt{\det(\mathbf{B}^T \cdot \mathbf{B})}$. For full rank lattices Minkowski's theorem implies there exists $\mathbf{x} \in \Lambda \setminus \{\mathbf{0}\}$ such that $\|\mathbf{x}\| \leq \sqrt{m} \cdot \det(\Lambda)^{1/m}$. For $i \in [d]$ write $\lambda_i(\Lambda) = \inf\{r > 0 : \mathrm{span}_{\mathbb{R}}(\mathsf{B}^m(r) \cap \Lambda) \cong \mathbb{R}^i\}$ and $\lambda^{\infty}_i(\Lambda) = \inf\{r > 0 : \mathrm{span}_{\mathbb{R}}(\mathsf{B}^m_{\infty}(r) \cap \Lambda) \cong \mathbb{R}^i\}$ for the ith minima in the Euclidean and infinity norm. We define $\mathsf{bl}(\Lambda) = \min\{\|\mathbf{B}\| : \mathbf{B} \text{ a basis of } \Lambda\}$, the minimum length of a basis of Λ. For full rank lattices we have $\lambda_m(\Lambda) \leq \mathsf{bl}(\Lambda) \leq \sqrt{m} \cdot \lambda_m(\Lambda)/2$ [CN97, Lem. 3]. We define $\mu(\Lambda) = \max\{\|\mathbf{x} - \mathbf{y}\| : \mathbf{x} \in \Lambda, \mathbf{y} \in \mathrm{span}_{\mathbb{R}}(\Lambda)\}$, the covering radius of a lattice. This quantity is the maximum distance a point in the span of the lattice can be from the lattice. We have $\lambda_m(\Lambda) \leq 2 \cdot \mu(\Lambda) \leq \sqrt{m} \cdot \lambda_m(\Lambda)$ [MG02, Thm. 7.9].

We also define two q-ary lattices given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with $m \geq n$:[8]

$$\Lambda^{\perp}_q(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0}\},$$
$$\Lambda_q(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x} \bmod q \in \mathbf{A}^T \cdot \mathbb{Z}_q^n\}.$$

They are (full) rank m lattices, contain $q\mathbb{Z}^m$, have $\det(\Lambda^{\perp}_q(\mathbf{A})) \leq q^n$ and $\det(\Lambda_q(\mathbf{A})) \geq q^{m-n}$ and are such that $\Lambda^{\perp}_q(\mathbf{A})^* = \frac{1}{q}\Lambda_q(\mathbf{A})$. We write $\Lambda^{\perp}_q(m, n)$ to denote the random variable $\Lambda^{\perp}_q(\mathsf{U}(\mathbb{Z}_q^{n \times m}))$ and similarly $\Lambda_q(m, n)$ to denote $\Lambda_q(\mathsf{U}(\mathbb{Z}_q^{n \times m}))$. Let $\varphi_{\mathbf{A}} : \mathbb{Z}^m \to \mathbb{Z}_q^n$, $\mathbf{u} \mapsto \mathbf{A} \cdot \mathbf{u}$. The next lemma tells us parameters (n, m, q) for which $\varphi_{\mathbf{A}}$ is likely surjective. For prime q we call $\mathbf{A}$ such that $\varphi_{\mathbf{A}}$ is surjective full rank. Let $F_{n,m,q} = \{\mathbf{A} \in \mathbb{Z}_q^{n \times m} : \mathbf{A} \text{ is full rank}\}$.

Proposition 2. *Let $n, m, q \in \mathbb{N}^+$ with $m \geq n$ and prime q then*

$$\Pr\left[\mathsf{U}(\mathbb{Z}_q^{n \times m}) \in F_{n,m,q}\right] \geq 1 - q^{n-m}.$$

Proof. Let $\mathbf{r}_0 = \mathbf{0} \in \mathbb{R}^{1 \times m}$ and $\mathbf{r}_i$ be the ith row of $\mathbf{A}$ for $i \in [n]$. A given $\mathbf{A}$ is full rank if and only if the row rank of $\mathbf{A}$ is n. Therefore

$$\begin{aligned}
\Pr\left[\mathsf{U}(\mathbb{Z}_q^{n \times m}) \in F_{n,m,q}\right] &= \Pr\left[\mathbf{r}_1 \notin \mathrm{span}_{\mathbb{Z}_q}(\mathbf{r}_0) \wedge \cdots \wedge \mathbf{r}_n \notin \mathrm{span}_{\mathbb{Z}_q}(\mathbf{r}_0, \ldots, \mathbf{r}_{n-1})\right] \\
&= 1 - \Pr\left[\mathbf{r}_1 \in \mathrm{span}_{\mathbb{Z}_q}(\mathbf{r}_0) \vee \cdots \vee \mathbf{r}_n \in \mathrm{span}_{\mathbb{Z}_q}(\mathbf{r}_0, \ldots, \mathbf{r}_{n-1})\right] \\
&\geq 1 - q^{-m} - q^{1-m} - \cdots - q^{n-1-m} \geq 1 - q^{n-m}
\end{aligned}$$

□

[8] While e.g. $\mathbf{A} \cdot \mathbf{x} = \mathbf{0} \bmod q$ is conventional, as the columns of $\mathbf{A}$ are elements of $(\mathbb{Z}_q^n, +)$ multiplication by an integer is unambiguous and the mod q redundant.

For prime q, $\mathbf{A}$ is full rank $\iff \det(\Lambda_q^\perp(\mathbf{A})) = q^n \iff \det(\Lambda_q(\mathbf{A})) = q^{m-n}$.

Generally, we aim to produce shorter bases rather than merely full rank linearly independent sets of lattice vectors. However, for some applications such as discrete Gaussian sampling the latter suffices. For this we rely on the following utility lemma.

Lemma 6 ([MG02, Lemma 7.1]). *There is a polynomial time algorithm that on input a lattice basis $\mathbf{B}$ and a full rank set of linearly independent lattice vectors $\mathbf{S} \subset \Lambda(\mathbf{B})$ such that $\|\mathbf{s}_1\| \leq \|\mathbf{s}_2\| \leq \cdots \leq \|\mathbf{s}_m\|$, outputs a basis $\mathbf{R}$ for $\Lambda(\mathbf{B})$ such that $\|\mathbf{r}_i\| \leq \max\{(\sqrt{i}/2), 1\} \cdot \|\mathbf{s}_i\|$ for all i. Moreover, the new basis satisfies $\operatorname{span}(\mathbf{r}_1, \ldots, \mathbf{r}_i) = \operatorname{span}(\mathbf{s}_1, \ldots, \mathbf{s}_i)$ and $\|\hat{\mathbf{r}}_i\| \leq \|\hat{\mathbf{s}}_i\|$ for all i.*

3.7 Gaussians

For $\mathbf{c} \in \mathbb{R}^m$ and $s \in \mathbb{R}_{>0}$ the (spherical) Gaussian function with parameter s and centre $\mathbf{c}$ is $\rho_{s,\mathbf{c}} \colon \mathbb{R}^m \to \mathbb{R}$, $\mathbf{x} \mapsto \exp\left(-\pi \cdot \|\mathbf{x} - \mathbf{c}\|^2 / s^2\right)$. The discrete Gaussian distribution over lattice $\Lambda \subset \mathbb{R}^m$ with parameter $s \in \mathbb{R}_{>0}$ and centre $\mathbf{c}$ is $D_{\Lambda,s,\mathbf{c}} \colon \Lambda \to \mathbb{R}$, $\mathbf{x} \mapsto \rho_{s,\mathbf{c}}(\mathbf{x})/\rho_{s,\mathbf{c}}(\Lambda)$, where $\rho_{s,\mathbf{c}}(\Lambda) = \sum_{\mathbf{x} \in \Lambda} \rho_{s,\mathbf{c}}(\mathbf{x})$. When $\mathbf{c} = \mathbf{0}$ we omit it.

We define some helper sets, $P_{\mathbf{A},\mathbf{t}}$, $P_{\beta,\mathbf{A},\mathbf{t}}$, $P_{\mathbf{A}}^+$ and $P_{\beta,\mathbf{A}}^+$. Here P stands for preimages; $P_{\mathbf{A},\mathbf{t}}$ is the preimages of $\mathbf{t}$ under $\mathbf{A}$, i.e. $\mathbf{u}$ such that $\mathbf{A} \cdot \mathbf{u} = \mathbf{t}$. A subscript β implies length bound β and a superscript $+$ implies only vectors with positive length.

Definition 7. *Let $(\mathbf{A}, \mathbf{t}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ and $\beta \in \mathbb{R}_{>0}$. Define*

$$P_{\mathbf{A},\mathbf{t}} = \begin{cases} \{\mathbf{u} \in \mathbb{Z}^m \colon \mathbf{A} \cdot \mathbf{u} = \mathbf{t}\} & \text{if } \mathbf{t} \in \operatorname{im}(\varphi_{\mathbf{A}}), \\ \emptyset & \text{otherwise,} \end{cases}$$

$$P_{\beta,\mathbf{A},\mathbf{t}} = P_{\mathbf{A},\mathbf{t}} \cap \mathsf{B}^m(\beta).$$

If $\mathbf{t} = \mathbf{0}$ we omit it. Let $P_{\mathbf{A}}^+ = P_{\mathbf{A}} \setminus \{\mathbf{0}\}$ and $P_{\beta,\mathbf{A}}^+ = P_{\beta,\mathbf{A}} \setminus \{\mathbf{0}\}$.

Note $P_{\mathbf{A}} = \Lambda_q^\perp(\mathbf{A})$. If $\mathbf{t} \in \operatorname{im}(\varphi_{\mathbf{A}})$ then $P_{\mathbf{A},\mathbf{t}} = \mathbf{u} + P_{\mathbf{A}}$ for any $\mathbf{u} \in P_{\mathbf{A},\mathbf{t}}$.

We also define the discrete Gaussian distribution over $\mathbb{Z}^m$ given auxiliary information $(\mathbf{A}, \mathbf{t}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ and possibly a length bound $\beta \in \mathbb{R}_{>0}$. The restricted distributions are over preimages of $\mathbf{t}$, namely $\mathbf{u} \in \mathbb{Z}^m$ such that $\mathbf{A} \cdot \mathbf{u} = \mathbf{t}$, and possibly such that $\|\mathbf{u}\| \leq \beta$.

Definition 8. *Let $(\mathbf{A}, \mathbf{t}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ with $\mathbf{t} \in \operatorname{im}(\varphi_{\mathbf{A}})$. Then*

$$D_{\mathbb{Z}^m,s,\mathbf{c}}[\mathbf{A}, \mathbf{t}] \colon P_{\mathbf{A},\mathbf{t}} \to \mathbb{R},\ \mathbf{u} \mapsto \rho_{s,\mathbf{c}}(\mathbf{u})/\rho_{s,\mathbf{c}}(P_{\mathbf{A},\mathbf{t}}).$$

If $\mathbf{t} = \mathbf{0}$ or $\mathbf{c} = \mathbf{0}$ we omit it.

Note $D_{\mathbb{Z}^m,s,\mathbf{c}}[\mathbf{A}] = D_{\Lambda_q^\perp(\mathbf{A}),s,\mathbf{c}}$ and $D_{\mathbb{Z}^m,s}[\mathbf{A}, \mathbf{t}] = \mathbf{u} + D_{\Lambda_q^\perp(\mathbf{A}),s,-\mathbf{u}}$ for any $\mathbf{u} \in P_{\mathbf{A},\mathbf{t}}$.

Lemma 7 (Implicit in [Lyu11, Lem. 4.4(iii)]). *For $\alpha \geq 1$, $s \in \mathbb{R}_{>0}$ and full rank $\Lambda \subset \mathbb{R}^m$*

$$\Pr\left[\|D_{\Lambda,s}\| \geq \alpha \cdot s\sqrt{\frac{m}{2\pi}}\right] < \alpha^m \cdot \exp\left(m \cdot (1 - \alpha^2)/2\right).$$

Note $\alpha^m \cdot \exp\left(m \cdot (1 - \alpha^2)/2\right) = \exp(m \cdot (1 + 2\ln(\alpha) - \alpha^2)/2) = \exp(m \cdot f(\alpha))$. If $\alpha > 1$ then $f(\alpha) < 0$ and if $f(\alpha) \leq -\ln 2$ then $\exp(m \cdot f(\alpha)) \leq 2^{-m}$. We have $\alpha \geq 1.93 \Rightarrow f(\alpha) \leq -\ln 2$.

Corollary 3. *For $s \in \mathbb{R}_{>0}$ and full rank $\Lambda \subset \mathbb{R}^m$*

$$\Pr\left[\|D_{\Lambda,s}\| \geq s\sqrt{m}\right] < 2^{-m}.$$

Smoothing parameter. The smoothing parameter with $\delta \in (0,1)$ of a full rank lattice Λ is $\eta_\delta(\Lambda) = \min\{s > 0 \colon \rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \delta\}$. We may bound $\eta_\delta(\Lambda)$ via m, δ and $\lambda_1^\infty(\Lambda^*)$.

Lemma 8 ([Pei08, Lem. 3.5]). *For a full rank m lattice Λ and $\delta \in (0,1)$*

$$\eta_\delta(\Lambda) \leq \frac{1}{\lambda_1^\infty(\Lambda^*)} \cdot \sqrt{\frac{\ln(2m \cdot (1 + 1/\delta))}{\pi}}.$$

Note that replacing ln with log in the upper bounds of Lemmas 8 and 10 (below) give inessentially less tight upper bounds on $\eta_\delta(\Lambda)$. We do so throughout to avoid multiple logarithmic bases. We also introduce the notation $\eta_\delta(m) = \sqrt{\log(2m(1+1/\delta))/\pi}$ to later simplify some expressions, noting $\eta_\delta(\mathbb{Z}^m) \leq \eta_\delta(m)$.

By bounding $\lambda_1^\infty(\Lambda_q^\perp(\mathbf{A})^*)$ below we therefore bound $\eta_\delta(\Lambda_q^\perp(\mathbf{A}))$ above as a function of only m and δ. This is achieved probabilistically. The lemma below considers the probability that an element of $\Lambda_q(m,n)$ contains a non zero point in an open cube of particular side length and then uses the relationship $\Lambda_q^\perp(\mathbf{A})^* = \frac{1}{q}\Lambda_q(\mathbf{A})$. It is a slightly more thorough exploration of [GPV07, Lem. 5.3] that corrects a small mistake. Indeed, if $q = 1 \bmod 4$ then the set Z defined in [GPV07, Lem. 5.3] has size (exactly) $((q+1)/2)^m$ rather than at most $(q/2)^m$. If $q = 2$ or $q = 3 \bmod 4$ then [GPV07, Lem. 5.3] is correct.

Lemma 9 (Implicit in [GPV07, Lem. 5.3]). *Let $n, q \in \mathbb{N}^+$ with q prime, $m \geq n$ and $f \colon \mathbb{N} \to [1, q/4)$. Then*

$$\begin{aligned}\Pr\left[\lambda_1^\infty(\Lambda_q^\perp(m,n)^*) < \frac{1}{4f(m)}\right] &< \frac{1}{2^m} \cdot \frac{q^n}{f^m(m)} \cdot \left(1 + \frac{2f(m)}{q}\right)^m \\ &< \left(\frac{3}{4}\right)^m \cdot \frac{q^n}{f^m(m)}.\end{aligned}$$

Proof. Let $\mathsf{C}^m(q,f) = (-q/(4f(m)), q/(4f(m)))^m$, $\mathsf{Z} = \mathsf{C}^m(q,f) \cap \mathbb{Z}^m$ and $N = |\mathsf{Z}|$. Note $N = |(-q/(4f(m)),\ q/(4f(m))) \cap \mathbb{Z}|^m$. For $x > 0$ we have $|(-x,x) \cap \mathbb{Z}| = 2\lceil x \rceil - 1 < 2x + 1$. Setting $x = q/(4f(m)) > 1$ gives $N < (q/(2f(m)) + 1)^m$.

Fix $\mathbf{v} \in \mathbb{Z}_q^n \setminus \{\mathbf{0}\}$ and note that no elements of Z are equivalent modulo q. For a uniform $\mathbf{A}$ the probability that $\mathbf{A}^T \cdot \mathbf{v} \in \mathsf{Z}$ is N/q^m. Taking a union bound over $\mathbf{v}$ we have

$$\begin{aligned}\Pr[\lambda_1^\infty(\Lambda_q(m,n)) < q/4f(m)] &\leq N \cdot (q^n - 1)/q^m \\ &< (q/(2f(m)) + 1)^m \cdot q^n/q^m \\ &= \frac{1}{2^m} \cdot \frac{q^n}{f^m(m)} \cdot \left(1 + \frac{2f(m)}{q}\right)^m.\end{aligned}$$

Noting that $\Lambda_q^\perp(\mathbf{A})^* = \frac{1}{q}\Lambda_q(\mathbf{A})$ concludes the proof. □

Note that if $f \geq q/4$ then we would be asking for the probability of the infinity norm first minimum to be less than $1/q$. Given that $\Lambda_q^\perp(\mathbf{A})^* = \frac{1}{q}\Lambda_q(\mathbf{A})$ and $\Lambda_q(\mathbf{A}) \subseteq \mathbb{Z}^m$ this probability is zero.

The factor $(1 + 2f(m)/q)^m$ captures the fact that $q/(4f(m))$ can deviate arbitrarily from the integers for general f, rather than just by quarters when $f = 1$. Ignoring this factor and the erroneous $q = 1 \bmod 4$ case, we recover the lemma of [GPV07, Lem. 5.3] with $m \geq 2n \log q$ and $f = 1$. The utility of f is that it allows one to consider smaller m with respect to n at the cost of a larger upper bound on the smoothing parameter, see Corollary 4.

Throughout let $C_{f,n} = q^n \cdot (3/4)^m \cdot f^{-m}(m)$, the upper bound of Lemma 9. We define the set of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ that are smoothed by a choice of f.

Definition 9. *Let $n, m, q \in \mathbb{N}^+$ with q prime and $m \geq n$. Let $f\colon \mathbb{N} \to [1, q/4)$ and define $\mathsf{SMOOTH}_f(n) \subseteq \mathbb{Z}_q^{n\times m}$ as the $\mathbf{A}$ such that for all $\delta \in (0,1)$*

$$\eta_\delta(\Lambda_q^\perp(\mathbf{A})) \leq 4 \cdot f(m) \cdot \sqrt{\frac{\log(2m \cdot (1+1/\delta))}{\pi}}.$$

Combining Lemmas 8 and 9 we have the following.

Corollary 4. *Let $n, m, q \in \mathbb{N}^+$ with q prime and $m \geq n$. Let $f\colon \mathbb{N} \to [1, q/4)$ then $|\mathsf{SMOOTH}_f(n)| \geq (1 - C_{f,n}) \cdot q^{mn}$, i.e. comprises at least a $1 - C_{f,n}$ fraction of $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$.*

In particular, for any $s(m) \in f(m) \cdot \omega(\sqrt{\log m})$ there exists a $\delta(m) \in \mathsf{negl}(m)$ such that for any $\mathbf{A} \in \mathsf{SMOOTH}_f(n)$ we have $\eta_{\delta(m)}(\Lambda_q^\perp(\mathbf{A})) \leq s(m)$. We also use an alternative bound on $\eta_\delta(\Lambda)$.

Lemma 10 (Smoothing parameter bound [MR07, Lemma 3.3]). *For any full rank m lattice Λ and $\delta > 0$,*

$$\eta_\delta(\Lambda) \leq \sqrt{\frac{\ln(2m \cdot (1+1/\delta))}{\pi}} \cdot \lambda_m(\Lambda).$$

In particular, for any $s(m) \in \omega(\sqrt{\log m}) \cdot \lambda_m(\Lambda)$ there exists a $\delta(m) \in \mathsf{negl}(m)$ such that we have $\eta_{\delta(m)}(\Lambda) \leq s(m)$. Given the introduction of the smoothing parameter we have a tail bound for discrete Gaussians with arbitrary centres.

Lemma 11 ([MR07, Lemma 4.4]). *For $\mathbf{c} \in \mathbb{R}^m$, $\delta \in (0,1)$ and full rank $\Lambda \subset \mathbb{R}^m$ if $s \geq \eta_\delta(\Lambda)$ then*

$$\Pr\left[\|D_{\Lambda,s,\mathbf{c}} - \mathbf{c}\| \geq s\sqrt{m}\right] < 2^{-m} \cdot \frac{1+\delta}{1-\delta}.$$

Lemma 12 links the min-entropy of a discrete Gaussian over a lattice to its dual and s.

Lemma 12 (Implicit in [PR05, Lem. 2.16]). *Let Λ be a full rank m lattice, $\mathbf{c} \in \mathbb{R}^m$ and $s \geq \eta_\delta(\Lambda)$ then $H_\infty(D_{\Lambda,s,\mathbf{c}}) \geq \log(s^m \cdot \det(\Lambda^*) \cdot (1-\delta))$.*

We use a bound on the Gaussian mass of lattice cosets.

Lemma 13 ([GPV07, Lemma 2.7]). *Let $\Lambda \subset \mathbb{R}^m$ be full rank, $\mathbf{c} \in \mathbb{R}^m$ and $s \geq \eta_\delta(\Lambda)$ for some $\delta \in (0,1)$ then*

$$\rho_{s,\mathbf{c}}(\Lambda) \in \left[\frac{1-\delta}{1+\delta}, 1\right] \cdot \rho_s(\Lambda).$$

We make use of an algorithm sampling $D_{\mathbb{Z},s}$ in our results.

Lemma 14 (Discrete Gaussian sampling over $\mathbb{Z}$ [BLP+13b, Sec. 5.1]). *For any $s \in \mathbb{R}_{>0}$ and $c \in \mathbb{R}$ there exists a space and time efficient randomised algorithm $\mathsf{PreSamp}$ that outputs a sample according to $D_{\mathbb{Z},s,c}$ with probability at least one half.*

In this work we often desire $N \in 2^{o(m \log m)}$ samples from $D_{\mathbb{Z}^m,s}$. In Corollary 5 all N output a sample from $D_{\mathbb{Z}^m,s}$ with probability superexponentially close to one and each sampling procedure remains efficient.

Corollary 5. *Let $N \in 2^{o(m \log m)}$. There exists a procedure Samp that makes $N \cdot m^3$ calls to $\mathsf{PreSamp}$ and outputs N samples according to $D_{\mathbb{Z}^m,s}$ with probability in $1 - 2^{-\Omega(m^2)}$.*

Proof. Let $\mathsf{PreSamp}^{m^2}$ denote the process of making m^2 calls to $\mathsf{PreSamp}$. The probability $\mathsf{PreSamp}^{m^2}$ fails to return a sample according to $D_{\mathbb{Z},s}$ is at most 2^{-m^2}. Therefore, by a union bound, $N \cdot m$ calls to $\mathsf{PreSamp}^{m^2}$ returns N samples according to $D_{\mathbb{Z}^m,s}$ except with probability at most $N \cdot m \cdot 2^{-m^2} \in 2^{-\Omega(m^2)}$. □

Lemma 15 ([GPV08, Theorem 4.1]). *There is a probabilistic polynomial time algorithm that, given a basis* $\mathbf{B}$ *of a rank* m *lattice* Λ*, a parameter* $s \geq \left\|\hat{\mathbf{B}}\right\| \cdot \omega(\sqrt{\log m})$ *and a center* $\mathbf{c} \in \mathbb{R}^m$ *outputs a sample from a distribution that is statistically close to* $D_{\Lambda,s,\mathbf{c}}$*.*

The following is a straightforward composition of Lemma 6 and Lemma 15.

Corollary 6. *There is a probabilistic polynomial time algorithm that, given a basis* $\mathbf{B}$ *of a rank* m *lattice* Λ*, a full rank set* $\mathbf{S} \subset \Lambda$*, a parameter* $s \geq \left\|\hat{\mathbf{S}}\right\| \cdot \omega(\sqrt{\log m})$ *and a center* $\mathbf{c} \in \mathbb{R}^m$ *outputs a sample from a distribution that is statistically close to* $D_{\Lambda,s,\mathbf{c}}$*.*

The following simple corollary allows us to establish that we can sample from $\Lambda_q^\perp(\mathbf{A})$ with some appropriately large parameter s.

Corollary 7. *There is a probabilistic polynomial time algorithm that, given a matrix* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$*, a parameter* $s \geq q \cdot \omega(\sqrt{\log m})$ *and a center* $\mathbf{c} \in \mathbb{R}^m$ *outputs a sample from a distribution that is statistically close to* $D_{\Lambda_q^\perp(\mathbf{A}),s,\mathbf{c}}$*.*

Proof. Note that $q\mathbf{I}_m$ is a full rank set of vectors $\mathbf{S}$ in $\Lambda_q^\perp(\mathbf{A})$ with $\left\|\tilde{\mathbf{S}}\right\| = q$. Given $\mathbf{A}$ we may construct a basis $\mathbf{B}$ of $\Lambda_q^\perp(\mathbf{A})$ via linear algebra. Apply Corollary 6. □

Given that $D_{\Lambda,s,\mathbf{c}} = \mathbf{c} + D_{\Lambda-\mathbf{c},s}$ setting $s_1 = s_2 \geq \sqrt{2} \cdot \eta_\delta(\Lambda)$ and $\mathbf{z} = (1\ 1)$ in [GMPW20, Thm. 4] gives the following.

