Hardness of SIS and LWE with Small Parameters

Daniele Micciancio[1⋆] and Chris Peikert[2⋆⋆]

[1] University of California, San Diego
[2] School of Computer Science, Georgia Institute of Technology

Abstract. The Short Integer Solution (SIS) and Learning With Errors (LWE) problems are the foundations for countless applications in lattice-based cryptography, and are provably as hard as approximate lattice problems in the worst case. An important question from both a practical and theoretical