Lemma 16 (Specialised from [GMPW20, Thm. 4]). *Let* $\Lambda \subseteq \mathbb{R}^m$ *be a lattice. Let* $s \geq \sqrt{2} \cdot \eta_\delta(\Lambda)$ *for some* $\delta \in \mathsf{negl}(n)$*. Let* $\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{R}^m$ *be arbitrary. If* $\mathbf{x}_1 \leftarrow D_{\Lambda,s,\mathbf{c}_1}$ *and* $\mathbf{x}_2 \leftarrow D_{\Lambda,s,\mathbf{c}_2}$ *then* $\mathbf{x}_1 + \mathbf{x}_2$ *is within statistical distance at most* 3δ *from* $D_{\Lambda,\sqrt{2}s,\mathbf{c}_1+\mathbf{c}_2}$*.*

3.8 Leftover hash lemma

We use the following specialisation of a more general leftover hash lemma theorem. In particular, we consider only diagonal covariance matrices, uniform $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and let $(K, R) = (\mathbb{Q}, \mathbb{Z})$.

Theorem 4 (Specialisation of [BL26, Thm. 6]). *Let* $n, m, q \in \mathbb{N}^+$ *with* q *prime and* $m \geq n$*. Let* $p_{\text{im}} = \Pr\left[\mathsf{U}(\mathbb{Z}_q^{n \times m}) \notin F_{n,m,q}\right]$*. For* $\delta \in (0, 1/2)$ *and* $s > 0$ *let* $p_\eta(\delta, s) = \Pr\left[\eta_\delta(\Lambda_q^\perp(m, n)) > s\right]$*. Let* $P \sim (\mathbf{A}, \mathbf{Ax})$ *for* $\mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n \times m})$ *and* $\mathbf{x} \leftarrow D_{\mathbb{Z}^m,s}$ *and* $U \sim \mathsf{U}(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n)$*, then*

$$\Delta(P, U) \leq 2\delta + p_{\text{im}} + p_\eta(\delta, s).$$

Corollary 8. *Let* $n, m, q \in \mathbb{N}^+$ *with* q *prime and* $m \geq n$ *and let* $f \colon \mathbb{N} \to [1, q/4)$*. For* $\delta \in (0, 1/2)$ *let* $s \geq 4f(m) \cdot \sqrt{\log(2m(1 + 1/\delta))/\pi}$ *then for* (P, U) *as in Theorem 4 we have*

$$\Delta(P, U) \leq 2\delta + 1/q^{m-n} + C_{f,n}.$$

Proof. Proposition 2 and Corollary 4 give upper bounds on p_{im} and $p_\eta(\delta, s)$ respectively. □

In particular, if $\delta \in \mathsf{negl}(m)$, $m \geq 2n$ and $f(m) \geq q^{n/m}$ then this statistical distance is $\mathsf{negl}(m)$.

3.9 Upper bounding $\mu(\Lambda_q^\perp(m,n))$

We appeal to [RSD24, Lem. 6.1] and give a simple probabilistic upper bound on the covering radius. As $\lambda_m(\Lambda) \leq 2\mu(\Lambda)$ we obtain bounds on the final minimum.

Corollary 9. *Let (m,q) be parametrised by n with $m \geq n$ and q prime. Let*

$$B(n) = 4(1+\sqrt{m/2\pi}) \cdot q^{n/m} \cdot \sqrt{\log(6m)/\pi}.$$

It holds that

$$\Pr\left[\mu(\Lambda_q^\perp(m,n)) \geq B(n)\right] \leq (3/4)^m.$$

Proof. Set $f(m) = q^{n/m}$ in Corollary 4 so that $C_{f,n} = (3/4)^m$ therein. Set $\delta = 1/2$ in Definition 9 so that $\eta_\delta(\Lambda_q^\perp(\mathbf{A})) \leq 4 \cdot q^{n/m} \cdot \sqrt{\log(6m)/\pi}$ except with probability $C_{f,n}$ for uniform $\mathbf{A}$. Finally, $s \geq \eta_{1/2}(\Lambda) \Rightarrow \rho_{1/s}(\Lambda^*) \leq 3/2$ and conclude by [RSD24, Lem. 6.1]. □

This generic probabilistic upper bound on μ is a square root logarithmic factor larger than the Gaussian heuristic suggests for the first minimum. For certain parameters, e.g. $m = 2n$ and sufficiently large $q \in \mathsf{poly}(n)$, one can prove a probabilistic upper bound on μ that is only a constant factor larger than the Gaussian heuristic via transference [Ban93, Thm. 2.2] and the statistical closeness of $\Lambda_q^\perp(m, m-n)$ and $\Lambda_q(m,n)$ [Mic11]. We do not require this result, and omit the details.

3.10 Computational problems

Two presumed hard computational problems on lattices are SVP and SIVP along with their approximate variants.

Definition 10 (SVP). *Let $\mathbf{B} \in \mathbb{R}^{n\times n}$ be full rank and $\gamma\colon \mathbb{N} \to \mathbb{R}_{\geq 1}$. A solution to the γ-*approximate shortest vector problem *(γ-SVP) is a vector $\mathbf{v} \in \Lambda(\mathbf{B}) \setminus \{\mathbf{0}\}$ such that $\|\mathbf{v}\| \leq \gamma(n) \cdot \lambda_1(\Lambda(\mathbf{B}))$.*

Definition 11 (SIVP). *Let $\mathbf{B} \in \mathbb{R}^{n\times n}$ be full rank and $\gamma\colon \mathbb{N} \to \mathbb{R}_{\geq 1}$. A solution to the γ-*approximate shortest independent vectors problem *(γ-SIVP) is a set of n linearly independent lattice vectors $S \subset \Lambda(\mathbf{B})$ such that $\|S\| \leq \gamma(n) \cdot \lambda_n(\Lambda(\mathbf{B}))$.*

In this work we focus on average-case analogues of the above problems over $\Lambda_q^\perp(m,n)$. In particular, an average-case analogue of SVP over $\Lambda_q^\perp(m,n)$ is known as the (homogeneous) short integer solution problem, introduced in the seminal work of Ajtai [Ajt96]. A solution is a short vector in $P_{\mathbf{A}}^+$ for given uniform $\mathbf{A}$. An inhomogeneous version was formalised later [Mic07]. A solution is a short vector in $P_{\mathbf{A},\mathbf{t}}$ for given uniform $(\mathbf{A},\mathbf{t})$. We also define natural average-case analogues of SIVP where an adversary must find either a short full rank set or a short generating set of $\Lambda_q^\perp(\mathbf{A})$. We call these the short independent integer solutions (SIIS) problem and the short integer generating set (SIGS) problem respectively.

Definition 12 (SIS ([Ajt96]), ISIS ([Mic07]), SIIS and SIGS). *Let $\mathsf{params} = (m,q,\beta)$ be parametrised by n with $m,q \in \mathbb{N}^+$, $m \geq n$, $q \geq 2$ and $\beta > 0$. The experiments of the $\mathsf{SIS}_{\mathsf{params}}$, $\mathsf{ISIS}_{\mathsf{params}}$, $\mathsf{SIIS}_{\mathsf{params}}$ and $\mathsf{SIGS}_{\mathsf{params}}$ problems are defined in Fig. 2.*

Remark 4. There are trivial reductions from SIS to SIIS and SIIS to SIGS as these problems are defined via (absolute) norm bounds β rather than approximation factors. In §9 we show if $2n \leq m \in o(n\log q)$ and $q \geq n^2$ then $\lambda_m(\Lambda_q^\perp(\mathbf{A})) \leq \lambda_1(\Lambda_q^\perp(\mathbf{A})) \cdot O(\sqrt{\log m})$ with overwhelming probability over the choice of $\mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n\times m})$. As such a γ-SIVP solution for $\Lambda_q^\perp(\mathbf{A})$ likely gives m distinct $O(\gamma \cdot \sqrt{\log m})$-SVP solutions.

$\mathsf{Exp\text{-}SIS}_{(m,q,\beta),\mathcal{A}}(1^n)$

1 : $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$
2 : $\mathbf{u} \leftarrow \mathcal{A}(\mathbf{A})$
3 : // **return** $[\![\mathbf{u} \in P^+_{\beta,\mathbf{A}}]\!]$
4 : **return** $[\![\mathbf{A}\cdot\mathbf{u} = \mathbf{0} \wedge 0 < \|\mathbf{u}\| \leq \beta]\!]$

$\mathsf{Exp\text{-}SIIS}_{(m,q,\beta),\mathcal{A}}(1^n)$

1 : $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$
2 : $\mathbf{U} \leftarrow \mathcal{A}(\mathbf{A})$
3 : // **return** $[\![\mathbf{U} \subset P_{\beta,\mathbf{A}} \wedge \mathsf{rank}(\mathbf{U}) = m]\!]$
4 : **return** $[\![\mathbf{A}\cdot\mathbf{U} = \mathbf{0} \wedge \|\mathbf{U}\| \leq \beta \wedge \mathsf{rank}(\mathbf{U}) = m]\!]$

$\mathsf{Exp\text{-}ISIS}_{(m,q,\beta),\mathcal{A}}(1^n)$

1 : $(\mathbf{A},\mathbf{t}) \leftarrow \mathbb{Z}_q^{n\times m} \times \mathbb{Z}_q^n$
2 : $\mathbf{u} \leftarrow \mathcal{A}(\mathbf{A},\mathbf{t})$
3 : // **return** $[\![\mathbf{u} \in P_{\beta,\mathbf{A},\mathbf{t}}]\!]$
4 : **return** $[\![\mathbf{A}\cdot\mathbf{u} = \mathbf{t} \wedge \|\mathbf{u}\| \leq \beta]\!]$

$\mathsf{Exp\text{-}SIGS}_{(m,q,\beta),\mathcal{A}}(1^n)$

1 : $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$
2 : $\mathbf{U} \leftarrow \mathcal{A}(\mathbf{A})$
3 : // **return** $[\![\mathbf{U} \subset P_{\beta,\mathbf{A}} \wedge \Lambda(\mathbf{U}) = \Lambda_q^\perp(\mathbf{A})]\!]$
4 : **return** $[\![\mathbf{A}\cdot\mathbf{U} = \mathbf{0} \wedge \|\mathbf{U}\| \leq \beta \wedge \Lambda(\mathbf{U}) = \Lambda_q^\perp(\mathbf{A})]\!]$

Figure 2. The experiments for SIS, ISIS, SIIS and SIGS.

An algorithm solving ISIS can be used to solve SIS with a longer length bound by considering the difference of a sampled $\mathbf{u}$ such that $\mathbf{A}\cdot\mathbf{u} = \mathbf{t}$ and a solution $\mathbf{u}'$ of the ISIS instance $(\mathbf{A}, \mathbf{t})$. We will show a reduction from a problem called one more inhomogeneous short integer solution [AKSY22] (omISIS) to a problem introduced in §7. We recall that omISIS does not enjoy a known reduction from a standard lattice problem.

Definition 13 (omISIS [AKSY22]). *Let* $\mathsf{params} = (m, q, \beta, s)$ *be parametrised by* n*. The experiment of the* $\mathsf{omISIS}^{\mathsf{syn,pre}}_{\mathsf{params}}$ *problem is defined in Fig. 3.*

$\mathsf{omISIS}^{\mathsf{syn,pre}}_{(m,q,\beta,s),\mathcal{A}}(1^n)$

$k = 0$
$S = \emptyset$
$\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$
$\{(\mathbf{y}_i, \mathbf{t}_i)\}_{i\in[k+1]} \leftarrow \mathcal{A}^{\mathsf{syn,pre}}(\mathbf{A})$
$b_0 = \wedge_i [\![\mathbf{y}_i \in P_{\beta,\mathbf{A},\mathbf{t}_i} \wedge \mathbf{t}_i \in S]\!]$
$b_1 = \wedge_{i\neq j} [\![\mathbf{y}_i \neq \mathbf{y}_j]\!]$
return $b_0 \wedge b_1$

$\mathsf{syn}()$

$\mathbf{t} \leftarrow \mathbb{Z}_q^n$
$S = S \cup \{\mathbf{t}\}$
return $\mathbf{t}$

$\mathsf{pre}(\mathbf{t})$

$\mathbf{u} \leftarrow D_{\mathbb{Z}^m,s}[\mathbf{A},\mathbf{t}]$
$k = k + 1$
return $\mathbf{u}$

Figure 3. The omISIS game.

3.11 Finding short or close vectors

We recall algorithms for solving hard lattice problems from the literature. The first is an algorithm that on input a basis $\mathbf{B}$ and some $\mathbf{t}$ in its span, outputs a point $\mathbf{s}$ that is equivalent modulo the lattice generated by $\mathbf{B}$. The closeness of lattice point $\mathbf{t} - \mathbf{s}$ is determined by the Gram–Schmidt basis associated to $\mathbf{B}$.

Lemma 17 ([Bab86]). *Let* $\mathbf{B} \in \mathbb{R}^{n\times r}$ *be a basis of* Λ*,* $\mathbf{t} \in \mathrm{span}(\mathbf{B}) \subset \mathbb{R}^n$ *and*

$$\mathcal{P}_{1/2}(\hat{\mathbf{B}}) = \left\{ \sum_{i=1}^{r} a_i \cdot \hat{\mathbf{b}}_i : a_i \in (-1/2, 1/2] \right\}.$$

Then there exists a polynomial time algorithm that finds $\mathbf{s} \in \mathcal{P}_{1/2}(\hat{\mathbf{B}})$ *such that* $\mathbf{s} = \mathbf{t} \bmod \Lambda$.

Enumeration. A series of works [Pho81, Kan83, FP85, SE94, MW15, ABF+20, ABLR21] solve the following problem: given basis $\mathbf{B} \in \mathbb{R}^{n \times r}$ and length bound R, find all $\mathbf{v} = \sum_{i=1}^{r} u_i \cdot \mathbf{b}_i$ with $u_i \in \mathbb{Z}$ such that $\|\mathbf{v}\| \leq R$. Returning a shortest non zero vector encountered solves the shortest vector problem. Thus, enumeration can be thought of as an exhaustive search in a ball where projection is used to reduce the search space. Several works have also explored techniques to reduce the search space further using probabilistic arguments and yielded exponential but not superexponential speed ups [GNR10, ABF+20]. Enumeration runs in $n^{n/(2e)+o(n)}$ time and polynomial memory [HS07]. It has been shown heuristically that when enumeration is used as the SVP oracle inside blockwise lattice reduction (see §3.12) the time is reduced to $n^{n/8+o(n)}$ [ABF+20]. Note that enumeration requires strongly superexponential time. In particular, lower bounds for (heuristic) enumeration-type algorithms were provided and discussed in [ANSS18, ABF+20], suggesting smaller leading constants compared to known algorithms but no avenue for beating the $n^{\Theta(n)}$ complexity.

Sieving. The works of [AKS01, MV10b, ADRS15, BGJ15, Laa15a, BDGL16, HK17] take as input a list $L \subset \Lambda$ of lattice points and search for integer combinations of these points that are short. Thus, sieving can be thought of as a 'collision-finding' type algorithm, inherently required to consider sufficiently many elements to encounter such solutions. This has been formalised as conceptualising sieving algorithms as a repeated application of a near-neighbour search algorithm (NNS) [BDGL16, AGPS20]. In particular, if the initial list is exponentially large in n, shorter elements can be found and SVP can be solved by performing this NNS process recursively. Each point in the initial list can be sampled at a cost polynomial in n [Kle00]. Hence the initial list can be sampled at a cost of $|L|^{1+o(1)}$ by repeating this process $|L|$ times.

Sieves that combine k points at a time are called k-sieves; 2-sieves take integer combinations of the form $\mathbf{u} \pm \mathbf{v}$ with $\mathbf{u}, \mathbf{v} \in L$ and $\mathbf{u} \neq \pm\mathbf{v}$. Heuristic sieving algorithms are analysed under the assumption, introduced in [NV08], that the points in L are independently and identically distributed uniformly in a thin spherical shell. As a further simplification, it is assumed that the shell is the unit sphere in $\mathbb{R}^n$. Then a pair $(\mathbf{u}, \mathbf{v})$ is reducible if and only if the angle between $\mathbf{u}$ and $\mathbf{v}$ satisfies $\theta(\mathbf{u}, \mathbf{v}) < \pi/3$, where $\theta(\mathbf{u}, \mathbf{v}) = \arccos\left(\langle \mathbf{u}, \mathbf{v} \rangle / (\|\mathbf{u}\| \cdot \|\mathbf{v}\|)\right)$, $\arccos(x) \in [0, \pi]$. Under these assumptions, we require $|L| \approx \sqrt{4/3}^n$ in order to see a neighbour, i.e. reductions. The asymptotically fastest sieve has a heuristic running time and memory cost of $2^{0.292n+o(n)}$ [BDGL16] on a classical computer. Provable sieve style algorithms run in time and memory $2^{n+o(n)}$ [ADRS15]. Lower bounds for (heuristic) sieving-type algorithms, relying in nearest-neighbour search techniques, were provided in [KL21], suggesting current leading constants are optimal and maintaining the complexity of $2^{\Theta(n)}$ time and $2^{\Theta(n)}$ memory.[9]

Between. As mentioned in the introduction, finding an algorithm that achieves single-exponential time with polynomial memory is a long standing open problem in the literature on lattice algorithms, as discussed in e.g. [MV10b, HPS11, ADRS15, LV20]. Moreover, while the time complexities of enumeration-type and sieving-type algorithm seem 'close' in being only a factor of $\log(n)$ away in the exponent, prior attempts to 'interpolate' between the two classes have not yielded algorithms that achieve single-exponential time and polynomial memory complexities,

[9] We note that the same work also provided a lower bound for sieving-type algorithms relying on nearest-neighbour type search for quantum computers with a leading constant of 0.265. In [CL21] a heuristic quantum algorithm with leading constant 0.257 is presented. This does not invalidate the claimed lower bound because that algorithm does not satisfy the conditions for the claimed lower bound. The algorithm in [CL21] still requires $2^{\Theta(n)}$ time and $2^{\Theta(n)}$ memory.

but algorithms that trade time for memory between the two endpoints of sieving and enumeration [BLS16, ACKS21].

BKZ. Instead of calling an algorithm finding (essentially) optimally short vectors in a lattice on the entire lattice, we may repeated call it on (projected) sublattices – called 'blocks' – to achieve worse but still non-trivial approximation factors.

Throughout, we write BKZ-b for the BKZ algorithm with blocksize b. We define the root Hermite factor of a basis $\mathbf{B} \in \mathbb{R}^{n \times r}$ as $(\|\mathbf{b}_1\|/\mathsf{vol}(\Lambda(\mathbf{B})^{1/r}))^{1/(r-1)}$. A common heuristic [Che13] is that for a random q-ary lattice and appropriate b, BKZ-b outputs a basis with root Hermite factor

$$\delta_b = \left(\frac{b}{2\pi e} \cdot (\pi\, b)^{1/b}\right)^{1/(2(b-1))}.$$

Note that $\delta_b \sim b^{1/(2b)}$ and $\log \delta_b \sim (\log b)/(2b)$. Given a sample from $\Lambda_q^\perp(m, n)$ the following proposition expresses the length of the output vector we expect to be output by BKZ-b.

Proposition 3 (BKZ quality). *Let $2n \leq m \in \mathsf{poly}(n)$, $d = \sqrt{n \log q / \log \delta_b}$, q prime and Λ be sampled according to $\Lambda_q^\perp(m, n)$, then BKZ-b heuristically outputs a basis with first vector of norm 2^ν where*

$$\nu = \begin{cases} \frac{m \log b}{2b} + \frac{n}{m} \log q & d \geq m, \\ \min\left(\sqrt{\frac{2 \log b \cdot n \log q}{b}}, \log q\right) & n < d < m, \\ \log q & \text{otherwise.} \end{cases}$$

Proof. We assume the volume of Λ is q^n, this is true with probability at least $1 - 1/q^n$. Heuristically [AD21, Def. 10] BKZ-b will output a basis vector of length $\ell(m) = \delta_b^{m-1} \cdot q^{n/m}$. We may consider a lattice of rank in $\{n, \ldots, m\}$ with volume q^n by putting any basis of Λ into Hermite normal form and removing the appropriate columns [MR09, CL15]. (Equivalently, remove columns of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ provided its row rank is unchanged.) The choice of d in the hypothesis is the minimising zero of $\partial \ell / \partial m$ [MR09]. If it is at least m then we remove no columns and $\nu = \log \ell(m)$. If it is at most n then we stop dropping columns when the lattice becomes $q\mathbb{Z}^n$, and we have the solution $q \cdot \mathbf{e}_1$. Else we compute $\nu = \log \ell = (d-1) \log \delta_b + (n/d) \cdot \log q$. Assuming b and d are large enough that $d - 1 \approx d$ and $\log \delta_b \approx (\log b)/(2b)$ are sensible, the result follows. □

For 'ordinary' parameters, with large enough $q \in \mathsf{poly}(n)$, we have $\nu = \sqrt{\frac{2 \log b \cdot n \log q}{b}}$, i.e. fall into the case where $d < m$ and the heuristic suggests BKZ-b outputs a vector shorter than q. In §3.12 below we give a description of what is heuristically achievable in this setting.

3.12 Finding short vectors in polynomial memory

We recall when BKZ achieves polynomial Hermite approximate factors. In particular, we consider the BKZ algorithm with enumeration or sieving as the SVP oracle, rely on established heuristic analyses for the runtime and output quality over average-case q-ary lattices, and compare to the dimension normalised volume. Given the same computational budget, these analyses suggest shorter vectors are output than their respective proven statements for worst-case lattices. As such they represent a stronger comparison.

Remark 5. The problem of approximately solving SIGS algorithmically is not commonly studied. However, any algorithm solving γ-SIGS solves γ-SIS.

Corollary 10 (BKZ with sieving). *Adopt the notation from Proposition 3. Write BKZ-Sieve$_{\mathsf{poly}(n)}$-b for BKZ-b with lattice-point sieving instantiating the underlying SVP oracle in dimension b and all parameters chosen such that the entire algorithm runs in $\mathsf{poly}(n)$ memory. Then for some $C > 0$, BKZ-Sieve$_{\mathsf{poly}(n)}$-b achieves*

$$\nu \sim \nu_{sieve,\mathsf{poly}(n)} \coloneqq C\sqrt{\log q \cdot n \log\log n / \log n} \notin O(\log m).$$

Proof. Using sieving as the underlying approximate SVP oracle, the BKZ algorithm runs in heuristic time $\mathsf{poly}(m) \cdot 2^{0.292\,b+o(b)}$ [BDGL16] and memory $2^{0.208\,b+o(b)}$.[10] Thus, for any constant $c > 0$ it runs in polynomial memory when $b \leq \log m^c \in O(\log n)$ and BKZ-Sieve$_{\mathsf{poly}(n)}$-b is BKZ-$O(\log m)$ instantiated with sieving. Applying Proposition 3 we obtain the claimed expression. □

Proposition 4 (BKZ with enumeration). *Adopt the notation from Proposition 3 and assume $n < q \in \mathsf{poly}(n)$. Let $f, g : \mathbb{N}^+ \to [1, \infty)$ such that $g \in o(\log n)$, $f \in \omega(g)$ and $f \leq \log q$. Let $m = n \log q / f(n)$. Write BKZ-Enum$_g$-b for BKZ-b with enumeration instantiating the underlying SVP oracle in dimension and all parameters chosen such that the overall running time of the algorithm is bounded by $2^{O(m \cdot g(n))}$. Then BKZ-Enum$_g$-b achieves*

$$\nu \sim \nu_{f,g} \coloneqq \sqrt{\frac{2\,\log b \cdot n \log q}{b}} = \sqrt{2\,\log b \cdot \log n \cdot \frac{f(n)}{g(n)}} \notin O(\log(m)) \ .$$

Proof. Using enumeration as the underlying approximate SVP oracle, the BKZ algorithm runs in heuristic time $\mathsf{poly}(d) \cdot b^{\frac{b}{8}+o(b)}$ [ABF+20] and polynomial memory. Set

$$b = (g(n)/f(n)) \cdot n \log q / \log n,$$

by inspection $b \log b \sim m \cdot g(n)$. Thus, BKZ-Enum$_g$-$b$ is BKZ-$\left(\frac{g(n)}{f(n)} \cdot n \log q / \log n\right)$ instantiated with enumeration. Applying Proposition 3 we obtain the expression. The final non-inclusion follows from $f(n)/g(n) \in \omega(1)$ and $\log b \geq \log n - \log\log n$. □

Thus, BKZ with enumeration can achieve only superpolynomial approximation factors when permitted to run in $2^{O(m \cdot g(n))}$ steps if $m = n \log q / f(n)$ and $f(n) \in \omega(g(n))$.

$m \in \Omega(n \log q)$. We consider the case of single exponential time BKZ-Enum, e.g. $g = 1$, and remove the condition $m \in o(n \log q)$ of Proposition 4. In this case, BKZ-Enum$_g$ can find polynomial length vectors. Note that when $m \in \Omega(n \log q)$ and $g = 1$ we have $f(n) \notin \omega(g)$ and therefore violate the conditions of Proposition 4.

BKZ-Enum$_g$-b runs in single-exponential time when $b \log b \in O(m)$. Let $c > 0$, then a real solution, if one exists, to $b \log b = cm$ is given by $c/W_0(cm)$. Here W_0 is the principal branch of the Lambert-W function and for large x we have $W_0(x) = \ln x - \ln\ln x + o(1)$. Therefore the solution $b \sim cm \log e / (\log m - \log\log m) \sim cm \log e / \log m$. The constant is inessential in the following and we may set $b = m / \log m$. Then $b = m / \log m = n \log q / (\log n + \log\log q) \approx n \log q / \log n$. Here we assume $\log q \in \mathsf{poly}(n)$. Then we obtain

$$\sqrt{\frac{2\,\log b \cdot n \log q}{b}} \sim \sqrt{\frac{\log b \cdot n \log q}{n \log q / \log n}} \sim \sqrt{\log b \cdot \log n} \in O(\log(m)).$$

That is, BKZ with enumeration achieves polynomial approximation in single exponential time *in the dimension of the lattice* m and $\mathsf{poly}(n)$ memory when $m \geq n \log q$. This is a different way of expressing the common knowledge that when solving SIS in this regime the BKZ block size b grows linearly with n rather than m.

[10] A more refined analysis is given in [Duc22].

Table 1. Overview of notation

$H_\infty(X)$	min-entropy of random variable X	3.1
$\Delta(X,Y)$	statistical distance $\frac{1}{2}\sum_{z\in\mathrm{Supp}(X)\cup\mathrm{Supp}(Y)}\lvert\Pr[X=z]-\Pr[Y=z]\rvert$	3.1
(τ,μ,ε)	running in time $\leq\tau(n)$ using memory $\leq\mu(n)$ and succeeding with probability $\geq\varepsilon(n)$	
κ-entropic	conditional min-entropy $\geq\kappa(n)$ conditioning on problem distribution	
$\mathrm{Ber}(p)$	Bernoulli with parameter p, i.e. $\Pr[\mathrm{Ber}(p)=1]=p$ and $\Pr[\mathrm{Ber}(p)=0]=1-p$	
$\mathrm{Bin}(N,p)$	Binomial distribution of N trials each with success probability p	
$R_\alpha(P\Vert Q)$	$\left(\sum_{x\in\mathrm{Supp}(P)}\frac{P(x)^\alpha}{Q(x)^{\alpha-1}}\right)^{\frac{1}{\alpha-1}}$, Rényi divergence of order α	
$R^\varepsilon_\infty(P\Vert Q)$	smooth Rényi divergence of infinite order	
$\lambda_i(\Lambda)$	ith minimum of lattice Λ	3.6
$\mathsf{bl}(\Lambda)$	minimum length of a basis of lattice Λ	
$\mu(\Lambda)$	covering radius of lattice Λ	
$\Lambda_q^\perp(\mathbf{A})$	$\{\mathbf{x}\in\mathbb{Z}^m\colon\mathbf{A}\cdot\mathbf{x}=\mathbf{0}\}$	
$\Lambda_q(\mathbf{A})$	$\{\mathbf{x}\in\mathbb{Z}^m\colon\mathbf{x}\bmod q\in\mathbf{A}^T\cdot\mathbb{Z}_q^n\}$	
$\Lambda_q^\perp(m,n)$	distribution $\Lambda_q^\perp(\mathsf{U}(\mathbb{Z}_q^{n\times m}))$	
$\Lambda_q(m,n)$	distribution $\Lambda_q(\mathsf{U}(\mathbb{Z}_q^{n\times m}))$	
$\varphi_{\mathbf{A}}$	map $\mathbb{Z}^m\to\mathbb{Z}_q^n$, $\mathbf{x}\mapsto\mathbf{A}\mathbf{x}$ for $\mathbf{A}\in\mathbb{Z}_q^{n\times m}$	3.6
$P_{\mathbf{A},\mathbf{t}}$	$\{\mathbf{u}\in\mathbb{Z}^m\colon\mathbf{A}\cdot\mathbf{u}=\mathbf{t}\}$, $\mathbf{t}=\mathbf{0}$ if omitted	
$P^+_{\mathbf{A}}$	$P_{\mathbf{A}}\setminus\{\mathbf{0}\}$	
$P_{\beta,\mathbf{A},\mathbf{t}}$	$P_{\mathbf{A},\mathbf{t}}\cap\mathsf{B}^m(\beta)$, $\mathbf{t}=\mathbf{0}$ if omitted	
$P^+_{\beta,\mathbf{A}}$	$P_{\beta,\mathbf{A}}\setminus\{\mathbf{0}\}$	
$\rho_{s,\mathbf{c}}(\mathbf{x})$	Gaussian function with parameter s and centre $\mathbf{c}$ evaluated at $\mathbf{x}$, $\mathbf{c}=\mathbf{0}$ if omitted	
$D_{\Lambda,s,\mathbf{c}}$	Gaussian distribution with parameter s and centre $\mathbf{c}$ over lattice Λ, $\mathbf{c}=\mathbf{0}$ if omitted	
$D_{\mathbb{Z}^m,s,\mathbf{c}}[\mathbf{A}]$	$D_{\Lambda_q^\perp(\mathbf{A}),s,\mathbf{c}}$, $\mathbf{c}=\mathbf{0}$ if omitted	
$D_{\mathbb{Z}^m,s}[\mathbf{A},\mathbf{t}]$	$\mathbf{u}+D_{\Lambda_q^\perp(\mathbf{A}),s,-\mathbf{u}}$, $\mathbf{u}\in P_{\mathbf{A},\mathbf{t}}$, requires $\mathbf{t}\in\mathrm{im}(\varphi_{\mathbf{A}})$	
$\eta_\delta(\Lambda)$	smoothing parameter of lattice Λ with parameter δ	
SVP	Shortest Vector Problem	
SIVP	Shortest Independent Vectors Problem	
SIS	Short Integer Solution	
ISIS	Inhomogeneous Short Integer Solution	
SIIS	Short Independent Integer Solution	
SIGS	Short Integer Generating Set	
omISIS	One More Inhomogeneous Short Integer Solution	
δ_b	$\left(\frac{b}{2\pi e}\cdot(\pi\,b)^{1/b}\right)^{1/(2(b-1))}$, root Hermite factor achieved by BKZ-b	3.11
$\mathsf{B}^m(r;\mathbf{x})$	ball of radius r centred at $\mathbf{x}\in\mathbb{R}^m$, $\mathbf{x}=\mathbf{0}$ if omitted	
$\mathsf{V}^m(r)$	volume of ball of radius r in $\mathbb{R}^m$	

4 Lower bounding the probability of sampling near-neighbours

We prove Lemma 19, a core lemma that lower bounds the probability of sampling near-neighbours from any distribution over a ball. We use this to reason about the success probability of the core subroutine in the algorithm constructed in Theorem 6.

4.1 An upper bound on the number of not close points in a ball

Let $r \geq 0$, $m \geq 1$ and $\mathbf{x} \in \mathbb{R}^m$. We write $\mathsf{B}^m(r; \mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^m \colon \|\mathbf{y} - \mathbf{x}\| \leq r\}$ for the (closed) ball of radius r around $\mathbf{x}$. If $\mathbf{x} = \mathbf{0}$ we omit it. We adopt the same conventions for $\mathsf{B}^m_\infty(r; \mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^m \colon \|\mathbf{y} - \mathbf{x}\|_\infty \leq r\}$. Denote by $\mathsf{V}^m(r) = \pi^{m/2} \cdot r^m / \Gamma(1 + m/2)$ the volume of $\mathsf{B}^m(r; \mathbf{x})$ for any $\mathbf{x}$. We upper bound the minimum number of points in $\mathsf{B}^m(\beta)$ required to guarantee the existence of two distinct points with distance at most β/γ for $\gamma > 1$.

Lemma 18. *Let $m \in \mathbb{N}^+$, $\beta > 0$ and $\gamma > 1$. For any $S \subset \mathsf{B}^m(\beta)$ with $|S| > 2^{m \log(1+2\gamma)}$, there exist distinct $\mathbf{x}, \mathbf{x}' \in S$ with $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma$.*

Proof. Consider the set

$$\mathcal{X} = \{X \subset \mathsf{B}^m(\beta) \colon \forall\, (\mathbf{x}_i, \mathbf{x}_j) \in X^2,\ \|\mathbf{x}_i - \mathbf{x}_j\| \in \{0\} \cup (\beta/\gamma, \infty)\}$$

consisting of all subsets $X \subset \mathsf{B}^m(\beta)$ such that all pairs of distinct points in X are more than distance β/γ apart. Let $r = \beta/2\gamma$, $X \in \mathcal{X}$ and $Y = \cup_{\mathbf{x} \in X} \mathsf{B}^m(r; \mathbf{x})$. Note $Y \subseteq \mathsf{B}^m(\beta + r)$ and for any distinct $\mathbf{x}, \mathbf{x}' \in X$ we have $\mathsf{B}^m(r; \mathbf{x}) \cap \mathsf{B}^m(r; \mathbf{x}') = \emptyset$. In words, Y consists of non-overlapping balls of radius r around each point in X. Therefore

$$|X| \cdot \mathsf{V}^m(r) = \mathsf{vol}(Y) \leq \mathsf{vol}(\mathsf{B}^m(\beta + r)) = \mathsf{V}^m(\beta + r).$$

Therefore $|X| \leq (1 + \beta/r)^m = (1 + 2\gamma)^m = 2^{m \log(1+2\gamma)}$. □

Let $\gamma = \gamma(m) > 1$. If $\gamma \in O(1)$ then $|S| \in 2^{O(m)}$. If $\log \gamma \in o(\log m)$, e.g. $\gamma \in \mathsf{polylog}\,(m)$, then $|S| \in 2^{o(m \log m)}$ is weakly superexponential. If $\log \gamma \in \Omega(\log m)$, e.g. $\gamma = m^\varepsilon$ with $\varepsilon > 0$, then $|S| \in 2^{\Omega(m \log m)}$ is strongly superexponential.

4.2 Randomisation via Motzkin–Straus

We make use of the Motzkin–Straus Theorem [MS65] to 'upgrade' Lemma 18 to a probabilistic statement (see §3.4 for more details).

Theorem 5 (Motzkin–Straus Theorem [MS65]). *Let $G = (V, E)$ be an undirected simple graph with $|V| = n$ vertices and adjacency matrix $\mathbf{M} \in \{0,1\}^{n \times n}$. Denote by $\Delta = \{\mathbf{x} \in \mathbb{R}^n \colon \sum_i x_i = 1 \wedge \forall i,\, x_i \geq 0\} \subset [0,1]^n$ the n-dimensional simplex. Then*

$$\max_{\mathbf{d} \in \Delta} \mathbf{d}^T \cdot \mathbf{M} \cdot \mathbf{d} = 1 - \frac{1}{\omega(G)}$$

where $\omega(G)$ is the clique number of G.

Lemma 19 (Probabilistic near-neighbours with failures). *Let $\beta > 0$, $\gamma > 1$ and D be a distribution over $\mathsf{B}^m(\beta) \cup \{\perp\}$ with finite support. For independent $\mathbf{x}, \mathbf{x}' \leftarrow D$,*

$$\Pr[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma] \geq \frac{(1 - D(\perp))^2}{2^{m \log(1+2\gamma)}}$$

where we extend $\|\cdot\|$ such that $\|\mathbf{x} - \mathbf{x}'\| = \infty$ if $\mathbf{x} = \perp$ or $\mathbf{x}' = \perp$.

Proof. Let D' be the distribution of D conditioned on not being $\perp$ and sample independent $\mathbf{x}, \mathbf{x}' \leftarrow D$. Let $E_{\not\perp}$ be the event $\mathbf{x} \neq \perp$ and $\mathbf{x}' \neq \perp$. Since $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma$ implies $E_{\not\perp}$ we have $\Pr_{\mathbf{x},\mathbf{x}' \leftarrow D}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma]$

$$\begin{aligned}
&= \Pr_{\mathbf{x},\mathbf{x}' \leftarrow D}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma \wedge E_{\not\perp}] \\
&= \Pr_{\mathbf{x},\mathbf{x}' \leftarrow D}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma \mid E_{\not\perp}] \cdot \Pr_{\mathbf{x},\mathbf{x}' \leftarrow D}[E_{\not\perp}] \\
&= \Pr_{\mathbf{x},\mathbf{x}' \leftarrow D'}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma] \cdot (1 - D(\perp))^2.
\end{aligned}$$

Note $\mathrm{Supp}(D')$ is finite and let $\mathbf{d} \in [0,1]^{|\mathrm{Supp}(D')|}$ represent the probability masses of D'. Define three undirected graphs with vertex set $V = \mathrm{Supp}(D') \subset \mathsf{B}^m(\beta)$:

1. $\hat{G}$: two vertices $\mathbf{y}, \mathbf{z}$ have an edge if and only if $\|\mathbf{y} - \mathbf{z}\| \leq \beta/\gamma$.
2. G: two vertices $\mathbf{y}, \mathbf{z}$ have an edge if and only if $0 < \|\mathbf{y} - \mathbf{z}\| \leq \beta/\gamma$.
3. $\bar{G}$: two vertices $\mathbf{y}, \mathbf{z}$ have an edge if and only if $\|\mathbf{y} - \mathbf{z}\| > \beta/\gamma$.

Let $\hat{\mathbf{M}}$ and $\bar{\mathbf{M}}$ be the adjacency matrices of $\hat{G}$ and $\bar{G}$, respectively. Note that $\hat{G}$ is not a simple graph since each vertex is connected to itself, i.e. the diagonal entries of $\hat{\mathbf{M}}$ are 1. Also note that G and $\bar{G}$ are complementary simple graphs. In particular, the diagonal entries of $\bar{\mathbf{M}}$ are 0. Observe that $\hat{\mathbf{M}} + \bar{\mathbf{M}} = \mathbf{1}$ implies that

$$\mathbf{d}^T \cdot (\hat{\mathbf{M}} + \bar{\mathbf{M}}) \cdot \mathbf{d} = \mathbf{d}^T \cdot \mathbf{1} \cdot \mathbf{d} = \left(\textstyle\sum_{\mathbf{x} \in \mathrm{Supp}(D')} D'(\mathbf{x})\right)^2 = 1$$

where $\mathbf{1}$ denotes the all-one matrix. We therefore have

$$\Pr_{\mathbf{x},\mathbf{x}' \leftarrow D'}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma] = \mathbf{d}^T \cdot \hat{\mathbf{M}} \cdot \mathbf{d} = 1 - \mathbf{d}^T \cdot \bar{\mathbf{M}} \cdot \mathbf{d}.$$

Since $\bar{G}$ is a simple undirected graph and $\mathbf{d}$ is an element of the n-simplex, by the Motzkin–Straus theorem (Theorem 5), we have

$$\mathbf{d}^T \cdot \bar{\mathbf{M}} \cdot \mathbf{d} \leq 1 - \frac{1}{\omega(\bar{G})}.$$

Since $\omega(\bar{G}) = \alpha(G)$, the independence number of G, we obtain

$$\Pr_{\mathbf{x},\mathbf{x}' \leftarrow D'}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/\gamma] \geq \frac{1}{\alpha(G)}.$$

Recall $\alpha(G)$ gives the maximum size of a subset $X \subseteq V$ such that no two distinct vertices of X are connected. In other words, it is the maximum size of a subset $X \subseteq \mathrm{Supp}(D') \subset \mathsf{B}^m(\beta)$ such that $\|\mathbf{x} - \mathbf{x}'\| > \beta/\gamma$ for any distinct $\mathbf{x}, \mathbf{x}' \in X$. Lemma 18 upper bounds the size of any such $X \subset \mathsf{B}^m(\beta)$, therefore $\alpha(G) \leq 2^{m \log(1+2\gamma)}$. □

5 Gaussian vectors generate the lattice

Given enough samples from a discrete Gaussian distribution over some rank n lattice Λ with large enough s, the sampled vectors form a generating set of Λ with high probability. Haviv and Regev [HR14] prove this fact for zero centred Gaussians using $2n^2\log(s\sqrt{n})$ samples and $s \geq \mathsf{bl}(\Lambda)$. We generalise the results in [HR14] in two ways. First, we allow each vector to be sampled from a discrete Gaussian distribution with an arbitrary centre. For example, the centres can all be different. Second, we reduce the number of samples needed by a factor of n to $100n\log(s\sqrt{n})$ while requiring $s \geq \mu(\Lambda) + \mathsf{bl}(\Lambda)$. Recall $\lambda_n(\Lambda)/2 \leq \mu(\Lambda) \leq \sqrt{n}\cdot\lambda_n(\Lambda)/2$ and $\lambda_n(\Lambda) \leq \mathsf{bl}(\Lambda) \leq \sqrt{n}\cdot\lambda_n(\Lambda)/2$. Therefore, if one uses the upper bounds then our required s doubles. In the worst case $\mu(\Lambda)+\mathsf{bl}(\Lambda)$ is a factor $\sqrt{n}$ larger than $\mathsf{bl}(\Lambda)$.

We recall the following useful lemma extracted from [Ban93, Lem. 1.3].

Lemma 20 (Implicit in [Ban93, Lem. 1.3]). *Let $s > 0$, $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$. We have*

$$\begin{aligned}&\exp(-\pi\|\mathbf{a}+\mathbf{b}\|^2/s^2) + \exp(-\pi\|\mathbf{a}-\mathbf{b}\|^2/s^2)\\ \geq\ & 2\cdot\exp(-\pi\|\mathbf{a}\|^2/s^2)\cdot\exp(-\pi\|\mathbf{b}\|^2/s^2).\end{aligned}$$

Proof. Recall that $\cosh(x) = \frac{1}{2}(\exp(x)+\exp(-x)) \geq 1$ for all $x \in \mathbb{R}$. We have

$$\begin{aligned}&\exp(-\pi\|\mathbf{a}+\mathbf{b}\|^2/s^2) + \exp(-\pi\|\mathbf{a}-\mathbf{b}\|^2/s^2)\\ =\ & \exp(-\pi\cdot(\mathbf{a}+\mathbf{b})^T\cdot(\mathbf{a}+\mathbf{b})/s^2) + \exp(-\pi\cdot(\mathbf{a}-\mathbf{b})^T\cdot(\mathbf{a}-\mathbf{b})/s^2)\\ =\ & \exp(-\pi\cdot(\mathbf{a}^T\cdot\mathbf{a}+2\,\mathbf{a}^T\cdot\mathbf{b}+\mathbf{b}^T\cdot\mathbf{b})/s^2) + \exp(-\pi\cdot(\mathbf{a}^T\cdot\mathbf{a}-2\,\mathbf{a}^T\cdot\mathbf{b}+\mathbf{b}^T\cdot\mathbf{b})/s^2)\\ =\ & \exp(-\pi\,\|\mathbf{a}\|^2/s^2)\cdot\exp(-\pi\,\|\mathbf{b}\|^2/s^2)\cdot\left(\exp(2\pi\,\mathbf{a}^T\cdot\mathbf{b}/s^2)+\exp(-2\pi\,\mathbf{a}^T\cdot\mathbf{b}/s^2)\right)\\ \geq\ & 2\cdot\exp(-\pi\,\|\mathbf{a}\|^2/s^2)\cdot\exp(-\pi\,\|\mathbf{b}\|^2/s^2).\end{aligned}$$

□

To begin, we first prove a lower bound for the Gaussian weight $\rho_{s,\mathbf{c}}(\Lambda+\mathbf{w})$ for any centre $\mathbf{c}$, generalising the result in [HR14, Claim 2.1] for $\mathbf{c} = \mathbf{0}$. We also provide a lower bound of $\rho_{s,\mathbf{c}}(\Lambda+\mathbf{w}) + \rho_{s,\mathbf{c}}(\Lambda-\mathbf{w})$ which might be more useful when $\mathbf{w}$ is short, although we do not need to use it in this work. The proof techniques follow from those in the proof of [HR14, Claim 2.1].

Lemma 21 (Generalising [HR14, Claim 2.1]). *For every lattice $\Lambda \subset \mathbb{R}^n$, Gaussian parameter $s > 0$, centre $\mathbf{c} \in \mathbb{R}^n$ and offset $\mathbf{w} \in \mathbb{R}^n$,*

$$\begin{aligned}\rho_{s,\mathbf{c}}(\Lambda+\mathbf{w}) &\geq \rho_{s,\mathbf{c}}(\mathbf{w})\cdot\rho_s(\Lambda) \geq \rho_{s,\mathbf{c}}(\mathbf{w})\cdot\rho_{s,\mathbf{c}}(\Lambda),\\ \rho_{s,\mathbf{c}}(\Lambda+\mathbf{w}) + \rho_{s,\mathbf{c}}(\Lambda-\mathbf{w}) &\geq 2\cdot\rho_s(\mathbf{w})\cdot\rho_{s,\mathbf{c}}(\Lambda).\end{aligned}$$

Proof. Following closely the proof of [HR14, Claim 2.1], we observe

$$\begin{aligned}\rho_{s,\mathbf{c}}(\Lambda+\mathbf{w}) &= \sum_{\mathbf{x}\in\Lambda}\exp(-\pi\|\mathbf{x}+\mathbf{w}-\mathbf{c}\|^2/s^2)\\ &= \frac{1}{2}\sum_{\mathbf{x}\in\Lambda}\left(\exp(-\pi\|\mathbf{x}+(\mathbf{w}-\mathbf{c})\|^2/s^2) + \exp(-\pi\|\mathbf{x}-(\mathbf{w}-\mathbf{c})\|^2/s^2)\right)\\ &\geq \sum_{\mathbf{x}\in\Lambda}\exp(-\pi\|\mathbf{x}\|^2/s^2)\cdot\exp(-\pi\|\mathbf{w}-\mathbf{c}\|^2/s^2)\\ &= \rho_{s,\mathbf{c}}(\mathbf{w})\cdot\rho_s(\Lambda)\end{aligned}$$

where the inequality follows from Lemma 20. The first claim then follows by observing that $\rho_s(\Lambda) \geq \rho_{s,\mathbf{c}}(\Lambda)$ [MR07, Lem. 2.9]. For the second claim, we group the terms differently and observe

$$\begin{aligned}
&\rho_{s,\mathbf{c}}(\Lambda + \mathbf{w}) + \rho_{s,\mathbf{c}}(\Lambda - \mathbf{w}) \\
&= \sum_{\mathbf{x}\in\Lambda} \exp(-\pi\|(\mathbf{x}-\mathbf{c}) + \mathbf{w}\|^2/s^2) + \exp(-\pi\|(\mathbf{x}-\mathbf{c}) - \mathbf{w}\|^2/s^2) \\
&\geq 2 \cdot \sum_{\mathbf{x}\in\Lambda} \exp(-\pi\|\mathbf{x}-\mathbf{c}\|^2/s^2) \cdot \exp(-\pi\|\mathbf{w}\|^2/s^2) \\
&= 2 \cdot \rho_s(\mathbf{w}) \cdot \rho_{s,\mathbf{c}}(\Lambda)
\end{aligned}$$

where the inequality follows from Lemma 20. □

Next, we prove an upper bound on the probability that, for a wide enough discrete Gaussian distribution with an arbitrary centre $\mathbf{c}$, a sample falls into a strict sublattice $\Psi \subsetneq \Lambda$. This again generalises the case of $\mathbf{c} = \mathbf{0}$ proven in [HR14, Lem. 5.1]. As in the proof of [HR14, Lem. 5.1], we prove there exists a lattice point $\mathbf{w} \in \Lambda \setminus \Psi$ which is close to $\mathbf{c}$. Extending the techniques in the proof of [HR14, Lem. 5.1], our first observation is that there must be a lattice point $\mathbf{u}$ in Λ within the covering radius of Λ from $\mathbf{c}$. If the point $\mathbf{u}$ falls outside of Ψ, then we are done. Otherwise, we apply a similar argument to the proof of [HR14, Lem. 5.1] to argue for the existence of a short vector $\mathbf{v} \in \Lambda \setminus \Psi$. Offsetting $\mathbf{u}$ by $\mathbf{v}$ yields a point outside Ψ which is not far from $\mathbf{c}$.

Lemma 22 (Generalising [HR14, Lem. 5.1]). *For any $\alpha > 0$, rank n lattice $\Lambda \subseteq \mathbb{R}^n$, strict sublattice $\Psi \subsetneq \Lambda$, Gaussian parameter $s \geq \alpha \cdot \sqrt{n} \cdot \lambda_n(\Lambda)$, and centre $\mathbf{c} \in \mathbb{R}^n$,*

$$\Pr[D_{\Lambda,s,\mathbf{c}} \in \Psi] \leq \frac{1}{1 + \exp(-\pi/\alpha^2)}.$$

Proof. Following [HR14, Lem. 5.1] we make some observations. For any $\mathbf{w} \in \Lambda \setminus \Psi$, which exists because Ψ is a strict sublattice of Λ, the cosets $\Psi + \mathbf{w}$ and Ψ are disjoint and both contained in Λ. Observe that

$$\rho_{s,\mathbf{c}}(\Lambda) \geq \rho_{s,\mathbf{c}}(\Psi) + \rho_{s,\mathbf{c}}(\Psi + \mathbf{w}) \geq (1 + \rho_{s,\mathbf{c}}(\mathbf{w})) \cdot \rho_{s,\mathbf{c}}(\Psi),$$

where the first inequality is due to $\Psi + \mathbf{w}$ and Ψ being disjoint, and the second inequality is due to the first bound in Lemma 21. Consequently,

$$\Pr[D_{\Lambda,s,\mathbf{c}} \in \Psi] = \frac{\rho_{s,\mathbf{c}}(\Psi)}{\rho_{s,\mathbf{c}}(\Lambda)} \leq \frac{1}{1 + \rho_{s,\mathbf{c}}(\mathbf{w})} = \frac{1}{1 + \exp(-\pi\|\mathbf{w}-\mathbf{c}\|^2/s^2)}.$$

By the definition of the covering radius μ there exists $\mathbf{u} \in \Lambda$ such that

$$\|\mathbf{u} - \mathbf{c}\| \leq \mu(\Lambda) \leq \frac{\sqrt{n}}{2}\lambda_n(\Lambda).$$

If $\mathbf{u} \notin \Psi$ then set $\mathbf{w} = \mathbf{u}$ and conclude. Next, we tackle the case where $\mathbf{u} \in \Psi$. Since Ψ is a strict sublattice of Λ there exists a vector $\mathbf{v} \in \Lambda \setminus \Psi$ with $\|\mathbf{v}\| \leq \mathsf{bl}(\Lambda) \leq \frac{\sqrt{n}}{2}\lambda_n(\Lambda)$. Setting $\mathbf{w} = \mathbf{u} + \mathbf{v} \in \Lambda \setminus \Psi$, we have

$$\|\mathbf{w} - \mathbf{c}\| = \|\mathbf{u} + \mathbf{v} - \mathbf{c}\| \leq \|\mathbf{u} - \mathbf{c}\| + \|\mathbf{v}\| \leq \frac{\sqrt{n}}{2}\lambda_n(\Lambda) + \frac{\sqrt{n}}{2}\lambda_n(\Lambda) = \sqrt{n} \cdot \lambda_n(\Lambda)$$

and therefore $\exp(-\pi\|\mathbf{w}-\mathbf{c}\|^2/s^2) \geq \exp(-\pi/\alpha^2)$.

The above suffices to complete the proof of the claim. We remark that a tighter bound exists when $\mathbf{u} \notin \mathrm{span}(\Psi)$. In this case $\mathrm{span}(\Psi) \subsetneq \mathrm{span}(\Lambda)$ which implies that there exists $\mathbf{v} \in \Lambda \setminus \Psi$ with $\|\mathbf{v}\| \leq \lambda_n(\Lambda)$. Setting $\mathbf{w} = \mathbf{u} + \mathbf{v} \in \Lambda \setminus \Psi$ we have

$$\begin{aligned}\|\mathbf{w} - \mathbf{c}\| = \|\mathbf{u} + \mathbf{v} - \mathbf{c}\| \leq \|\mathbf{u} - \mathbf{c}\| + \|\mathbf{v}\| &\leq \frac{\sqrt{n}}{2} \cdot \lambda_n(\Lambda) + \lambda_n(\Lambda) \\ &\leq \left(\frac{\sqrt{n}}{2} + 1\right) \cdot \lambda_n(\Lambda). \qquad \square\end{aligned}$$

In the above we note that neither s or the upper bound depend on the strict sublattice Ψ and that the upper bound converges to one half as s grows.

Given the above tools we can argue that $100n\log(s\sqrt{n})$ Gaussian samples with large enough s and arbitrary centres suffice to form a generating set of Λ. As in [HR14] we proceed in two steps. First we argue that $k = 50n$ samples suffice to form a set containing n linearly independent vectors. This improves upon n^2 [HR14, Cor. 5.1]. Second, we show that $\ell = 50n\log(s\sqrt{n})$ samples suffice to improve such a rank n set to a generating set. This improves upon $n^2\log(s\sqrt{n})$ [HR14, Lem. 5.2]. In both cases the improvement comes from the use of tighter tail bounds. We note that the following Corollary 11 is also proven in [DvW22, Lem. 3.4].

Corollary 11 (Generalising and tightening of [HR14, Cor. 5.1]). *For any rank n lattice Λ, sequence of k centres $\mathbf{c}_1, \ldots, \mathbf{c}_k \in \mathbb{R}^n$ and $s \geq \sqrt{n} \cdot \lambda_n(\Lambda)$ where $k \geq 2\,(1+e^{\pi}) \cdot n$, the probability that $(\mathbf{x}_1, \ldots, \mathbf{x}_k)$ with $\mathbf{x}_i \leftarrow D_{\Lambda,s,\mathbf{c}_i}$ contains no subset of n linearly independent vectors is $2^{-\Omega(n)}$.*

Proof. For $i \in [k]$ let $\Psi_i = \Lambda \cap \mathrm{span}(\mathbf{x}_1, \ldots, \mathbf{x}_i)$ and $n_i = \mathrm{rank}(\Psi_i)$. Clearly $n \geq n_i \geq n_{i-1} \geq 0$ for $i > 1$. Note that if $\mathrm{span}(\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}) \subsetneq \mathrm{span}(\Lambda)$ then Ψ_{i-1} is a strict sublattice of Λ. In this case, applying Lemma 22 with $\alpha = 1$, we have $\mathbf{x}_i \notin \mathrm{span}(\mathbf{x}_1, \ldots, \mathbf{x}_{i-1})$ with probability at least $\epsilon = 1 - \frac{1}{1+e^{-\pi}} = \frac{1}{1+e^{\pi}}$. In other words, conditioned on $n_{i-1} < n$, we have $n_i = n_{i-1} + 1$ with probability at least ϵ, independently of i. We show that $n_k = n$ except with probability $2^{-\Omega(n)}$.

To simplify the analysis, we define a sequence of i.i.d. Bernoulli random variables $G_1, \ldots, G_k$ such that $\Pr[G_i = 1] = \epsilon$ and $\Pr[G_i = 0] = 1 - \epsilon$ for all i. The events $G_i = 0$ and $G_i = 1$ correspond to the events $n_i = n_{i-1} < n$ and $n_i = \min\{n_{i-1} + 1, n\}$ respectively. If $\sum_{i=1}^{k} G_i \geq n$ except with probability $2^{-\Omega(n)}$ then $n_k = n$ except with probability $2^{-\Omega(n)}$. By Lemma 2

$$\Pr\left[\sum_{i=1}^{k} G_i \leq (1-\delta)k\epsilon\right] \leq e^{-k\epsilon\delta^2/2}.$$

Setting $\delta = 1 - \frac{n}{k\epsilon} \geq \frac{1}{2}$ so that $(1-\delta)k\epsilon = n$, we have

$$\Pr\left[\sum_{i=1}^{k} G_i < n\right] \leq \Pr\left[\sum_{i=1}^{k} G_i \leq n\right] \leq e^{-k\epsilon/8} \leq e^{-n/4} \leq 2^{-\Omega(n)}. \qquad \square$$

Lemma 23 (Generalising and tightening [HR14, Lem. 5.2]). *For any $\beta > 0$, rank n lattice Λ with $\det(\Lambda) \geq 1$, sequence of $r = k + \ell$ centres $\mathbf{c}_1, \ldots, \mathbf{c}_r \in \mathbb{R}^n$ with $\max_{i\in[r]}\|\mathbf{c}_i\| \leq \beta$ and $s \geq \max\{\sqrt{n} \cdot \lambda_n(\Lambda), \eta_\delta(\Lambda)\}$ for some $\delta \leq 2^{-\Omega(n)}$, where $k = 2\,(1+e^{\pi}) \cdot n$, $\ell \geq 2\,(1+e^{\pi}) \cdot n \cdot \log(\beta + s\sqrt{n})$, the probability that $(\mathbf{x}_1, \ldots, \mathbf{x}_r)$ with $\mathbf{x}_i \leftarrow D_{\Lambda,s,\mathbf{c}_i}$ does not generate Λ is $2^{-\Omega(n)}$. If $\mathbf{c}_1 = \ldots = \mathbf{c}_r = \mathbf{0}$, then the requirement $s \geq \eta_\delta(\Lambda)$ can be dropped.*

Proof. For $0 \leq i \leq \ell$ let $\Psi_i \subseteq \Lambda$ be the sublattice generated by the first $k+i$ samples, i.e. $(\mathbf{x}_1, \ldots, \mathbf{x}_k, \ldots, \mathbf{x}_{k+i})$. Via Lemma 11 (or Corollary 3 for the zero-centred case) and a union bound, except with probability $2^{-\Omega(n)}$, the first k vectors have norm at most $\beta + s\sqrt{n}$. Furthermore, by Corollary 11

and except with probability $2^{-\Omega(n)}$, the first k vectors contain n linearly independent vectors. We now assume these two events occur. Thus, by Hadamard's inequality and the assumption that $\det(\Lambda) \geq 1$, the index of Ψ_0 in Λ satisfies

$$|\Lambda : \Psi_0| = \frac{\det(\Psi_0)}{\det(\Lambda)} \leq (\beta + s\sqrt{n})^n.$$

Let $h_i = \log |\Lambda : \Psi_i| \geq 0$ denote the logarithm of the index of Ψ_i in Λ. We have $0 \leq h_i \leq h_{i-1}$ for $i \in [\ell]$. Applying Lemma 22 with $\alpha = 1$, conditioned on $h_{i-1} > 0$ we have $0 \leq h_i \leq h_{i-1} - 1$ with probability at least $\epsilon = 1 - \frac{1}{1+e^{-\pi}} = \frac{1}{1+e^{\pi}}$. We show that $h_\ell = 0$ except with probability $2^{-\Omega(n)}$.

To simplify the analysis we define a sequence of i.i.d. Bernoulli random variables $G_1, \ldots, G_\ell$ such that $\Pr[G_i = 1] = \epsilon$ and $\Pr[G_i = 0] = 1 - \epsilon$ for all i. The events $G_i = 0$ and $G_i = 1$ correspond to the events $h_i = h_{i-1} > 0$ and $h_i \leq \max\{h_{i-1} - 1, 0\}$ respectively. If $\sum_{i=1}^{\ell} G_i \geq h_0$ except with probability $2^{-\Omega(n)}$ then $h_\ell = 0$ except with probability $2^{-\Omega(n)}$. By Lemma 2

$$\Pr\left[\sum_{i=1}^{\ell} G_i \leq (1-\delta) \cdot \ell \cdot \epsilon\right] \leq e^{-\ell \cdot \epsilon \cdot \delta^2/2}.$$

Setting $\delta = 1 - \frac{h_0}{\ell \cdot \epsilon} \geq \frac{1}{2}$ so that $(1-\delta) \cdot \ell \cdot \epsilon = h_0$, we have

$$\Pr\left[\sum_{i=1}^{\ell} G_i < h_0\right] \leq \Pr\left[\sum_{i=1}^{\ell} G_i \leq h_0\right] \leq e^{-\ell \cdot \epsilon/8} \leq e^{-n \log(\beta + s\sqrt{n})/4} \leq 2^{-\Omega(n)}. \qquad \square$$

Remark 6. By Corollary 9 we have a probabilistic upper bound on the last minima of $\Lambda_q^\perp(m, n)$. We can therefore choose $s \geq \sqrt{m} \cdot \lambda_m(\Lambda)$ according to this upper bound to ensure that with overwhelming probability the $k + \ell$ samples of Lemma 23 will be sufficient.

6 Rejection sampling

This section is dedicated to establishing Corollary 12 which allows us to use rejection sampling on samples from $D_{\Lambda_q^\perp(\mathbf{A}),s,\mathbf{c}_i}$ to obtain samples from $D_{\Lambda_q^\perp(\mathbf{A}),s}$ (zero-centred). We are interested in a regime where $\mathbf{c}_i$ changes for each sample and where $\|\mathbf{c}_i\| \leq s\sqrt{2m}$. Below, we establish that there are choices of parameters where rejection sampling runs in single-exponential time in m and polynomial memory.

The next lemma bounds the Rényi divergence of finite order between two Gaussian distributions with zero and non-zero centres, respectively. It generalises [DFPS22, Lemma A.15], which is stated for $\mathbb{Z}^m$, to general lattices Λ.

Lemma 24 (Generalisation of [DFPS22, Lemma A.15]). *Let $\Lambda \subset \mathbb{R}^m$ be a lattice, $s \geq \eta_\delta(\Lambda)$, $\alpha \in (1,\infty)$ and $\mathbf{c} \in \mathbb{R}^m$. It holds that*

$$R_\alpha(D_{\Lambda,s}\|D_{\Lambda,s,\mathbf{c}}) \in \left[\left(\frac{1-\delta}{1+\delta}\right),1\right] \cdot \left[\left(\frac{1-\delta}{1+\delta}\right)^{1/(\alpha-1)},1\right] \cdot \exp\left(\frac{\alpha\pi}{s^2}\cdot\|\mathbf{c}\|^2\right).$$

If $\mathbf{c} \in \Lambda$ we may set the first interval to $\{1\}$. If $(\alpha-1)\cdot\mathbf{c} \in \Lambda$ we may set the second interval to $\{1\}$.

Proof. First, we observe the identity

$$\alpha\cdot\|\mathbf{x}\|^2 - (\alpha-1)\cdot\|\mathbf{x}-\mathbf{c}\|^2 = \|\mathbf{x}+(\alpha-1)\cdot\mathbf{c}\|^2 - \alpha\cdot(\alpha-1)\cdot\|\mathbf{c}\|^2.$$

First, assume $\mathbf{c} \in \Lambda$. Then, we show by direct calculation that $R_\alpha(D_{\Lambda,s}\|D_{\Lambda,s,\mathbf{c}})$

$$\begin{aligned}
&= \left(\sum_{\mathbf{x}\in\Lambda}\frac{D_{\Lambda,s}(\mathbf{x})^\alpha}{D_{\Lambda,s,\mathbf{c}}(\mathbf{x})^{\alpha-1}}\right)^{\frac{1}{\alpha-1}} \\
&= \left(\sum_{\mathbf{x}\in\Lambda}\frac{\left(\frac{\rho_s(\mathbf{x})}{\rho_s(\Lambda)}\right)^\alpha}{\left(\frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(\Lambda)}\right)^{\alpha-1}}\right)^{\frac{1}{\alpha-1}} = \left(\sum_{\mathbf{x}\in\Lambda}\frac{1}{\rho_s(\Lambda)}\cdot\frac{\rho_s(\mathbf{x})^\alpha}{\rho_{s,\mathbf{c}}(\mathbf{x})^{\alpha-1}}\right)^{\frac{1}{\alpha-1}} \qquad (\star) \\
&= \left(\frac{\sum_{\mathbf{x}\in\Lambda}\exp\left(-\frac{\pi}{s^2}(\alpha\cdot\|\mathbf{x}\|^2-(\alpha-1)\cdot\|\mathbf{x}-\mathbf{c}\|^2)\right)}{\rho_s(\Lambda)}\right)^{\frac{1}{\alpha-1}} \\
&= \left(\frac{\sum_{\mathbf{x}\in\Lambda}\exp\left(-\frac{\pi}{s^2}(\|\mathbf{x}+(\alpha-1)\cdot\mathbf{c}\|^2-\alpha(\alpha-1)\cdot\|\mathbf{c}\|^2)\right)}{\rho_s(\Lambda)}\right)^{\frac{1}{\alpha-1}} \\
&= \left(\frac{\sum_{\mathbf{x}\in\Lambda}\exp\left(-\frac{\pi}{s^2}(\|\mathbf{x}+(\alpha-1)\cdot\mathbf{c}\|^2)\right)\cdot\exp\left(\frac{\pi}{s^2}\cdot\alpha\cdot(\alpha-1)\cdot\|\mathbf{c}\|^2\right)}{\rho_s(\Lambda)}\right)^{\frac{1}{\alpha-1}} \\
&= \exp\left(\frac{\alpha\pi}{s^2}\cdot\|\mathbf{c}\|^2\right)\cdot\left(\frac{\sum_{\mathbf{x}\in\Lambda}\exp\left(-\frac{\pi}{s^2}(\|\mathbf{x}+(\alpha-1)\cdot\mathbf{c}\|^2)\right)}{\rho_s(\Lambda)}\right)^{\frac{1}{\alpha-1}} \\
&= \exp\left(\frac{\alpha\pi}{s^2}\cdot\|\mathbf{c}\|^2\right)\cdot\left(\frac{\rho_{s,-(\alpha-1)\cdot\mathbf{c}}(\Lambda)}{\rho_s(\Lambda)}\right)^{\frac{1}{\alpha-1}}.
\end{aligned}$$

Let $\mathbf{c}' = -(\alpha-1)\cdot\mathbf{c}$. If $\mathbf{c}' \in \Lambda$ then $\rho_{s,\mathbf{c}'}(\Lambda) = \rho_s(\Lambda)$. Else by Lemma 13 we have $\rho_{s,\mathbf{c}'}(\Lambda)/\rho_s(\Lambda) \in [(1-\delta)/(1+\delta), 1]$. Now, consider $\mathbf{c} \notin \Lambda$. Then, by Lemma 13 the equality $(\star)$ becomes

$$\begin{aligned}
\left(\sum_{\mathbf{x}\in\Lambda} \frac{\left(\frac{\rho_s(\mathbf{x})}{\rho_s(\Lambda)}\right)^{\alpha}}{\left(\frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(\Lambda)}\right)^{\alpha-1}}\right)^{\frac{1}{\alpha-1}} &= \left(\frac{1}{\rho_s(\Lambda)}\cdot\left(\frac{\rho_{s,\mathbf{c}}(\Lambda)}{\rho_s(\Lambda)}\right)^{\alpha-1}\cdot\sum_{\mathbf{x}\in\Lambda}\frac{\rho_s(\mathbf{x})^{\alpha}}{\rho_{s,\mathbf{c}}(\mathbf{x})^{\alpha-1}}\right)^{\frac{1}{\alpha-1}} \\
&\in \left(\frac{1}{\rho_s(\Lambda)}\cdot\left[\left(\frac{1-\delta}{1+\delta}\right),1\right]^{\alpha-1}\cdot\sum_{\mathbf{x}\in\Lambda}\frac{\rho_s(\mathbf{x})^{\alpha}}{\rho_{s,\mathbf{c}}(\mathbf{x})^{\alpha-1}}\right)^{\frac{1}{\alpha-1}} \\
&= \left[\left(\frac{1-\delta}{1+\delta}\right),1\right]\cdot\left(\frac{1}{\rho_s(\Lambda)}\cdot\sum_{\mathbf{x}\in\Lambda}\frac{\rho_s(\mathbf{x})^{\alpha}}{\rho_{s,\mathbf{c}}(\mathbf{x})^{\alpha-1}}\right)^{\frac{1}{\alpha-1}}
\end{aligned}$$

The proof then proceeds as before. □

Based on the above, the next lemma bounds the smooth Rényi divergence of infinite order between two Gaussian distributions with zero and non-zero centres, respectively. It generalises [DFPS22, Lemma C.2], which proved it for $\mathbb{Z}^m$, to general lattices Λ.

Lemma 25 (Generalisation of [DFPS22, Lemma C.2]). *Let $\Lambda \subset \mathbb{R}^m$ be a lattice, $\varepsilon \in (0,1)$, $s > 0$ and $\mathbf{c} \in \mathbb{R}^m$. It holds that*

$$R_\infty^\varepsilon(D_{\Lambda,s}\|D_{\Lambda,s,\mathbf{c}}) \le \exp\left(\pi\cdot\frac{\|\mathbf{c}\|^2}{s^2} + 2\sqrt{\pi\ln\frac{1}{\varepsilon}}\cdot\frac{\|\mathbf{c}\|}{s}\right).$$

Proof. We follow the proof strategy of [DFPS22, Lemma C.2]. First note if $\mathbf{c} = \mathbf{0}$ the bound is true, so assume $\mathbf{c} \neq \mathbf{0}$. Combining Lemmas 5 and 24, we have

$$R_\infty^\varepsilon(D_{\Lambda,s}\|D_{\Lambda,s,\mathbf{c}}) \le \frac{\exp\left(\frac{\alpha\pi}{s^2}\cdot\|\mathbf{c}\|^2\right)}{\varepsilon^{1/(\alpha-1)}} = \exp\left(\frac{\alpha\pi}{s^2}\cdot\|\mathbf{c}\|^2 + \frac{1}{\alpha-1}\ln\frac{1}{\varepsilon}\right).$$

Setting $\alpha = 1 + \frac{s}{\|\mathbf{c}\|}\sqrt{\frac{\ln(1/\varepsilon)}{\pi}}$ yields the claimed bound. □

We state a variant of [DFPS22, Lemma 2.2]. First, our variant considers running the rejection sampling procedure T times rather than potentially indefinitely. This allows us to bound the statistical distance of the output to the ideal target distribution with a simple hybrid argument.

Lemma 26 (Truncation and generalisation of [DFPS22, Lemma 2.2]). *Let $T \in \mathbb{N}$, $M \ge 1$ and $\varepsilon \in [0, 1/2]$ be such that $R_\infty^\varepsilon(P\|Q_i) \le M$ for all $i \in [T]$. Consider the algorithms in Fig. 4. It holds that*

$$\forall i \in [T],\ 1 - \frac{1}{M} \le \Pr\left[\mathsf{RejSamp}_{P,Q_i,M}(1^m) = \perp\right] \le 1 - \frac{1-\varepsilon}{M}.$$

Moreover,

$$\begin{gathered}
\forall i \in [T],\ \Delta(\mathsf{RejSamp}_{P,Q_i,M}(1^m),\ \ \mathsf{RejSamp}^*_{P,M}(1^m)) \le \varepsilon/M \quad \text{and} \\
\Delta(\mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}(1^m),\ \mathsf{RejSampLoop}^*_{P,M,T}(1^m)) \le T\cdot\varepsilon/M.
\end{gathered}$$

$\underline{\mathsf{RejSamp}_{P,Q,M}(1^m)}$

$x \leftarrow Q$

with probability $\min\left\{1, \frac{P(x)}{M \cdot Q(x)}\right\}$:

  return x

return $\perp$

$\underline{\mathsf{RejSamp}^*_{P,M}(1^m)}$

$x \leftarrow P$

with probability $\frac{1}{M}$:

  return x

return $\perp$

$\underline{\mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}(1^m)}$

for $i \in [T]$ **do**

  $z \leftarrow \mathsf{RejSamp}_{P,Q_i,M}(1^m)$

  if $z \neq \perp$ **then return** z

return $\perp$

$\underline{\mathsf{RejSampLoop}^*_{P,M,T}(1^m)}$

for $i \in [T]$ **do**

  $z \leftarrow \mathsf{RejSamp}^*_{P,M}(1^m)$

  if $z \neq \perp$ **then return** z

return $\perp$

Figure 4. Rejection sampling algorithms where P and Q are probability distributions, $M \geq 1$ and $T \in \mathbb{N}$.

Proof. We follow the proof strategy of [DFPS22, Lemma 2.2]. Adapting their notation, we introduce the following shorthands for the probability mass functions:

$$\mathcal{A}^{\mathsf{real}}_{Q_i}(x) \coloneqq \mathsf{RejSamp}_{P,Q_i,M}(1^m)(x), \qquad \mathcal{A}^{\mathsf{ideal}}(x) \coloneqq \mathsf{RejSamp}^*_{P,M}(1^m)(x),$$
$$\mathcal{B}^{\mathsf{real}}_{(Q_i)_{i=1}^T}(x) \coloneqq \mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}(1^m)(x), \qquad \mathcal{B}^{\mathsf{ideal}}(x) \coloneqq \mathsf{RejSampLoop}^*_{P,M,T}(1^m)(x).$$

For each $i \in [T]$, define the normalisation constant

$$C_i \coloneqq \sum_{x \in \mathrm{Supp}(Q_i)} \min\{P(x), M \cdot Q_i(x)\}.$$

Note that $C_i \leq \sum_{x \in \mathrm{Supp}(Q_i)} P(x) \leq 1$. Let $X_i = \{x \in \mathrm{Supp}(P) \colon P(x) \leq MQ_i(x)\}$ then since $R^{\varepsilon}_{\infty}(P \| Q_i) \leq M$, by definition $P(X_i) \geq 1 - \varepsilon$ and note $X_i \subseteq \mathrm{Supp}(Q_i)$. Therefore $C_i \geq P(X_i) \geq 1 - \varepsilon$. Define the probability mass function P_i with support equal to $\mathrm{Supp}(Q_i)$ as

$$P_i(x) \coloneqq \min\{P(x), M \cdot Q_i(x)\}/C_i.$$

We have that $\Pr_{x \leftarrow P_i}[P_i(x) \leq (M/C_i) \cdot Q_i(x)] = 1$, i.e. $R_{\infty}(P_i \| Q_i) \leq M/C_i$. Therefore, $\mathcal{A}^{\mathsf{real}}_{Q_i}$ is a perfect rejection sampling algorithm with source Q_i and target P_i. $\mathcal{A}^{\mathsf{real}}_{Q_i}$ can be expressed as

$$\mathcal{A}^{\mathsf{real}}_{Q_i}(x) = \begin{cases} 1 - \frac{C_i}{M} \in \left[1 - \frac{1}{M}, 1 - \frac{1-\varepsilon}{M}\right] & x = \perp \\ \frac{C_i}{M} \cdot P_i(x) = \min\left\{\frac{P(x)}{M}, Q_i(x)\right\} & x \neq \perp. \end{cases}$$

This proves the claim about $\Pr\left[\mathsf{RejSamp}_{P,Q_i,M}(1^m) = \perp\right]$.

We next bound the statistical distances. First, observe that for any $i \in [T]$

$$\begin{aligned}\Delta(\mathcal{A}_{Q_i}^{\mathsf{real}}, \mathcal{A}^{\mathsf{ideal}}) &= \frac{1}{2}\sum_{x\neq\perp}\left|\min\left\{\frac{P(x)}{M}, Q_i(x)\right\} - \frac{P(x)}{M}\right| + \frac{1}{2}\left|\left(1-\frac{C_i}{M}\right) - \left(1-\frac{1}{M}\right)\right| \\ &\leq \frac{1}{2}\sum_{x\neq\perp}\left|\max\left\{0, \frac{P(x)}{M} - Q_i(x)\right\}\right| + \frac{\varepsilon}{2M} \\ &\leq \frac{1}{2}\sum_{x:\frac{P(x)}{M}>Q_i(x)}\left(\frac{P(x)}{M} - Q_i(x)\right) + \frac{\varepsilon}{2M} \\ &\leq \frac{\varepsilon}{2M} + \frac{\varepsilon}{2M} = \frac{\varepsilon}{M}.\end{aligned}$$

Finally, consider $\Delta(\mathcal{B}^{\mathsf{real}}_{(Q_i)_{i=1}^T}, \mathcal{B}^{\mathsf{ideal}})$. Write $\mathbf{v}_{\mathsf{real}}$ and $\mathbf{v}_{\mathsf{ideal}}$ as the vectors $(\mathcal{A}^{\mathsf{real}}_{Q_1}, \ldots, \mathcal{A}^{\mathsf{real}}_{Q_T})$ and $(\mathcal{A}^{\mathsf{ideal}}, \ldots, \mathcal{A}^{\mathsf{ideal}})$. By a hybrid argument $\Delta(\mathbf{v}_{\mathsf{real}}, \mathbf{v}_{\mathsf{ideal}}) \leq T \cdot \varepsilon/M$. Finally, define f which outputs the first entry of $\mathbf{v}$ not equal to $\perp$, or $\perp$ if there is no such element. By the data processing inequality $\Delta(f(\mathbf{v}_{\mathsf{real}}), f(\mathbf{v}_{\mathsf{ideal}})) \leq \Delta(\mathbf{v}_{\mathsf{real}}, \mathbf{v}_{\mathsf{ideal}}) \leq T \cdot \varepsilon/M$. Noting that $\mathcal{B}^{\mathsf{real}}_{(Q_i)_{i=1}^T} = f(\mathbf{v}_{\mathsf{real}})$ and $\mathcal{B}^{\mathsf{ideal}} = f(\mathbf{v}_{\mathsf{ideal}})$ (as random variables) finishes the proof. □

We are now ready to state our main rejection sampling result.

Corollary 12 (Exponential time rejection sampling). *Let $m \in \mathbb{N}$, $\varepsilon \in (0, 1/2]$, $\Lambda \subseteq \mathbb{R}^m$ be a lattice,*

$$M \geq \exp\left(2\pi m + 2\sqrt{2\pi m \ln\frac{1}{\varepsilon}}\right) \geq 1, \qquad T \geq \frac{M}{\sqrt{\varepsilon}},$$

$s > 0$ and $\mathbf{c}_1, \ldots, \mathbf{c}_T \in \mathbb{R}^m$ with $\|\mathbf{c}_i\| \leq s\sqrt{2m}$ for all $i \in [T]$. Let $P = D_{\Lambda,s}$ and let $Q_i = D_{\Lambda,s,\mathbf{c}_i}$ for $i \in [T]$. Then

$$\Pr\left[\mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}(1^m) = \perp\right] \leq \sqrt{\varepsilon} \quad \text{and} \tag{1}$$

$$\Delta(\mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}(1^m), \mathsf{RejSampLoop}^*_{P,M,T}(1^m)) \leq T \cdot \varepsilon/M. \tag{2}$$

Moreover, $\mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}(1^m)$ has additive $\mathsf{poly}(m)$ time and memory overhead compared to sampling from Q_i sequentially T times.

Proof. By Lemma 25 and hypothesis on $m, \mathbf{c}, \varepsilon, s$ and M, $R^{\varepsilon}_{\infty}(D_{\Lambda,s}\|Q_i) \leq M$. For Eq. (1), we observe by the first part of Lemma 26 that

$$\begin{aligned}&\forall i \in [T],\ \Pr\left[\mathsf{RejSamp}_{P,Q_i,M}(1^m) = \perp\right] \leq 1 - \frac{1-\varepsilon}{M} \text{ therefore,} \\ &\Pr\left[\mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}(1^m) = \perp\right] \leq \left(1 - \frac{1-\varepsilon}{M}\right)^T \leq \frac{1}{1 + \frac{1-\varepsilon}{M}\cdot T} \leq \sqrt{\varepsilon}\end{aligned}$$

To bound $(1-(1-\varepsilon)/M)^T$ we use $(1+x)^r \leq 1/(1-rx)$ for $x \in [-1, 1/r)$ and $r \geq 0$ (set $r = T > 0$ and $x = -(1-\varepsilon)/M \in (-1, 0)$). Then $1 + \frac{1-\varepsilon}{M}\cdot T \geq 1 + \frac{1-\varepsilon}{\sqrt{\varepsilon}} \geq \frac{1}{\sqrt{\varepsilon}}$ as $\varepsilon \leq 1$. Eq. (2) directly follows from the second part of Lemma 26.

The computational overhead in $\mathsf{RejSamp}_{P,Q_i,M}(1^m)$ is to compute $\min\left\{1, \frac{P(x)}{M\cdot Q_i(x)}\right\}$, which reduces to computing $\exp(x)$. This can be accomplished in finite precision using $\mathsf{poly}(m)$ bits running in $\mathsf{poly}(m)$ time and memory. □

Finally, we capture the subtlety that in our proofs we receive samples distributed close to, rather than exactly as, $D_{\Lambda,s,\mathbf{c}}$ due to our use of convolution lemmas. In particular, we perform rejection sampling with source distribution Q'_i, an unknown distribution for which $\Delta(Q_i, Q'_i)$ is upper bounded. We cannot compute $\min\{1, P(x)/MQ'_i(x)\}$ as Q'_i is unknown. Instead we consider the algorithms of Fig. 5. In particular the lefthand algorithms $\mathsf{RejSamp}'_{P,Q,M}$ and $\mathsf{RejSampLoop}'_{P,(Q_i)_{i=1}^T,M}$ are distinct from Fig. 4. Crucially, $\mathsf{RejSamp}'_{P,Q,M}$ samples from Q' but still rejects with respect to the computable value $\min\{1, P(x)/MQ(x)\}$.

$\mathsf{RejSamp}'_{P,Q,M}(1^m)$

$x \leftarrow Q'$

with probability $\min\left\{1, \frac{P(x)}{M \cdot Q(x)}\right\}$:

  return x

return $\perp$

$\mathsf{RejSamp}^*_{P,M}(1^m)$

$x \leftarrow P$

with probability $\frac{1}{M}$:

  return x

return $\perp$

$\mathsf{RejSampLoop}'_{P,(Q_i)_{i=1}^T,M}(1^m)$

for $i \in [T]$ **do**

  $z \leftarrow \mathsf{RejSamp}'_{P,Q_i,M}(1^m)$

  if $z \neq \perp$ **then return** z

return $\perp$

$\mathsf{RejSampLoop}^*_{P,M,T}(1^m)$

for $i \in [T]$ **do**

  $z \leftarrow \mathsf{RejSamp}^*_{P,M}(1^m)$

  if $z \neq \perp$ **then return** z

return $\perp$

Figure 5. Rejection sampling algorithms where P and Q are probability distributions, Q' is a probability distribution 'close' to Q, $M \geq 1$ and $T \in \mathbb{N}$.

Corollary 13 (Exponential time rejection sampling, close distributions). *Adopt the hypothesis of Corollary 12 and further let Q'_i be such that $\Delta(Q_i, Q'_i) \leq \varepsilon/M$. Then*

$$\Pr\left[\mathsf{RejSampLoop}'_{P,(Q_i)_{i=1}^T,M}(1^m) = \perp\right] \leq 2\sqrt{\varepsilon} \quad \textit{and} \tag{3}$$

$$\Delta(\mathsf{RejSampLoop}'_{P,(Q_i)_{i=1}^T,M}(1^m), \mathsf{RejSampLoop}^*_{P,M,T}(1^m)) \leq 2T \cdot \varepsilon/M. \tag{4}$$

Proof. $\mathsf{RejSamp}_{P,Q_i,M}$ and $\mathsf{RejSamp}'_{P,Q_i,M}$ apply the same statistical test to Q_i and Q'_i respectively, namely: take a sample x and accept or reject with respect to the function $P(x)/MQ_i(x)$. Therefore, by data processing $\Delta(\mathsf{RejSamp}_{P,Q_i,M}, \mathsf{RejSamp}'_{P,Q_i,M}) \leq \varepsilon/M$ and for all $i \in [T]$, recalling $\Pr\left[\mathsf{RejSamp}_{P,Q_i,M}(1^m) = \perp\right] \leq 1 - \frac{1-\varepsilon}{M}$, we have

$$\Pr\left[\mathsf{RejSamp}'_{P,Q_i,M}(1^m) = \perp\right] \leq 1 - \frac{1-2\varepsilon}{M} \text{ therefore,}$$

$$\Pr\left[\mathsf{RejSampLoop}'_{P,(Q_i)_{i=1}^T,M}(1^m) = \perp\right] \leq \left(1 - \frac{1-2\varepsilon}{M}\right)^T \leq \frac{1}{1 + \frac{1-2\varepsilon}{M} \cdot T} \leq 2\sqrt{\varepsilon}$$

where we used $1 + \frac{1-2\varepsilon}{\sqrt{\varepsilon}} \geq \frac{1}{2\sqrt{\varepsilon}}$ for $\varepsilon \leq 1/2$. A hybrid gives

$$\Delta(\mathsf{RejSampLoop}_{P,(Q_i)_{i=1}^T,M}, \mathsf{RejSampLoop}'_{P,(Q_i)_{i=1}^T,M}) \leq T\varepsilon/M$$

which combined with Eq. (2) in Corollary 12 gives Eq. (4). □

We give a more compact self-contained version of Corollary 13.

Lemma 27 (Exponential time rejection sampling, compact). *Let $m \in \mathbb{N}$, $\varepsilon \in (0, 1/2]$, $\Lambda \subseteq \mathbb{R}^m$ be a lattice,*

$$M \geq \exp\left(2\pi m + 2\sqrt{2\pi m \ln \frac{1}{\varepsilon}}\right) \geq 1, \qquad T \geq \frac{M}{\sqrt{\varepsilon}},$$

$s > 0$ and $\mathbf{c}_1, \ldots, \mathbf{c}_T \in \mathbb{R}^m$ with $\|\mathbf{c}_i\| \leq s\sqrt{2m}$ for all $i \in [T]$. Let $P = D_{\Lambda,s}$ and let $Q_i = D_{\Lambda,s,\mathbf{c}_i}$ for $i \in [T]$. Let Q_i' be such that $\Delta(Q_i, Q_i') \leq \varepsilon/M$. Given access to an oracle for Q_i', there exists an algorithm $\mathcal{A}$ that fails with probability at most $2\sqrt{\varepsilon}$, that samples from a distribution within statistical distance $2T \cdot \varepsilon/M$ to P when it succeeds and has additive $\mathsf{poly}(m)$ time and memory overhead compared to sampling from Q_i' sequentially T times.

Note that $\mathcal{A}$ in Lemma 27 is $\mathsf{RejSampLoop}'$ in Corollary 13.

7 Hinted SIS problems

We formalise new problems that are hinted versions of (I)SIS. One such problem is H-ISIS and in Appendix A we reduce omISIS to H-ISIS. Of importance for our main results in §8 is a reduction from a problem related to H-SIS to H-ISIS that provides the core subroutine for Theorem 6.

7.1 Definitions of hinted problems

We define problems called H-SIS, H-ISIS, H-SIIS and H-SIGS which are hinted variants of SIS, ISIS, SIIS and SIGS. For example, a H-ISIS solution is still a short preimage under $\mathbf{A}$ of a uniform $\mathbf{t} \in \mathbb{Z}_q^n$, like ISIS, but the adversary is also given k elements of $\Lambda_q^\perp(\mathbf{A})$ as hints.

Definition 14 (H-SIS, H-ISIS, H-SIIS and H-SIGS[11]). *For each hinted problem let* $\mathsf{params} = (k, m, q, \beta, \mathsf{Dist})$ *be parametrised by* n *with* $k \in \mathbb{N}$, $m, q \in \mathbb{N}^+$, $m \geq n$, $q \geq 2$, $\beta > 0$ *and where* Dist *is a function that maps* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *to a distribution over* $\Lambda_q^\perp(\mathbf{A})^k$. *The experiments for the* $\mathsf{H\text{-}SIS}_{\mathsf{params}}$, $\mathsf{H\text{-}ISIS}_{\mathsf{params}}$, $\mathsf{H\text{-}SIIS}_{\mathsf{params}}$ *and* $\mathsf{H\text{-}SIGS}_{\mathsf{params}}$ *problems are defined in Fig. 6.*

Note that there are trivial reductions from H-SIS to H-SIIS, and from H-SIIS to H-SIGS.

Remark 7. As discussed above, alternative formulations of H-SIS, H-ISIS, H-SIIS and H-SIGS could require $\mathbf{u}$ or $\mathbf{U}$ to have norms relative to the hints rather than an absolute bound β, similar to IncGDD in [MR07]. Indeed, when constructing adversaries from H-ISIS adversaries this is more natural. Moreover, it would better echo the central 'message' of this work: that improving an already short set of vectors is presumably hard. However, we opted for an absolute bound to match conventions of related problems in the literature.

The difficulty of H-SIS, H-ISIS, H-SIIS and H-SIGS depends crucially on Dist. In particular, if β is sufficiently larger than the maximum length of the hints and their real span is $\mathbb{R}^m$ then the problems can be efficiently solved using e.g. Babai's nearest plane (Lemma 17).

Proposition 5. *Let* $\mathsf{params} = (k, m, q, \beta, \mathsf{Dist})$ *be parametrised by* n *such that* $n, q \in \mathbb{N}^+$ *with* q *prime,* $q^{n-m} \in \mathsf{negl}(n)$ *and*

$$\Pr\left[\mathrm{span}_{\mathbb{R}}(\mathsf{Dist}(\mathsf{U}(\mathbb{Z}_q^{n\times m}))) = \mathbb{R}^m \wedge \left\|\mathsf{Dist}(\mathsf{U}(\mathbb{Z}_q^{n\times m}))\right\| \leq 2\beta/\sqrt{m}\right] \notin \mathsf{negl}(n)$$

then there exists a PPT $\mathcal{A}$ *that solves* $\mathsf{H\text{-}ISIS}_{\mathsf{params}}$ *with non-negligible probability.*

Proof. Proposition 2 ensures $P_{\mathbf{A},\mathbf{t}} \neq \emptyset$ with overwhelming probability. Let $\mathcal{A}$ find $\mathbf{v} \in P_{\mathbf{A},\mathbf{t}}$ via Gaussian elimination and form a basis of a sublattice of $\Lambda_q^\perp(\mathbf{A})$. Start with $\mathbf{U} = (\mathbf{u}_1), i = 2$ and iterate the following. If $\mathbf{u}_i \notin \mathrm{span}_{\mathbb{R}}(\mathbf{U})$ append $\mathbf{u}_i$ to $\mathbf{U}$. Increment i. If $\mathsf{rank}(\mathbf{U}) = m$ then return it. With non-negligible probability $\mathbf{U}$ is a basis of a full rank sublattice of $\Lambda_q^\perp(\mathbf{A})$ and $\|\mathbf{U}\| \leq 2\beta/\sqrt{m}$. Lemma 17 with input $(\mathbf{U}, \mathbf{v})$ outputs $\mathbf{u} \in \mathcal{P}_{1/2}(\hat{\mathbf{U}})$ such that $\mathbf{u} = \mathbf{v} \bmod \Lambda_q^\perp(\mathbf{A})$. Therefore $\mathbf{u} \in P_{\mathbf{A},\mathbf{t}}$. Since $\|\hat{\mathbf{U}}\| \leq \|\mathbf{U}\|$ we have $\|\mathbf{u}\| \in P_{\beta,\mathbf{A},\mathbf{t}}$. □

[11] Suggested pronunciations are 'k-hint-i-sis', 'k-hint-sis', 'k-hint-cease' and 'k-hint-sigs' where 'sig' is as in 'signature'.

$\mathsf{Exp\text{-}H\text{-}SIS}_{(k,m,q,\beta,\mathsf{Dist}),\mathcal{A}}(1^n)$

$\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$
$(\mathbf{u}_1,\ldots,\mathbf{u}_k) \leftarrow \mathsf{Dist}(\mathbf{A})$
$\mathbf{u} \leftarrow \mathcal{A}(\mathbf{A},\mathbf{u}_1,\ldots,\mathbf{u}_k)$
// **return** $[\![\mathbf{u} \in P^+_{\beta,\mathbf{A}}]\!]$
return $[\![\mathbf{A}\cdot\mathbf{u} = \mathbf{0} \wedge 0 < \|\mathbf{u}\| \leq \beta]\!]$

$\mathsf{Exp\text{-}H\text{-}SIIS}_{(k,m,q,\beta,\mathsf{Dist}),\mathcal{A}}(1^n)$

$\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$
$(\mathbf{u}_1,\ldots,\mathbf{u}_k) \leftarrow \mathsf{Dist}(\mathbf{A})$
$\mathbf{U} \leftarrow \mathcal{A}(\mathbf{A},\mathbf{u}_1,\ldots,\mathbf{u}_k)$
// **return** $[\![\mathbf{U} \subset P_{\beta,\mathbf{A}} \wedge \mathsf{rank}(\mathbf{U}) = m]\!]$
return $[\![\mathbf{A}\cdot\mathbf{U} = \mathbf{0} \wedge \|\mathbf{U}\| \leq \beta \wedge \mathsf{rank}(\mathbf{U}) = m]\!]$

$\mathsf{Exp\text{-}H\text{-}ISIS}_{(k,m,q,\beta,\mathsf{Dist}),\mathcal{A}}(1^n)$

$(\mathbf{A},\mathbf{t}) \leftarrow \mathbb{Z}_q^{n\times m} \times \mathbb{Z}_q^n$
$(\mathbf{u}_1,\ldots,\mathbf{u}_k) \leftarrow \mathsf{Dist}(\mathbf{A})$
$\mathbf{u} \leftarrow \mathcal{A}(\mathbf{A},\mathbf{t},\mathbf{u}_1,\ldots,\mathbf{u}_k)$
// **return** $[\![\mathbf{u} \in P_{\beta,\mathbf{A},\mathbf{t}}]\!]$
return $[\![\mathbf{A}\cdot\mathbf{u} = \mathbf{t} \wedge \|\mathbf{u}\| \leq \beta]\!]$

$\mathsf{Exp\text{-}H\text{-}SIGS}_{(k,m,q,\beta,\mathsf{Dist}),\mathcal{A}}(1^n)$

$\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$
$(\mathbf{u}_1,\ldots,\mathbf{u}_k) \leftarrow \mathsf{Dist}(\mathbf{A})$
$\mathbf{U} \leftarrow \mathcal{A}(\mathbf{A},\mathbf{u}_1,\ldots,\mathbf{u}_k)$
// **return** $[\![\mathbf{U} \subset P_{\beta,\mathbf{A}} \wedge \Lambda(\mathbf{U}) = \Lambda_q^\perp(\mathbf{A})]\!]$
return $[\![\mathbf{A}\cdot\mathbf{U} = \mathbf{0} \wedge \|\mathbf{U}\| \leq \beta \wedge \Lambda(\mathbf{U}) = \Lambda_q^\perp(\mathbf{A})]\!]$

Figure 6. The experiments for H-SIS, H-ISIS, H-SIIS and H-SIGS.

7.2 'Sieving with centres'

The standard SIS to ISIS reduction samples a short and entropic $\mathbf{u} \in \mathbb{Z}^m$, asks an ISIS adversary to find a preimage $\mathbf{u}'$ for $\mathbf{t} = \mathbf{A}\cdot\mathbf{u}$ and outputs $\mathbf{u}^* = \mathbf{u} - \mathbf{u}'$. However, to solve H-SIGS rather than H-SIS, which we require for our iterative sieve approach in §8, we pursue an alternative approach for which the core subroutine is Lemma 28 below.

The intuition is to output $(\mathbf{u},\mathbf{u}')$ rather than $\mathbf{u}^* = \mathbf{u} - \mathbf{u}'$ and consider a different condition for success where the pair $(\mathbf{u},\mathbf{u}')$ is a valid solution if $\mathbf{u}' \in P_{\beta,\mathbf{A},\mathbf{A}\cdot\mathbf{u}}$, i.e. $\mathbf{u}'$ is a H-ISIS solution for $(\mathbf{A},\mathbf{A}\cdot\mathbf{u})$. In particular, other than $\mathbf{u}$ sharing a coset with $\mathbf{u}'$ it may be entirely arbitrary.

Lemma 28 (Sieving with centres subroutine). *Let $\tau_{\mathcal{A}},\mu_{\mathcal{A}},\varepsilon_{\mathcal{A}},\mathsf{params}_{\mathcal{A}} = (k,m,q,\beta,\mathsf{Dist})$ be parametrised by n with q prime and $m \geq n$. For $\delta \in (0,1/2)$ and $f\colon \mathbb{N} \to [1,q/4)$ let $s \geq 4f(m)\cdot\eta_\delta(m)$. If there is a $(\tau_{\mathcal{A}},\mu_{\mathcal{A}},\varepsilon_{\mathcal{A}})$ adversary $\mathcal{A}$ against $\mathsf{H\text{-}ISIS}_{\mathsf{params}_{\mathcal{A}}}$ then there exists a $(\tau_{\mathcal{B}},\mu_{\mathcal{B}},\varepsilon_{\mathcal{B}})$ algorithm $\mathcal{B}$ that on input uniform $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and $(\mathbf{u}_1,\ldots,\mathbf{u}_k) \leftarrow \mathsf{Dist}(\mathbf{A})$ outputs $(\mathbf{u},\mathbf{u}') \in \mathbb{Z}^m \times \mathbb{Z}^m$ with $\mathbf{u}' \in P_{\beta,\mathbf{A},\mathbf{A}\cdot\mathbf{u}}$ and*

$$\tau_{\mathcal{B}} \approx \tau_{\mathcal{A}}, \quad \mu_{\mathcal{B}} \approx \mu_{\mathcal{A}}, \text{ and } \varepsilon_{\mathcal{B}} \geq \varepsilon_{\mathcal{A}} - (2\delta + 1/q^{m-n} + C_{f,n}).$$

If $\delta \in \mathsf{negl}(m)$, $m \geq cn$ for some $c > 1$ and $f(m) = q^{n/m}$ then $\varepsilon_{\mathcal{B}} = \varepsilon_{\mathcal{A}} - \mathsf{negl}(m)$.

Proof. We construct $\mathcal{B}$. On input $(\mathbf{A},(\mathbf{u}_1,\ldots,\mathbf{u}_k))$ first $\mathcal{B}$ samples $\mathbf{u} \leftarrow D_{\mathbb{Z}^m,s}$ via e.g. Corollary 5 with $N = 1$. Then $\mathcal{B}$ receives $\mathbf{u}' \leftarrow \mathcal{A}(\mathbf{A},\mathbf{A}\cdot\mathbf{u},(\mathbf{u}_1,\ldots,\mathbf{u}_k))$ and outputs $(\mathbf{u},\mathbf{u}')$. First, $\tau_{\mathcal{B}}$ and $\mu_{\mathcal{B}}$ are clear. If $(\mathbf{u},\mathbf{u}')$ is output by $\mathcal{B}$ then $\Pr[\mathbf{u}' \in P_{\beta,\mathbf{A},\mathbf{A}\cdot\mathbf{u}}]$ is the success probability of $\mathcal{A}$ on H-ISIS instances with targets some statistical distance from uniform. Via the leftover hash lemma (Corollary 8) we have $\Pr[\mathbf{u}' \in P_{\beta,\mathbf{A},\mathbf{A}\cdot\mathbf{u}}] \geq \varepsilon_{\mathcal{A}} - (2\delta + 1/q^{m(1-n/m)} + C_{f,n})$ and conclude. □

In Theorem 6 we consider outputs $(\mathbf{u}_i,\mathbf{u}_i')$ and select close $\mathbf{u}_i',\mathbf{u}_j'$ with $i \neq j$. Given that correct $\mathbf{u}_i - \mathbf{u}_i'$ are distributed as $D_{\Lambda_q^\perp(\mathbf{A}),s,-\mathbf{u}_i'}$, i.e. a distribution with centre $-\mathbf{u}_i'$, we call this approach 'sieving with centres'.

8 Polynomial memory sieve

We construct an NNS (or 'sieve-style') algorithm that finds short lattice vector using only polynomial memory assuming the existence of H-ISIS adversaries. Our main results are stated in Lemma 29 and Theorem 7. Our construction relies notably on (a) Lemma 19, which lower bounds the probability of finding close centres by repeatedly running the H-ISIS adversary; (b) Lemma 23, which shows that non zero-centred Gaussian lattice vectors generate the lattice and (c) Lemma 27, which shows we can output zero-centred Gaussian lattice vectors. We may realise the latter two using Lemma 28, which assumes an H-ISIS adversary and a sampler for $D_{\mathbb{Z}^m,s}$. We ultimately set s according to β and (shrink factor) γ.

Theorem 6 essentially states that if there exists a polynomial time H-ISIS adversary finding solutions of a given norm then there exists a single-exponential time and polynomial memory H-SIGS adversary which finds a generating set of norm a constant factor shorter.[12] This H-SIGS adversary outputs vectors that follow non zero-centred discrete Gaussian distributions. In Lemma 29 we provide a 'clean' version of Theorem 6 which uses rejection sampling to construct a H-SIGS adversary which outputs vectors following zero-centred discrete Gaussian distributions. To obtain generating sets with much shorter norm we repeatedly apply Lemma 29 by feeding the generating set that a H-SIGS adversary finds into the next H-SIGS adversary as hints. This critically relies on there existing 'chains' of adversaries that cooperate in the sense of accepting the outputs of a previous adversary as their hint inputs, i.e. all adversaries succeeding on the same $\mathbf{A}$, as captured in Theorem 7.

8.1 Sampling a shorter generating set

We find many vectors of length β/γ that follow (up to a small distance) discrete Gaussian distributions. We take inspiration from the Gauss sieve [MV10b], which constructs a database of N vectors and checks each pair, in time $O(N^2)$, to see if their difference is short. Here, we check the differences of $O(N)$ pairs of vectors in a streaming fashion in time $O(N)$. To do so we use Lemma 19, a consequence of Motzkin–Straus (Theorem 5), on the outputs $(\cdot, \mathbf{u}')$ of $\mathcal{B}$ in Lemma 28.

Theorem 6 (H-SIGS $\leq$ H-ISIS). *Let* $\mathsf{params} = (k, m, q, \beta, \mathsf{Dist})$, $\gamma > 1$, $\delta \in (0, 1/2)$, $(\tau_{\mathcal{A}}, \mu_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$ *and* $(\tau_{\mathsf{Samp}}, \mu_{\mathsf{Samp}})$ *be parametrised by* n *with* params *as in Fig. 6. Let* $\mathsf{params}_{\mathcal{C}} = (k, m, q, \beta/\gamma, \mathsf{Dist})$, $f\colon \mathbb{N} \to [1, q/4)$ *and* $\nu(f, n) = 2\delta + 1/q^{m-n} + C_{f,n}$.

If $\mathcal{A}$ *is a* $(\tau_{\mathcal{A}}, \mu_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$ *adversary against* $\mathsf{H\text{-}ISIS}_{\mathsf{params}}$, q *is prime,* Samp *is a* τ_{Samp} *time and* μ_{Samp} *memory algorithm that samples* $D_{\mathbb{Z}^m,s}$ *and*

- $s \geq 4f(m) \cdot \eta_\delta(m)$,
- $\sqrt{(m/2)} \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A})) \leq s \leq \beta/(2\gamma\sqrt{2m})$,
- $s \leq 2^{\mathsf{poly}(m)}$

then there exists a $(\tau_{\mathcal{C}}, \mu_{\mathcal{C}}, \varepsilon_{\mathcal{C}})$ *adversary* $\mathcal{C}$ *against* $\mathsf{H\text{-}SIGS}_{\mathsf{params}_{\mathcal{C}}}$ *such that for large enough* n

- $\tau_{\mathcal{C}} \leq 2RN \cdot (\tau_{\mathcal{A}} + \tau_{\mathsf{Samp}})$,
- $\mu_{\mathcal{C}} \leq 2R \cdot (\mu_{\mathcal{A}} + \mu_{\mathsf{Samp}})$,
- $\varepsilon_{\mathcal{C}} \geq (\varepsilon_{\mathcal{A}} - \nu(f, n))/2 \cdot (1 - \mathsf{negl}(m))$,

where N, R *are integers with*

$$R = \left\lceil 2(1 + e^\pi) \cdot m \cdot (1 + \log(s\sqrt{m})) \right\rceil \in \mathsf{poly}(m),$$
$$N = \left\lceil 16m \cdot 2^{m \log(1+4\gamma)} / (\varepsilon_{\mathcal{A}} - \nu(f, n))^2 \right\rceil.$$

In particular, asymptotically we have $\tau_{\mathcal{C}} \in 2^{O(m \log \gamma)} \cdot (\tau_{\mathcal{A}} + \tau_{\mathsf{Samp}})/(\varepsilon_{\mathcal{A}} - \nu(f, n))^2$ *and* $\varepsilon_{\mathcal{C}} \in \Omega(\varepsilon_{\mathcal{A}} - \nu(f, n))$.

[12] While we highlight the single-exponential regime here the formal statements are more general, see 'Parameters' and especially Remark 8 below.

Proof. We construct $\mathcal{C}$ on H-SIGS instance $p = (\mathbf{A}, \mathbf{u}_1, \ldots, \mathbf{u}_k)$. First, we define subroutines. Let $\mathcal{B}(p)$ sample $\mathbf{v} \leftarrow \mathsf{Samp}$ and receive $\mathbf{v}' \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{A}\cdot\mathbf{v}, \mathbf{u}_1, \ldots, \mathbf{u}_k)$. If $\mathbf{v}' \in P_{\beta,\mathbf{A},\mathbf{A}\cdot\mathbf{v}}$ then $\mathcal{B}(p)$ returns $(\mathbf{v}, \mathbf{v}')$ else it returns $(\perp, \perp)$. Note that $\mathcal{B}$ is as in Lemma 28 except with explicitly marked failures. Let $\mathcal{I}(p)$ call $\mathcal{B}(p)$ twice. If either call returns $(\perp, \perp)$ then $\mathcal{I}(p)$ returns $\perp$. Else let $(\mathbf{v}_i, \mathbf{v}'_i) \leftarrow \mathcal{B}(p)$ for $i \in \{\ell, r\}$. If $\|\mathbf{v}'_\ell - \mathbf{v}'_r\| \le \beta/2\gamma$ then $\mathcal{I}(p)$ returns $(\mathbf{v}_\ell, \mathbf{v}'_\ell), (\mathbf{v}_r, \mathbf{v}'_r)$, else it returns $\perp$. Let $\mathcal{O}(p)$ call $\mathcal{I}(p)$ up to N times. If any call does not return $\perp$ then $\mathcal{O}(p)$ returns the output of this call, else if all calls return $\perp$ then $\mathcal{O}(p)$ returns $\perp$. Given these subroutines, $\mathcal{C}$ first sets $\mathbf{U} = ()$ and makes R calls $\mathcal{O}(p)$. Each call that does not return $\perp$ returns $(\mathbf{v}_\ell, \mathbf{v}'_\ell), (\mathbf{v}_r, \mathbf{v}'_r)$ and $\mathcal{C}$ appends $(\mathbf{v}_\ell - \mathbf{v}_r) - (\mathbf{v}'_\ell - \mathbf{v}'_r)$ to $\mathbf{U}$. Finally, $\mathcal{C}$ returns $\mathbf{U}$.

Immediately $\tau_\mathcal{C} \le 2RN \cdot (\tau_\mathcal{A} + \tau_{\mathsf{Samp}})$ and $\mu_\mathcal{C} \le 2R \cdot (\mu_\mathcal{A} + \mu_{\mathsf{Samp}})$. Let $\varepsilon_\mathcal{B}$ denote the probability that $\mathcal{B}$ does not return $(\perp, \perp)$ over p sampled according to $\mathsf{params}_\mathcal{C}$ and let $\varepsilon_{\mathcal{B},p}$ denote the probability that $\mathcal{B}(p)$ does not return $(\perp, \perp)$ for fixed p sampled according to $\mathsf{params}_\mathcal{C}$. We have $\Pr[\varepsilon_{\mathcal{B},p} \ge \varepsilon_\mathcal{B}/2] \ge \varepsilon_\mathcal{B}/2$ by Proposition 1 and therefore that

$$\begin{aligned}\Pr[\mathcal{C}(p) \text{ wins H-SIGS}] &\ge \Pr[\mathcal{C}(p) \text{ wins H-SIGS} \mid \varepsilon_{\mathcal{B},p} \ge \varepsilon_\mathcal{B}/2] \cdot \Pr[\varepsilon_{\mathcal{B},p} \ge \varepsilon_\mathcal{B}/2] \\ &\ge \Pr[\mathcal{C}(p) \text{ wins H-SIGS} \mid \varepsilon_{\mathcal{B},p} \ge \varepsilon_\mathcal{B}/2] \cdot \varepsilon_\mathcal{B}/2.\end{aligned}$$

By Lemma 28 we have $\varepsilon_\mathcal{B} \ge \varepsilon_\mathcal{A} - \nu(f, n)$. We now assume $\varepsilon_{\mathcal{B},p} \ge \varepsilon_\mathcal{B}/2$. Let $\varepsilon_\mathcal{I} = \Pr[\mathcal{I}(p) \ne \perp]$, or equivalently the probability that the second outputs of two calls to $\mathcal{B}(p)$ have $\|\mathbf{v}'_\ell - \mathbf{v}'_r\| \le \beta/2\gamma$. By applying Lemma 19 to the second output of $\mathcal{B}(p)$ we have $\varepsilon_\mathcal{I} \ge \varepsilon_{\mathcal{B},p}^2/2^{m\log(1+4\gamma)}$. Setting $\alpha = 1$, $\beta = \varepsilon_\mathcal{I}$ and $t = m$ in Lemma 4 we have

$$N = 4m/\varepsilon_\mathcal{I} \le 16m \cdot 2^{m\log(1+4\gamma)}/(\varepsilon_\mathcal{A} - \nu(f, n))^2$$

is sufficient for N calls to $\mathcal{I}$ to have one success with probability at least $1 - 2^{-m}$. Therefore $\varepsilon_\mathcal{O} = \Pr[\mathcal{O}(p) \ne \perp] \ge 1 - 2^{-m}$. The probability all R calls to $\mathcal{O}$ do not fail is at least $1 - R \cdot (1 - (1 - 2^{-m})) = 1 - R \cdot 2^{-m}$. Therefore $\mathbf{U}$ has R columns with probability at least $1 - R \cdot 2^{-m}$.

We now describe the distributions and norms of these columns. Given a pair $(\mathbf{v}, \mathbf{v}')$ output by $\mathcal{B}$ with $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$ and $\mathbf{v}' \in P_{\beta,\mathbf{A},\mathbf{A}\cdot\mathbf{v}}$ the distribution of $\mathbf{v}$ follows $\mathbf{v}' + D_{\Lambda_q^\perp(\mathbf{A}),s,-\mathbf{v}'}$ so that $\mathbf{v} - \mathbf{v}' \leftarrow D_{\Lambda_q^\perp(\mathbf{A}),s,-\mathbf{v}'}$. The difference of a pair of such pairs, $(\mathbf{v}_\ell - \mathbf{v}_r) - (\mathbf{v}'_\ell - \mathbf{v}'_r)$ is statistically close to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s,\mathbf{v}'_r-\mathbf{v}'_\ell}$ which has norm concentrated below $\|\mathbf{v}'_r - \mathbf{v}'_\ell\| + s\sqrt{2m}$. We accepted only close $(\mathbf{v}'_\ell, \mathbf{v}'_r)$ and chose s bounded above such that this sum is at most β/γ.

Indeed, by assumption $s \ge \sqrt{m/2} \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A}))$ so there exists some $\delta'(m) \in \mathsf{negl}(m)$ such that $\eta_{\delta'}(\Lambda_q^\perp(\mathbf{A})) \le s$ via Lemma 10. We may ensure that $s \ge \sqrt{2}\eta_\delta(\Lambda_q^\perp(\mathbf{A}))$ by letting $\delta = 1/(\sqrt{(1 + 1/\delta')/(2m)} - 1)$. If $\delta' \in \mathsf{negl}(m)$ then $\delta \in \mathsf{negl}(m)$. Then columns of $\mathbf{U}$ have negligible statistical distance to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s,\mathbf{v}'_r-\mathbf{v}'_\ell}$ by Lemma 16. This distribution has norm at most $s\sqrt{2m}$ except with probability $2^{-m}(1+\delta)/(1-\delta)$ by Lemma 11 and we have set $s \le \beta/(2\gamma\sqrt{2m})$. Therefore the probability that $\mathbf{U} \subset P_{\beta/\gamma,\mathbf{A}}$ is at least $1 - R \cdot \mathsf{negl}(m)$.

Finally, each of the R columns of $\mathbf{U}$ is negligibly close to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s,\mathbf{c}}$ for some $\mathbf{c}$ of norm at most $\beta/2\gamma$. Therefore, the probability that $\mathbf{U}$ generates $\Lambda_q^\perp(\mathbf{A})$ is at least $1 - R \cdot \mathsf{negl}(m)$ by Lemma 23, noting $\sqrt{2}s \ge \sqrt{m} \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A}))$. Therefore, given $R \in \mathsf{poly}(m)$ we have

$$\varepsilon_\mathcal{C} \ge (1 - \mathsf{negl}(m)) \cdot \varepsilon_\mathcal{B}/2 \ge (1 - \mathsf{negl}(m)) \cdot (\varepsilon_\mathcal{A} - \nu(f, n))/2. \qquad \square$$

Parameters. If $\nu(f, n) \in \mathsf{negl}(m)$ and $\varepsilon_\mathcal{A} \ge 1/\mathsf{poly}(m)$ then $\varepsilon_\mathcal{C} \ge 1/\mathsf{poly}(m)$. If additionally $\tau_\mathcal{A}, \tau_{\mathsf{Samp}} \in \mathsf{poly}(m)$ then $\tau_\mathcal{C} \in 2^{O(m\log\gamma)}$. If $\log\gamma \in o(\log m)$ then $\tau_\mathcal{C} \in 2^{o(m\log m)}$, e.g. $\gamma \in O(1)$ gives $\tau_\mathcal{C} \in 2^{O(m)}$. Finally, if $\mu_\mathcal{A}, \mu_{\mathsf{Samp}} \in \mathsf{poly}(m)$ then $\mu_\mathcal{C} \in \mathsf{poly}(m)$ and given the above we have constructed a weakly superexponential time, polynomial memory and non negligible success probability algorithm for H-SIGS on $\mathsf{params}_\mathcal{C}$.

The N calls to Samp are each realised in polynomial time via Corollary 5 and we choose γ. We therefore have only the size of $\nu(f,n)$ and the ability to select a satisfying s to make explicit.

If $m \geq cn$ for $c > 1$ then $1/q^{m-n} \leq 1/q^{c'm} \leq 1/2^{c'm}$ for some $c' > 0$. Furthermore, if $f(m) \geq q^{n/m}$ then $C_{f,n} \leq (3/4)^m$ and $f(m) \leq q/4$. It is then enough to set $\delta \in \mathsf{negl}(m)$ to have $\nu(f,n) \in \mathsf{negl}(m)$. (In common settings $q \in \Omega(n)$ and $m \geq n \log q$, in which case $1/q^{m-n}$ is superexponentially small and constant f is sufficient.)

Finally, if $B(n) = 4(1+\sqrt{m/2\pi}) \cdot q^{n/m} \cdot \sqrt{\log(6m)/\pi}$ then Corollary 9 gives that $\lambda_m(\Lambda_q^\perp(\mathbf{A})) \leq 2B(n)$ except with probability $(3/4)^m$. Therefore, if $\beta \in 2^{\mathsf{poly}(m)}$ is not too large, but $\beta \geq 4B(n) \cdot \gamma m$ and $\beta \geq \sqrt{128}\gamma\sqrt{m} \cdot f(m) \cdot \eta_\delta(m)$ are large enough, then setting $s = \beta/(2\gamma\sqrt{2m})$ is satisfying except with exponentially small probability. Setting $f(m) = q^{n/m}$ and $\delta \in 2^{-o(m^2)}$, we have that $\beta \geq 4B(n) \cdot \gamma m$ is sufficient for large enough n. Given $\mathsf{negl}(m) \cap 2^{-o(m^2)} \neq \emptyset$ we can select satisfying parameters.

Remark 8. Theorem 6 is particularly interesting when Dist is such that, with overwhelming probability over $\mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n \times m})$ and $\mathbf{U} \leftarrow \mathsf{Dist}(\mathbf{A})$, we have $\|\mathbf{U}\|_{\min} \geq \beta/\gamma_\uparrow$ and $\gamma = \gamma_\uparrow \cdot \gamma_\downarrow$ with $1 < \gamma_\uparrow, \gamma_\downarrow \leq \mathsf{polylog}\,(m)$. Then our reduction turns a $\mathsf{H\text{-}ISIS}$ adversary which returns solutions at most a growth factor $\gamma_\uparrow$ longer than the shortest hint, into a $\mathsf{H\text{-}SIGS}$ adversary which returns solutions at least a shrink factor $\gamma_\downarrow$ shorter than the shortest hint. Demanding $\gamma_\uparrow, \gamma_\downarrow \leq \mathsf{polylog}\,(m)$ implies $\log \gamma \in o(\log m)$.

8.2 Sampling a clean shorter generating set

In Theorem 6, the$\mathcal{C}$ that is constructed outputs vectors following distributions $\mathsf{negl}(n)$ close to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s,\mathbf{c}_i}$ with $\|\mathbf{c}_i\| \leq \beta/2\gamma$. Lemma 29 'cleans up' these distributions to be $\mathsf{negl}(n)$ close to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$, i.e. to zero-centred Gaussians, with a multiplicative time overhead of $\exp(m)$. This follows from composing the proof of Theorem 6 with rejection sampling via Lemma 27.

Lemma 29 (Clean $\mathsf{H\text{-}SIGS} \leq \mathsf{H\text{-}ISIS}$). *Adopt the notation of Theorem 6. Let*

$$\sqrt{Cm/\pi} \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A})) \leq s = \beta/(2\gamma\sqrt{2m})$$

for some $C \geq 4(1+\pi/\ln 2)$. Then there exists a $(\tau_\mathcal{D}, \mu_\mathcal{D}, \varepsilon_\mathcal{D})$ adversary $\mathcal{D}$ for $\mathsf{H\text{-}SIGS}_{\mathsf{params}_\mathcal{C}}$ such that for all large enough n

$$\tau_\mathcal{D} \leq \exp(4\pi m) \cdot 2^m \cdot \tau_\mathcal{C}, \quad \mu_\mathcal{D} \leq \mu_\mathcal{C} + \mathsf{poly}(m), \quad \varepsilon_\mathcal{D} \geq \varepsilon_\mathcal{C} - \mathsf{negl}(m)\ .$$

The output of $\mathcal{D}$ is formed of vectors with $\mathsf{negl}(m)$ statistical distance from $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$.

Proof. As $C/\pi \geq 1/2$ the choice of s is consistent with Theorem 6. This proof is similar to its proof. Instead of collecting samples close to $D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s,\mathbf{c}_i}$ until we form a generating set, we perform rejection sampling on them and collect zero-centred outputs until they form a generating set. This requires many more calls to $\mathcal{O}$ and accept or reject their outputs.

Following the proof of Theorem 6 the distribution $D'_{i,j}$ of $(\mathbf{v}_i - \mathbf{v}_j) - (\mathbf{v}'_i - \mathbf{v}'_j)$ is negligibly close to $Q_\imath = D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s,\mathbf{v}'_j - \mathbf{v}'_i}$ (for index $\imath = \imath(i,j)$ labelled as a function of i,j). The $\mathsf{negl}(n)$ distance comes from Lemma 16 and is at most 3δ for $s \geq \sqrt{\ln(2m \cdot (1+1/\delta))/\pi} \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A})) \geq \eta_\delta(\Lambda_q^\perp(\mathbf{A}))$ (Lemma 10). Note for any constant $C > 0$ such that $s \geq \sqrt{Cm/\pi} \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A}))$ we have $\delta \leq 2^{-C \cdot m}$ for large n.

Consider Lemma 27 with $Q'_{\imath(i,j)} \coloneqq D'_{i,j}$ and $P = D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$. Set $\varepsilon = 2^{-2m}$, $M \geq \exp(4\pi m)$ and note $M/\varepsilon = \exp(4\pi m) \cdot 2^{2m} \in 2^{O(m)}$. Set $C \geq 4(1+\pi/\ln 2)$ so that $3\delta \leq \varepsilon/M$ and note $\left\|\mathbf{c}_{\imath(i,j)}\right\| = \left\|\mathbf{v}'_i - \mathbf{v}'_j\right\| \leq s\sqrt{2m}$ since $s\sqrt{2m} = \beta/2\gamma$ by assumption. We therefore satisfy the

conditions of Lemma 27. The pair $\left((Q'_i)_{1\le i\le T}, P^T\right)$ consumes $T = M/\sqrt{\varepsilon} = \exp(4\pi m)\cdot 2^m \in 2^{O(m)}$ samples, each from some Q'_i, and allows us to output a sample with distance $2\sqrt{\varepsilon} = 2\cdot 2^{-m}$ to P in polynomial memory. The running time $\tau_{\mathcal{D}}$ increases by a factor of $\exp(4\pi m)\cdot 2^m$ relative to $\tau_{\mathcal{C}}$.

Finally, again by Lemma 27 we have that rejection sampling T times outputs $\perp$ with probability at most $2\sqrt{\varepsilon} \in \mathsf{negl}(m)$. To apply Lemma 23 (now for $\mathbf{c}_i = \mathbf{0}$, which is still permissible) we must collect $R \in \mathsf{poly}(m)$ outputs not equal to $\perp$. Thus all R rejection sampling routines output a sample exponentially close to P except with negligible probability and $\varepsilon_{\mathcal{D}} \ge \varepsilon_{\mathcal{C}} - \mathsf{negl}(m)$. □

8.3 Chaining adversaries

In Lemma 29, if the hints have lower bounded length as in Remark 8, then we construct an algorithm $\mathcal{A}_1$ that takes as input $\mathbf{A}$ and $\mathbf{U}_0$ and outputs $\mathbf{U}_1$ such that $\Lambda(\mathbf{U}_1) = \Lambda_q^{\perp}(\mathbf{A})$ and $\|\mathbf{U}_1\| \le \|\mathbf{U}_0\|_{\min}/\gamma_{\downarrow}$. We may thus hope to run a related algorithm $\mathcal{A}_2$ that takes as input $\mathbf{A}$ and $\mathbf{U}_1$ and outputs some $\mathbf{U}_2$ such that $\Lambda(\mathbf{U}_2) = \Lambda_q^{\perp}(\mathbf{A})$ and $\|\mathbf{U}_2\| \le \|\mathbf{U}_1\|_{\min}/\gamma_{\downarrow}$, and so on.

In general, we cannot derive the existence of $\mathcal{A}_{i+1}$ from the existence of $\mathcal{A}_i$. Moreover, even if we assume such a chain, we need to account for the fact that each $\mathcal{A}_i$ may succeed only with some probability ε_i over the randomness of the problem instance and its internal randomness. We thus start by establishing the degree to which we may rerandomise the problem instances $p_i = (\mathbf{A}, \mathbf{U}_i)$.

In Proposition 6 we observe that if we can sample one clean shorter generating set in single exponential time in Lemma 29 then we can sample $w \in \mathsf{poly}(m)$ clean shorter generating sets by repeating our algorithm roughly $w/\varepsilon_{\mathcal{D}}$ times. This allows us to randomise the hints $\mathbf{U}_i$ we give to $\mathcal{A}_{i+1}$ by picking w sufficiently large.[13]

This establishes that we must assume a chain of adversaries and a large enough set $A^{\checkmark} \subseteq \mathbb{Z}_q^{n\times m}$ favourable to all adversaries in the chain: Theorem 7 then formalises the chaining assumption and the required parameters.

The following follows from applying Corollary 2 with $k = w$ and $p = \varepsilon_{\mathcal{D}}$ to the output of Lemma 29.

Proposition 6 (\$-Hint). *Adopt the notation from Lemma 29. Then there exists an adversary $\mathcal{D}^+$ outputting w solutions for $\mathsf{H\text{-}SIGS}_{\mathsf{params}_{\mathcal{C}}}$ with probability $\varepsilon_{\mathcal{D}^+} \ge 1 - 2^{-w}$. Moreover, the vectors output by $\mathcal{D}$ each have $\mathsf{negl}(m)$ distance to $D_{\Lambda_q^{\perp}(\mathbf{A}),\sqrt{2}s}$, $\mathcal{D}$ runs in time $\tau_{\mathcal{D}^+} \le 4\cdot w/\varepsilon_{\mathcal{D}}\cdot\tau_{\mathcal{D}}$ and uses memory $\mu_{\mathcal{D}^+} \le w\cdot\mu_{\mathcal{D}}$.*

In Theorem 7 we assume a chain $\mathcal{A}_1, \ldots, \mathcal{A}_i, \ldots \mathcal{A}_z$ of adversaries where $\mathbf{A}$ is favourable to the chain, and essentially that each achieves a factor $\gamma_{\downarrow}$ improvement and succeeds with probability p_i.[14] The chain then achieves a factor $\gamma_{\downarrow}^z$ improvement and succeeds with probability $\prod_{i=1}^z p_i$. Given our assumption and Proposition 6, we can ensure that $\prod_{i=1}^z p_i \ge 1/\mathsf{poly}(m)$.

Below, the reader should think of the family of $\mathcal{A}_i$ constructed as in Lemma 29 from some other algorithm $\mathcal{A}^{\mathsf{H\text{-}ISIS}}$ that takes hints from $D_{\Lambda_q^{\perp}(\mathbf{A}),s_i}$ and outputs ISIS solutions of norm $\gamma_{\uparrow}\cdot\sqrt{m}\cdot s_i$ with success probabilities in $1/\mathsf{poly}(m)$. In particular, by Proposition 6 these algorithms are Gaussian samplers, as required in Theorem 7, for any polynomial number of outputs. Note that from Lemma 29, we may assume $\delta_i \in \mathsf{negl}(m)$ below.

Theorem 7 (Gaussian sampler self-reduction). *Let non negative integer $z \in \mathsf{poly}(n)$,*

$$\{\mathsf{params}_i\}_{i=1}^z = \left\{\left(k, m, q, \beta_i, D_{\Lambda_q^{\perp}(\mathbf{A}),s_i}\right)\right\}_{i=1}^z,\ \{(\tau_i, \mu_i, \varepsilon_i)\}_{i=1}^z$$

[13] This can be generalised to $w \in 2^{O(m)}$ which would require us to access the outputs of $\mathcal{A}_i$ in a streaming fashion, to avoid storing an exponential number of generating sets. We leave this for future work.

[14] While we state the lemma for arbitrary families of s_i, the most interesting application is where $s_{i+1} = s_i/\gamma_{\downarrow}$ for some fixed $\gamma_{\downarrow} \in \mathsf{polylog}(n)$.

and $A^{\checkmark} \subseteq \mathbb{Z}_q^{n \times m}$ *be parametrised by* n*. Let* $\beta_i = \sqrt{m} \cdot s_{i+1}$ *and let* $\mathcal{A}_i$ *be* $(\tau_i, \mu_i, \varepsilon_i)$ *adversaries for* $\mathsf{H\text{-}SIGS}_{\mathsf{params}_i}$ *with the additional guarantee that for all* $i \in [z]$ *they can be parameterised to output an arbitrary polynomial number of generating sets, each consisting of vectors from* Dist'_{i+1} *which is within* $\delta_{i+1} \in \mathsf{negl}(m)$ *distance to* $D_{\Lambda_q^\perp(\mathbf{A}),s_{i+1}}$*. Suppose*

$$\Pr\left[\mathbf{A} \in A^{\checkmark}\right] \geq 1/\mathsf{poly}(m)$$

and for all $i \in [z] \setminus \{1\}$ *we have*

$$\Pr\left[\mathsf{Exp\text{-}H\text{-}SIGS}_{\mathsf{params}_i,\mathcal{A}_i}(1^n) = 1 \;\middle|\; \mathbf{A} \in A^{\checkmark}\right] = p_i \in 1/\mathsf{poly}(m)$$

Then there exists a (τ, μ, ε) *adversary* $\mathcal{B}$ *for* $\mathsf{H\text{-}SIGS}_{k,m,q,\beta_z,D_{\Lambda_q^\perp(\mathbf{A}),s_1}}$ *with*

$$\tau \leq \sum_{i \in [z]} \mathsf{poly}(m) \cdot \tau_i, \quad \mu \leq \max\{\mu_i \cdot \mathsf{poly}(m)\}_{i \in [z]}, \quad \varepsilon \geq \varepsilon_1/\mathsf{poly}(m) \ .$$

Our proof essentially runs each $\mathcal{A}_i$ on sufficiently many instances to ensure an overwhelming probability of success per $\mathcal{A}_i$ such that their product over $i \in [z]$ is still overwhelming.

Proof. The result is immediate when $z \leq 1$ so consider $z \geq 2$. We have

$$\Pr\left[\mathsf{Exp\text{-}H\text{-}SIGS}_{(k,m,q,\beta_i,\mathsf{Dist}'_i),\mathcal{A}_i}(1^n) = 1 \;\middle|\; \mathbf{A} \in A^{\checkmark}\right] \geq p_i - k \cdot \delta_i$$

as probability of $\mathcal{A}_i$ succeeding with Dist'_i. Let $\hat{p}_i = p_i - k \cdot \delta_i \geq 1/\mathsf{poly}(m)$, as $k \in \mathsf{poly}(m)$ and $w_i = \lceil n/\hat{p}_{i+1} \rceil \in \mathsf{poly}(m)$ for $i \in [z-1]$. Let $\hat{p} = \min\{\hat{p}_i\}_i$ and $w = \lceil n/\hat{p} \rceil \in \mathsf{poly}(m)$. We run $\mathcal{A}_1$ with parameters chosen to produce $w_1 \leq w$ generating sets. This succeeds with probability $\varepsilon_1 \cdot (1 - 2^{-w_1}) \in \varepsilon_1 \cdot (1 - \mathsf{negl}(m))$. We abort if this step does not succeed.

Then for each $i \in [z] \setminus \{1\}$ we run $\mathcal{A}_i$ on $w_{i-1} \leq w$ generating sets. Each run is parametrised to output $w_i \leq w$ generating sets, i.e. hints for $\mathcal{A}_{i+1}$. If $\mathbf{A} \in A^{\checkmark}$, which holds with probability in $1/\mathsf{poly}(m)$ by assumption, the probability that at least one run succeeds is

$$1 - (1 - \hat{p}_i)^{w_i} \geq 1 - \exp(-\hat{p}_i \cdot w_i) \geq 1 - \exp(-\hat{p} \cdot w).$$

Since these probabilities are independent we obtain an overall probability of at least

$$\Pi_{i=2}^{z}(1 - \exp(-\hat{p}_i \cdot w)) \geq 1 - z \cdot \exp(-\hat{p} \cdot w) \geq 1 - z \cdot \exp(-n) \geq 1 - \mathsf{negl}(n)$$

for the chain $\mathcal{A}_2, \ldots, \mathcal{A}_z$ succeeding. The time and memory complexity follow from $w \in \mathsf{poly}(m)$ and Proposition 6. □

9 Bounds on λ_1 for $\Lambda_q^\perp(m,n)$

The output of our algorithm is a function of $\mu(\Lambda_q^\perp(\mathbf{A}))$ and we would like to understand the quality of this output as a function of $\lambda_1(\Lambda_q^\perp(\mathbf{A}))$. Thus, in this section we recall and prove probabilistic bounds for this magnitude.

Definition 15. *Let $(m,r) \in \mathbb{N} \times \mathbb{R}_{\geq 1}$ then denote*

- $X^m(r) = (\mathbb{Z}^m \setminus \{\mathbf{0}\}) \cap \mathsf{B}^m(r)$,
- $X_q^m(r) = (\mathbb{Z}^m \setminus q\mathbb{Z}^m) \cap \mathsf{B}^m(r)$,
- $N^m(r) = |X^m(r)|$,
- $N_q^m(r) = \left|X_q^m(r)\right|$.

Here $X^m(r)$ is the non zero integer points in a closed ball of radius r, $X_q^m(r)$ is the lifts of non zero elements of $\mathbb{Z}_q^m$ in the same ball, and $N^m(r)$ and $N_q^m(r)$ are their cardinalities. As $X_q^m(r) \subseteq X^m(r)$ we have $N_q^m(r) \leq N^m(r)$ with $X_q^m(r) = X^m(r)$ if and only if $r < q$. When $r < 1$ we have $X^m(r) = X_q^m(r) = \emptyset$ and $N_q^m(r) = N^m(r) = 0$.

We make use of two sets of bounds on $N^m(r) + 1$, i.e. the number of integer points in a ball including zero. The first is elementary. In particular, the upper bounds are upper bounds on $N^m(r)$.

Lemma 30. *Let $r \geq 1$. We have*

$$\mathsf{V}^m(r - \sqrt{m}/2) \leq N^m(r) + 1 \leq \mathsf{V}^m(r + \sqrt{m}/2).$$

Lemma 31 ([RSD24, Claim. 8.2]). *Let $1 \leq r \leq \sqrt{m}$. We have*

$$(2m/\lfloor r^2 \rfloor)^{\lfloor r^2 \rfloor} \leq N^m(r) + 1 \leq (2e^3 m/\lfloor r^2 \rfloor)^{\lfloor r^2 \rfloor}.$$

We prove a probabilistic lower bound on the first minimum of a uniform kernel lattice in terms of $N^m(r)$, that is we study $\lambda_1(\Lambda_q^\perp(m,n))$ as follows.

Lemma 32. *Let (m,q) be parametrised by n with $m \geq n$ and q prime. Let $1 \leq r < q$ then*

$$\Pr\left[\lambda_1(\Lambda_q^\perp(m,n)) \leq r\right] \leq N^m(r)/q^n.$$

Proof. Let $\mathbf{x} \in \mathbb{Z}^m \setminus q\mathbb{Z}^m$ then $\mathsf{U}(\mathbb{Z}_q^{n\times m}) \cdot \mathbf{x} \sim \mathsf{U}(\mathbb{Z}_q^n)$. Then

$$\Pr\left[\mathbf{x} \in \Lambda_q^\perp(m,n)\right] = \Pr\left[\mathsf{U}(\mathbb{Z}_q^n) = \mathbf{0}\right] = 1/q^n.$$

Moreover,

$$\begin{aligned}
\Pr\left[\lambda_1(\Lambda_q^\perp(m,n)) \leq r\right] &= \Pr\left[\mathbf{0} \in \mathsf{U}(\mathbb{Z}_q^{n\times m}) \cdot X^m(r)\right] \\
&= \Pr\left[\mathbf{0} \in \mathsf{U}(\mathbb{Z}_q^{n\times m}) \cdot X_q^m(r)\right] \\
&\leq N_q^m(r)/q^n \\
&= N^m(r)/q^n.
\end{aligned}$$

The first equality is by definition, the second as $r < q$, the inequality is by the union bound over $\mathbf{x} \in \mathbb{Z}^m \setminus q\mathbb{Z}^m$ and the last equality again as $r < q$. □

If $r < 1$ then the probability is zero and if $r \geq q$ then the probability is one. It remains to bound $N^m(r)$ and ensure it is sufficiently strongly dominated by q^n.

Lemma 33. *Let (m, q) be parametrised by n with $m \geq n$ and q prime. Let $\alpha > 0$ and*

$$\beta = \sqrt{m}\left(\frac{\alpha q^{n/m}}{\sqrt{2\pi e}} - \frac{1}{2}\right).$$

If $1 \leq r \leq \beta$ and $r < q$ then

$$\Pr\left[\lambda_1(\Lambda_q^\perp(m, n)) \leq r\right] \leq \alpha^m.$$

Furthermore, if $m = (n \log q)/f(n)$ for $f(n) \in [1, \log q]$ and $f(n) \geq \log(\sqrt{2\pi e}/\alpha)$ then

$$\beta \geq \frac{\alpha\sqrt{m} \cdot q^{n/m}}{2\sqrt{2\pi e}} \geq \sqrt{m}/2$$

and $[1, \beta]$ is non empty for large enough n.

Proof. By hypothesis we satisfy the hypothesis of Lemma 32 and thus the probability is at most $N^m(r)/q^n$. It suffices to show $(N^m(r))^{1/m} \leq \alpha \cdot q^{n/m}$. By Lemma 30 we have $N^m(r) \leq \mathsf{V}^m(r + \sqrt{m}/2)$, where the latter is given by $\pi^{m/2} \cdot (r + \sqrt{m}/2)^m / \Gamma(1 + m/2)$. Via $\Gamma(1 + m/2)^{1/m} \geq \sqrt{m/2e}$ and $r \leq \beta$ we have

$$(N^m(r))^{1/m} \leq \sqrt{2\pi e/m} \cdot (r + \sqrt{m}/2) \leq \alpha \cdot q^{n/m}.$$

Let $m = (n \log q)/f(n)$. If $x > 0$ then $x^{1/\log x} = 2$ therefore $q^{n/m} = q^{f(n)/\log q} = 2^{f(n)}$. Therefore $\alpha q^{n/m}/\sqrt{2\pi e} \geq 1$. □

The following corollary shows satisfying r exist in Lemma 33 for $m \in o(n \log q)$ and some $m \notin o(n \log q)$, in particular some $m \in \Omega(n \log q)$. That is, $\beta \geq \sqrt{m}/2$ so that $1 \leq r \leq \beta$ and $r < q$ has satisfying r for large enough n.

Corollary 14. *Let (m, q, f) be parametrised by n with $f(n) \in [1, \log q]$, $m = (n \log q)/f(n)$ and q prime. Let $\alpha > 0$. If $m \in o(n \log q)$ then $f(n) \in \omega(1)$. If $q \geq \sqrt{2\pi e}/\alpha$, for example any $q \in \omega(1)$, then set $f(n) = \log(\sqrt{2\pi e}/\alpha)$ and $m \in \Omega(n \log q)$.*

Finally, we give a corollary of Lemma 33 that ensures $\beta < q$ when $m \in [cn, n \log q]$ for some $c > 1$ by determining a minimum satisfying $q \in \mathsf{poly}(n)$.

Corollary 15. *Let (m, q) be parametrised by n with $m \geq n$ and q prime. Let $c > 1$ and $cn \leq m \leq n \log q$. Let α, β be as in Lemma 33. If $q \geq n^\gamma$ for $\gamma > c/(2(c-1))$ and n is sufficiently large then $\beta < q$.*

Proof. Let $C = \alpha/\sqrt{2\pi e}$ then since $\beta < Cq^{n/m}\sqrt{m}$ and $q^{1-n/m} \geq q^{(c-1)/c}$ it is enough to show $C\sqrt{m} < q^{(c-1)/c}$. By hypothesis $m \leq n \log q$ and $n \leq q^{1/\gamma}$. It is therefore enough to show for large enough n that $C\sqrt{q^{1/\gamma} \log q} < q^{(c-1)/c}$. Conclude by hypothesis on γ. □

For example, if $m = 2n$ we may take $q = n^2$, and then we have $\Pr\left[\lambda_1(\Lambda_q^\perp(2n, n)) \leq \beta\right] \leq \alpha^m$ for sufficiently large n. Provided $m \in [cn, (n \log q)/f(n)]$ for $f(n) \geq \log(\sqrt{2\pi e}/\alpha)$ and $c > 1$, and that q is large enough, the quantity β is some constant fraction of the Gaussian heuristic for almost all lattices in $\Lambda_q^\perp(m, n)$. Indeed the Gaussian heuristic is $\mathsf{vol}(\Lambda)^{1/m}/(\mathsf{V}^m(1))^{1/m}$, where $1/(\mathsf{V}^m(1))^{1/m} \sim \sqrt{m/2\pi e}$ and $\mathsf{vol}(\Lambda_q^\perp(m, n)) = q^n$ with probability at least $1 - 1/q^{m-n}$. That is, provided m grows constant factors faster than n and slower than $n \log q$ the probability we violate the Gaussian heuristic by more than a constant factor is in $\mathsf{negl}(n)$.

To make progress beyond some $m \notin o(n \log q)$ we require $r \in o(\sqrt{m})$. This is not required for this work, but may find use elsewhere. While Lemmas 30 and 31 are similar for $r \in \Theta(\sqrt{m})$, the latter is significantly tighter for $r \in o(\sqrt{m})$. We therefore simplify the form of the upper bound of Lemma 31 and then examine it with respect to q^n.

Proposition 7. *Let $m \in \mathbb{N}^+$ and $2 \le r \le \sqrt{m}$. Let $C = 4e^3$ then*

$$(2e^3 m/\lfloor r^2 \rfloor)^{\lfloor r^2 \rfloor} \le \left(\frac{Cm}{r^2}\right)^{r^2}.$$

Proof. For any r we have $r^2 \ge \lfloor r^2 \rfloor$. If further $r \ge 2$ we have $r^2/2 \le \lfloor r^2 \rfloor$. □

Lemma 34. *Let (m, q) be parametrised by n with $m \ge n$ and q prime. Let $0 < \alpha < 1$, write $m = \phi \cdot n \ln q$ and let $\beta = \frac{C\phi}{\alpha}$ for $C = 4e^3$. If $4\phi \ge \alpha$ and r is such that $2 \le r \le \sqrt{m}$, $r \le \sqrt{\frac{\alpha}{2 \log \beta} \cdot n \log q}$ and $r < q$ then*

$$\Pr\left[\lambda_1(\Lambda_q^\perp(m, n)) \le r\right] \le 2^{-(1-\alpha) \cdot n \log q}.$$

Proof. By hypothesis we satisfy the hypothesis of Lemma 32 and thus the probability is at most $N^m(r)/q^n$. Let $x = Cm/r^2 \ge C \ge e$. By Proposition 7 we have

$$(N^m(r))^{1/m} \le \left(\frac{Cm}{r^2}\right)^{r^2/m} = x^{C/x} = e^{(C \ln x)/x}.$$

If $x \ge e$ then $(\ln x)/x \in (0, 1/e]$ and is decreasing. Ensuring that $(\ln x)/x \le 1/\beta$ implies $N^m(r)/q^n \le 2^{-(1-\alpha)n \log q}$. We seek the minimum x that satisfies $(\ln x)/x \le 1/\beta$.

It is sufficient to set $x \ge -\beta \cdot W_{-1}(-1/\beta)$ where W_{-1} is the (-1)-branch of the Lambert W function. Indeed, first note

$$\frac{\ln x}{x} \le 1/\beta \iff -(x/\beta)e^{-(x/\beta)} \ge -1/\beta.$$

By hypothesis $\beta \ge e^3$ so $-1/\beta \in [-1/e, 0)$. This implies real solutions to $we^w = -1/\beta$ are given solely by $w_0 = W_0(-1/\beta)$ and $w_{-1} = W_{-1}(-1/\beta)$. We have $w_{-1} \le w_0 < 0$ so taking $x \ge -\beta w_{-1} \ge -\beta w_0$ is sufficient. By [Cha13, Theorem 1] for all $u > 0$ we have

$$-W_{-1}(-e^{-u-1}) \le 1 + \sqrt{2u} + u$$

with $1+\sqrt{2u}+u \le 1+2u$ for $u \ge 2$. Consider the substitution $u = \ln(\beta)-1$ so that $-W_{-1}(-e^{-u-1}) = -W_{-1}(-1/\beta)$. We have $u \ge 2 \iff \beta \ge e^3 \iff 4\phi \ge \alpha$. Therefore, to satisfy $\frac{\ln x}{x} \le \frac{1}{\beta}$, it suffices to pick $x \ge 2\beta \ln \beta \ge \beta \cdot (1 + 2u) \ge -\beta \cdot W_{-1}(-1/\beta)$. Substituting $x = Cm/r^2$ and $\beta = C\phi/\alpha$ yields the upper bound of r in the hypothesis. □

Note that Lemma 34 can make statements about $\lambda_1(\Lambda_q^\perp(m, n))$ for $m \in \Theta(n \log q)$ and *some* functions $m \in \omega(n \log q)$. Indeed, if $m \in o(n \log q)$ then $\phi \in o(1)$ and no satisfying constant α can be found. Similarly, if $m \ge q^n \cdot n \ln q$ then $\log \beta \ge n \log q$ so $r < 1$ is required. However, along with Lemma 33 a non trivial statement is achieved for all reasonable m, as explained below. First, for $m \in o(n \log q)$ and some $m \in \Omega(n \log q)$ Lemma 33 provides a lower bound on $\lambda_1(\Lambda_q^\perp(m, n))$ that is a constant factor away from the Gaussian heuristic. Second, for Lemma 34, if $m \ge q^n \cdot n \ln q$ then the probability of observing a zero column in $\mathbf{A}$ is at least constant, and hence so is the probability a canonical vector is in $\Lambda_q^\perp(\mathbf{A})$. That is, Lemma 34 gives a non trivial statement for reasonable $m \in \Omega(n \log q)$.

10 Conclusion

To conclude, assuming a chain of adversaries for a given $\mathbf{A}$ we are able to sequentially call each adversary to output some polynomial number of generating sets of norm bound at most β_i. These generating sets are formed of vectors sampled from (close to) a zero-centred discrete Gaussian distribution over $\Lambda_q^\perp(\mathbf{A})$ with decreasing widths s_i. We can start this chain by sampling hints using a Hermite normal form basis of $\Lambda_q^\perp(\mathbf{A})$ and an appropriate s_1, as detailed in Corollary 7. This chain terminates whenever we can no longer select a satisfying width s per the statement of Lemma 29, in particular when $\beta_z = 2\gamma m\sqrt{2C/\pi} \cdot \lambda_m(\Lambda_q^\perp(\mathbf{A}))$ for some $C > 0$. With overwhelming probability we have $\lambda_m(\Lambda_q^\perp(\mathbf{A})) \leq 8(1 + \sqrt{m/2\pi}) \cdot q^{n/m}\sqrt{\log(6m)/\pi}$ by Corollary 9. In particular, selecting constant γ there exists a final norm bound $\beta_z \in O(q^{n/m} \cdot \sqrt{m^3 \log m})$. To satisfy the conditions of Conjecture 1, and therefore violate the content of this conjecture, let $q \in \mathsf{poly}(n)$ and $2n \leq m \in o(n \log q)$ then there exist $q \in O((n \log n)^6)$ such that $q \geq 8n\sqrt{m}\beta_z$. This is far from tight, and subsequent works [MP13] also reduce the required q. We have exhibited an algorithm, given our assumptions, in the supposedly empty class. Given the same conditions, we show in §9 that this β_z is a factor of $O(m\sqrt{\log m})$ from $\lambda_1(\Lambda_q^\perp(\mathbf{A}))$ with overwhelming probability.

Finally, we recall that we can either instantiate the chain of algorithms $\{\mathcal{A}_i\}_{i\in[z]}$ required for our chain to invalidate Conjecture 1 or there exits at least one $j \in [z]$ such that $\mathcal{A}_j$ does not exist, i.e. there exists β_j for hints following $D_{\Lambda_q^\perp(\mathbf{A}),s_j}$ where H-ISIS is hard.

Acknowledgements

We thank Damien Stehlé and Julien Devevey for helpful discussions on rejection sampling. We thank Cyrill Graf for finding mistakes in an earlier version of this work. We thank Matthew Jenssen for discussions on packing points into balls. We thank previous anonymous reviewers for insightful comments that improved this paper. Martin R. Albrecht's work was supported by EPSRC grant EP/Y02432X/1. Russell W. F. Lai's work was supported by the Research Council of Finland grant 358951 and the Future Makers Program of the Technology Industries of Finland Centennial Foundation via the project SPICED.

References

ABF+20. Martin R. Albrecht, Shi Bai, Pierre-Alain Fouque, Paul Kirchner, Damien Stehlé, and Weiqiang Wen. Faster enumeration-based lattice reduction: Root hermite factor $k^{1/(2k)}$ time $k^{k/8+o(k)}$. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 186–212. Springer, Cham, August 2020.

ABLR21. Martin R. Albrecht, Shi Bai, Jianwei Li, and Joe Rowell. Lattice reduction with approximate enumeration oracles - practical algorithms and concrete performance. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 732–759, Virtual Event, August 2021. Springer, Cham.

ACKS21. Divesh Aggarwal, Yanlin Chen, Rajendra Kumar, and Yixin Shen. Improved (provable) algorithms for the shortest vector problem via bounded distance decoding. In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021*, volume 187 of *LIPIcs*, pages 4:1–4:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

ACL+22. Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri Aravinda Krishnan Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 102–132. Springer, Cham, August 2022.

AD21. Martin R. Albrecht and Léo Ducas. Lattice attacks on NTRU and LWE: A history of refinements. In Joppe W. Bos and Martijn Stam, editors, *Computational Cryptography: Algorithmic Aspects of Cryptology*, London Mathematical Society Lecture Note Series, pages 15–40. Cambridge University Press, 2021.

ADH+19. Martin R. Albrecht, Léo Ducas, Gottfried Herold, Elena Kirshanova, Eamonn W. Postlethwaite, and Marc Stevens. The general sieve kernel and new records in lattice reduction. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 717–746. Springer, Cham, May 2019.

ADRS15. Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in 2^n time using discrete Gaussian sampling: Extended abstract. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 733–742. ACM Press, June 2015.

ADS15. Divesh Aggarwal, Daniel Dadush, and Noah Stephens-Davidowitz. Solving the closest vector problem in 2^n time - the discrete Gaussian strikes again! In Venkatesan Guruswami, editor, *56th FOCS*, pages 563–582. IEEE Computer Society Press, October 2015.

AGPS20. Martin R. Albrecht, Vlad Gheorghiu, Eamonn W. Postlethwaite, and John M. Schanck. Estimating quantum speedups for lattice sieves. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 583–613. Springer, Cham, December 2020.

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

Ajt98. Miklós Ajtai. The shortest vector problem in L2 is NP-hard for randomized reductions (extended abstract). In *30th ACM STOC*, pages 10–19. ACM Press, May 1998.

AKS01. Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *33rd ACM STOC*, pages 601–610. ACM Press, July 2001.

AKSY22. Shweta Agrawal, Elena Kirshanova, Damien Stehlé, and Anshu Yadav. Practical, round-optimal lattice-based blind signatures. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 39–53. ACM Press, November 2022.

ALS21. Divesh Aggarwal, Zeyong Li, and Noah Stephens-Davidowitz. A $2^{n/2}$-time algorithm for $\sqrt{n}$-SVP and $\sqrt{n}$-Hermite SVP, and an improved time-approximation tradeoff for (H)SVP. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 467–497. Springer, Cham, October 2021.

ANS18. Yoshinori Aono, Phong Q. Nguyen, and Yixin Shen. Quantum lattice enumeration and tweaking discrete pruning. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 405–434. Springer, Cham, December 2018.

ANSS18. Yoshinori Aono, Phong Q. Nguyen, Takenobu Seito, and Junji Shikata. Lower bounds on lattice enumeration with extreme pruning. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 608–637. Springer, Cham, August 2018.

AS17. Jacob Alperin-Sheriff. NIST's PQC Standardization: Suggested avenues for lattice-based research. Talk, slides available at http://crypto-events.di.ens.fr/LATCA/program/alperin-sheriff.pdf, May 2017.

Bab86. László Babai. On lovász' lattice reduction and the nearest lattice point problem. *Comb.*, 6(1):1–13, 1986.

Ban93. W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–636, 1993.

BBC+20. Daniel J. Bernstein, Billy Bob Brumley, Ming-Shing Chen, Chitchanok Chuengsatiansup, Tanja Lange, Adrian Marotzke, Bo-Yuan Peng, Nicola Tuveri, Christine van Vredendaal, and Bo-Yin Yang. NTRU Prime. Technical report, National Institute of Standards and Technology, 2020. available at https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions.

BBTV24. Nina Bindel, Xavier Bonnetain, Marcel Tiepelt, and Fernando Virdia. Quantum lattice enumeration in limited depth. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VI*, volume 14925 of *LNCS*, pages 72–106. Springer, Cham, August 2024.

BDGL16. Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In Robert Krauthgamer, editor, *27th SODA*, pages 10–24. ACM-SIAM, January 2016.

Ber16. Daniel Bernstein. Re: Inaccurate security claims in NTRUprime. Cryptanalytic algorithms mailing list, May 2016. https://groups.google.com/g/cryptanalytic-algorithms/c/BoSRL0uHIjM/m/eB4G-dscCAAJ.

BF11. Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 1–16. Springer, Berlin, Heidelberg, March 2011.

BGJ15. Anja Becker, Nicolas Gama, and Antoine Joux. Speeding-up lattice sieving without increasing the memory, using sub-quadratic nearest neighbor search. Cryptology ePrint Archive, Report 2015/522, 2015.

BL26. Katharina Boudgoust and Oleksandra Lapiha. Leftover hash lemma(s) over cyclotomic rings. In Joan Daemen and Emmanuel Thomé, editors, *Advances in Cryptology – EUROCRYPT 2026*, pages 62–91, Cham, 2026. Springer Nature Switzerland.

BLNS23. Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Alessandro Sorniotti. A framework for practical anonymous credentials from lattices. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 384–417. Springer, Cham, August 2023.

BLP+13a. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 575–584. ACM Press, June 2013.

BLP+13b. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. https://arxiv.org/abs/1306.0281, 2013.

BLS16. Shi Bai, Thijs Laarhoven, and Damien Stehlé. Tuple lattice sieving. *LMS Journal of Computation and Mathematics*, 19(A):146–162, 2016.

Cha13. Ioannis Chatzigeorgiou. Bounds on the lambert function and their application to the outage analysis of user cooperation. *IEEE Communications Letters*, 17(8):1505–1508, 2013.

Che13. Yuanmi Chen. *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. PhD thesis, Paris 7, 2013.

CL15. Jung Hee Cheon and Changmin Lee. Approximate algorithms on lattices with small determinant. Cryptology ePrint Archive, Report 2015/461, 2015.

CL21. André Chailloux and Johanna Loyer. Lattice sieving via quantum random walks. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 63–91. Springer, Cham, December 2021.

CLW25. Valerio Cini, Russell W. F. Lai, and Ivy K. Y. Woo. Lattice-based obfuscation from NTRU and equivocal LWE. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 39–72. Springer, Cham, August 2025.

CN97. Jin-yi Cai and Ajay Nerurkar. An improved worst-case to average-case connection for lattice problems. In *38th FOCS*, pages 468–477. IEEE Computer Society Press, October 1997.

Das16. Shagnik Das. A brief note on estimates of binomial coefficients. http://discretemath.imp.fu-berlin.de/DMI-2016/notes/binomials.pdf, 2016. Accessed: 2024-04-03.

DFPS22. Julien Devevey, Omar Fawzi, Alain Passelègue, and Damien Stehlé. On rejection sampling in Lyubashevsky's signature scheme. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 34–64. Springer, Cham, December 2022.

DKLW25. Adrien Dubois, Michael Klooß, Russell W. F. Lai, and Ivy K. Y. Woo. Lattice-based proof-friendly signatures from vanishing short integer solutions. In Tibor Jager and Jiaxin Pan, editors, *PKC 2025, Part I*, volume 15674 of *LNCS*, pages 452–486. Springer, Cham, May 2025.

Duc18. Léo Ducas. Shortest vector from lattice sieving: A few dimensions for free. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 125–145. Springer, Cham, April / May 2018.

Duc22. Léo Ducas. Estimating the hidden overheads in the BDGL lattice sieving algorithm. In Jung Hee Cheon and Thomas Johansson, editors, *PQCrypto 2022*, pages 480–497. Springer, Cham, September 2022.

DvW22. Léo Ducas and Wessel P. J. van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 643–673. Springer, Cham, May / June 2022.

FMN24. Giacomo Fenzi, Hossein Moghaddas, and Ngoc Khanh Nguyen. Lattice-based polynomial commitments: Towards asymptotic and concrete efficiency. *Journal of Cryptology*, 37(3):31, July 2024.

FP85. U. Fincke and M. Pohst. Improved methods for calculating vectors of short length in a lattice, including a complexity analysis. *Mathematics of Computation*, 44(170):463–463, May 1985.

GMPW20. Nicholas Genise, Daniele Micciancio, Chris Peikert, and Michael Walter. Improved discrete gaussian and subgaussian analysis for lattice cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 623–651. Springer, Cham, May 2020.

GNR10. Nicolas Gama, Phong Q. Nguyen, and Oded Regev. Lattice enumeration using extreme pruning. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 257–278. Springer, Berlin, Heidelberg, May / June 2010.

GPV07. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. Cryptology ePrint Archive, Report 2007/432, 2007.

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.

HJL25. Yao-Ching Hsieh, Aayush Jain, and Huijia Lin. Lattice-based post-quantum iO from circular security with random opening assumption. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 3–38. Springer, Cham, August 2025.

HK17. Gottfried Herold and Elena Kirshanova. Improved algorithms for the approximate k-list problem in euclidean norm. In Serge Fehr, editor, *PKC 2017, Part I*, volume 10174 of *LNCS*, pages 16–40. Springer, Berlin, Heidelberg, March 2017.

HPS11. Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Algorithms for the shortest and closest lattice vector problems. In Yeow Meng Chee, Zhenbo Guo, San Ling, Fengjing Shao, Yuansheng Tang, Huaxiong Wang, and Chaoping Xing, editors, *Coding and Cryptology - Third International Workshop, IWCC 2011*, volume 6639 of *Lecture Notes in Computer Science*, pages 159–190. Springer, 2011.

HR12. Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory of Computing*, 8(1):513–531, 2012. Preliminary version in *Proceedings of STOC '07*.

HR14. Ishay Haviv and Oded Regev. On the lattice isomorphism problem. In Chandra Chekuri, editor, *25th SODA*, pages 391–404. ACM-SIAM, January 2014.

HS07. Guillaume Hanrot and Damien Stehlé. Improved analysis of kannan's shortest lattice vector algorithm. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 170–186. Springer, Berlin, Heidelberg, August 2007.

Jaq24. Samuel Jaques. Memory adds no cost to lattice sieving for computers in 3 or more spatial dimensions. *CiC*, 1(3):6, 2024.

Kan83. Ravi Kannan. Improved algorithms for integer programming and related lattice problems. In *15th ACM STOC*, pages 193–206. ACM Press, April 1983.

Kho05. Subhash Khot. Hardness of approximating the shortest vector problem in lattices. *Journal of the ACM*, 52(5):789–808, 2005. Preliminary version in *Proceedings of FOCS '04*.

KL21. Elena Kirshanova and Thijs Laarhoven. Lower bounds on lattice sieving and information set decoding. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 791–820, Virtual Event, August 2021. Springer, Cham.

Kle00. Philip N. Klein. Finding the closest lattice vector when it's unusually close. In David B. Shmoys, editor, *11th SODA*, pages 937–941. ACM-SIAM, January 2000.

Laa15a. Thijs Laarhoven. *Search problems in cryptography: From fingerprinting to lattice sieving*. PhD thesis, Eindhoven University of Technology, 2015.

Laa15b. Thijs Laarhoven. Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 3–22. Springer, Berlin, Heidelberg, August 2015.

LMvdP13. Thijs Laarhoven, Michele Mosca, and Joop van de Pol. Solving the shortest vector problem in lattices faster using quantum search. In Philippe Gaborit, editor, *PQCrypto 2013*, pages 83–101. Springer, Berlin, Heidelberg, June 2013.

Loè77. M. Loève. *Probability Theory I*. Springer New York, NY, 4 edition, 1977.

LP21. David Lazard and Chris Peikert. Deterministic Falcon Signatures. https://github.com/algorand/falcon/blob/main/falcon-det.pdf, November 2021. accessed 25 Apr 2023.

LPSS14. San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld. Hardness of k-LWE and applications in traitor tracing. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 315–334. Springer, Berlin, Heidelberg, August 2014.

LV20. Alex Lombardi and Vinod Vaikuntanathan. Fiat-Shamir for repeated squaring with applications to PPAD-hardness and VDFs. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 632–651. Springer, Cham, August 2020.

Lyu11. Vadim Lyubashevsky. Lattice signatures without trapdoors. Cryptology ePrint Archive, Report 2011/537, 2011.

MG02. Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, March 2002.

Mic01. Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM Journal on Computing*, 30(6):2008–2035, 2001. Preliminary version in *Proceedings of FOCS '98*.

Mic07. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complex.*, 16(4):365–411, 2007.

Mic11. Daniele Micciancio. *Foundations of Security Analysis and Design VI: FOSAD Tutorial Lectures*, chapter The Geometry of Lattice Cryptography, pages 185–210. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

Mic12. Daniele Micciancio. Inapproximability of the Shortest Vector Problem: Toward a deterministic reduction. *Theory of Computing*, 8(22):487–512, 2012.

MP13. Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 21–39. Springer, Berlin, Heidelberg, August 2013.

MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.

MR09. Daniele Micciancio and Oded Regev. Lattice-based cryptography. In Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen, editors, *Post-Quantum Cryptography*, pages 147–191. Springer, Heidelberg, Berlin, Heidelberg, New York, 2009.

MS65. T. S. Motzkin and E. G. Straus. Maxima for graphs and a new proof of a theorem of turán. *Canadian Journal of Mathematics*, 17:533–540, 1965.

MU05. Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.

MV10a. Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations. In Leonard J. Schulman, editor, *42nd ACM STOC*, pages 351–358. ACM Press, June 2010.

MV10b. Daniele Micciancio and Panagiotis Voulgaris. Faster exponential time algorithms for the shortest vector problem. In Moses Charika, editor, *21st SODA*, pages 1468–1480. ACM-SIAM, January 2010.

MW15. Daniele Micciancio and Michael Walter. Fast lattice point enumeration with minimal overhead. In Piotr Indyk, editor, *26th SODA*, pages 276–294. ACM-SIAM, January 2015.

NIS23. NIST. FAQ on Kyber512. https://csrc.nist.gov/csrc/media/Projects/post-quantum-cryptography/documents/faq/Kyber-512-FAQ.pdf, December 2023.

NV08. Phong Q. Nguyen and Thomas Vidick. Sieve algorithms for the shortest vector problem are practical. *J. Mathematical Cryptology*, 2(2):181–207, 2008.

OPF+22. Mike Ounsworth, Chris Peikert, Scott Fluhrer, Uri Blumenthal, Jeff Burdges, Panos Kampanakis, Phillip Hallam-Baker, Derek Atkins, Brent Kimberley, Tony Arcieri, Greg Maxwell, and Samuel Lavery. Design rationale for keyed message digests in SPHINCS+, Dilithium, FALCON? https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/cIsc6tUY9Rw/m/E0fPG7QkAQAJ, September 2022. accessed 25 Apr 2023.

Pei08. Chris Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, May 2008.

Pei09. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009.

PFH+22. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022.

Pho81. Michael Phost. On the computation of lattice vectors of minimal length, successive minima and reduced bases with applications. *SIGSAM Bulletin*, 15:37–44, 1981.

PR05. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. *Electron. Colloquium Comput. Complex.*, TR05-158, 2005.

PS00. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

RS92. V.K. Rohatgi and Gábor J. Székely. An inverse markov-chebyshev inequality. *Periodica Polytechnica Civil Engineering*, 36(4):455–458, 1992.

RSD24. Oded Regev and Noah Stephens-Davidowitz. A reverse Minkowski theorem. *Annals of Mathematics*, 199(1):1–49, 2024.

Sch24. John Schanck. An Update on Lattice Cryptanalysis vol. 2. Invited talk delivered at RWPQC'24, March 2024. https://na.eventscloud.com/website/65452/presentations-and-video-/.

SE94. C.P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.*, 66:181–199, 1994.

Wee22. Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Cham, May / June 2022.

Wee24. Hoeteck Wee. Circuit ABE with $\mathsf{poly}(\mathsf{depth}, \lambda)$-sized ciphertexts and keys from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 178–209. Springer, Cham, August 2024.

WW23. Hoeteck Wee and David J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 385–416. Springer, Cham, April 2023.

A Reduction from One-More-ISIS

We argue for the utility of H-ISIS via a reduction from omISIS, i.e. if H-ISIS is easy (with a specific Dist) then omISIS is easy. Our proof uses a technical lemma that we prove first.

A.1 Multi instance repetition gadget

Below, we rely on being able to query an adversary on many problem instances to obtain valid solutions to several distinct problem instances. This is in contrast to querying an adversary many times on the same problem instance to extract several valid solutions to the same problem instance. The lemma below establishes the feasibility of this approach.

Lemma 35. *Let* $\mathsf{Prob} = (\mathcal{P}, \mathcal{D})$ *be a problem,* $k \in \mathbb{N}^+$ *and* $\mathcal{A}$ *be a* (τ, μ, ε) *adversary against* Prob*. Let* O *be an oracle which returns problem instances sampled according to* Prob*. If* $\varepsilon > 0$ *then there exists an algorithm* $\mathcal{E}^{\mathsf{O}}$ *that either returns* $(p_i, x_i)_{i=1}^k$ *or aborts. It runs in time* $\tau' = 4k\tau/\varepsilon + o(\tau)$*, memory* $\mu' = \mu + o(\mu)$ *and all* x_i *are correct outputs for* p_i *with probability* $\varepsilon' \geq 1/2$*.*

Proof. Algorithm $\mathcal{E}^{\mathsf{O}}$ repeats the following $N \geq k$ times, for $N = 4k/\varepsilon$. Receive $p \leftarrow \mathsf{O}$ and $x \leftarrow \mathcal{A}(p)$. When $\mathcal{A}$ requires randomness, $\mathcal{E}$ provides its random tape. If x is a correct solution to p then $\mathcal{E}$ stores (p, x) in a set GOOD. If $|\mathsf{GOOD}| = k$ then $\mathcal{E}$ returns it. If this does not occur within N repetitions, $\mathcal{E}$ aborts. The probability $\mathcal{E}$ returns GOOD is $\Pr[X \geq k]$ for $X \sim \mathrm{Bin}(N, \varepsilon)$. We appeal to Corollary 2 with $p = \varepsilon$ and $c = 1 - 2^{-k} \geq 1/2$. □

We apply this argument in Lemma 36. Problem instances are provided by a 'syn' oracle (recall Fig. 3). In general applications of Lemma 35 one must ensure that generating problem instances is efficient compared to solving them.

A.2 Reduction

Lemma 36 (omISIS to H-ISIS). *Let*

- $\tau_{\mathcal{A}}, \mu_{\mathcal{A}}, \varepsilon_{\mathcal{A}}, \mathsf{params}_{\mathcal{A}} = (k, m, q, \beta, \mathsf{Dist}), s$ *be parametrised by* n*,*
- Dist *map* $\mathbf{A}$ *to* $D_{\mathbb{Z}^m, s}\,[\mathbf{A}]^k$*,*
- $\mathsf{params}_{\mathcal{B}} = (m, q, \beta, s)$*.*

If there exists a $(\tau_{\mathcal{A}}, \mu_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$ *adversary* $\mathcal{A}$ *against* $\mathsf{H\text{-}ISIS}_{\mathsf{params}_{\mathcal{A}}}$ *then there exists a* $(\tau_{\mathcal{B}}, \mu_{\mathcal{B}}, \varepsilon_{\mathcal{B}})$ *adversary* $\mathcal{B}$ *against* $\mathsf{omISIS}_{\mathsf{params}_{\mathcal{B}}}$*. Adversary* $\mathcal{B}$ *makes* k *queries to its* pre *oracle and at most* $8(k+1)/\varepsilon_{\mathcal{A}}$ *queries to its* syn *oracle with* $\tau_{\mathcal{B}} \approx 8(k+1) \cdot \tau_{\mathcal{A}}/\varepsilon_{\mathcal{A}}$*,* $\mu_{\mathcal{B}} \approx \mu_{\mathcal{A}}$ *and* $\varepsilon_{\mathcal{B}} \geq \delta \cdot \varepsilon_{\mathcal{A}}/4$ *for* δ *given in the proof.*

Proof. We define $\mathcal{B} = \mathcal{B}^{\mathsf{syn},\mathsf{pre}}(\mathbf{A})$. Let $\mathsf{GOOD} = \emptyset$. First $\mathcal{B}$ makes k calls to $\mathsf{pre}(\mathbf{0})$ and receives $\mathbf{u}_1, \ldots, \mathbf{u}_k$. Then $\mathcal{B}$ repeats the following up to $r = 8 \cdot (k+1)/\varepsilon_{\mathcal{A}}$ times. It calls syn to receive $\mathbf{t}$ then receives $\mathbf{y} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{t}, \mathbf{u}_1, \ldots, \mathbf{u}_k)$. When $\mathcal{A}$ requires randomness, $\mathcal{B}$ provides its own random tape. If $\mathbf{y} \in P_{\beta,\mathbf{A},\mathbf{t}}$ then $\mathcal{B}$ adds $(\mathbf{y}, \mathbf{t})$ to GOOD. If $|\mathsf{GOOD}| = k+1$ then $\mathcal{B}$ returns GOOD.

A pair $(\mathbf{y}, \mathbf{t})$ such that $\mathbf{y} \in P_{\beta,\mathbf{A},\mathbf{t}}$ can already be an element of GOOD only if $\mathbf{t}$ repeats. Let $N = q^n$. If $r \in \{0, 1, \ldots, N\}$ then there is no repeated $\mathbf{t}$ with probability $\delta = \binom{N}{r} \cdot \frac{r!}{N^r}$. If $r > N$ then $\delta = 0$.

Let $\mathsf{aux} = (\mathbf{u}_1, \ldots, \mathbf{u}_k)$. We determine the probability that $\mathcal{A}$ with fixed $(\mathbf{A}, \mathsf{aux})$, as provided by $\mathcal{B}$, succeeds with probability at least $\varepsilon/2$ on uniform $\mathbf{t}$, i.e. we consider the success probability $\varepsilon_{\mathbf{A},\mathsf{aux}}$ of $\mathcal{A}_{\mathbf{A},\mathsf{aux}}(\cdot) = \mathcal{A}(\mathbf{A}, \cdot, \mathsf{aux})$. By an inverse Markov inequality (Proposition 1) with X the random variable that returns $\varepsilon_{\mathbf{A},\mathsf{aux}}$ and $(a, c) = (\varepsilon_{\mathcal{A}}/2, 1)$ we have $\Pr[\varepsilon_{\mathbf{A},\mathsf{aux}} \geq \varepsilon_{\mathcal{A}}/2] \geq \varepsilon_{\mathcal{A}}/2$. We

search for $k+1$ repetitions assuming success probability $\varepsilon_{\mathcal{A}}/2$ via Lemma 35. Note that syn is the problem generating oracle O.

We avoid repeated $\mathbf{t}$ with probability at least δ and we avoid $(\mathbf{A}, \mathsf{aux})$ such that $\varepsilon_{\mathbf{A},\mathsf{aux}} < \varepsilon_{\mathcal{A}}/2$ with probability at least $\varepsilon_{\mathcal{A}}/2$. These events are independent, and hence both happen with probability at least $\delta \cdot \varepsilon_{\mathcal{A}}/2$. Given this, after $8 \cdot (k+1)/\varepsilon_{\mathcal{A}}$ calls to $\mathcal{A}$, $\mathcal{B}$ receives $k+1$ correct responses with probability at least $1/2$. Therefore, $\mathcal{B}$ outputs GOOD with probability $\varepsilon_{\mathcal{B}} \geq \delta \cdot \varepsilon_{\mathcal{A}}/4$. □

The value of δ is a function of (N, r) which are functions of n via $(q, k, \varepsilon_{\mathcal{A}})$. If $r \leq N^{\gamma}$, or equivalently $8 \cdot (k+1) \leq q^{\gamma n} \cdot \varepsilon_{\mathcal{A}}$, for $\gamma \in [0, 1/2)$ then by Corollary 1 we have $\delta \sim 1$.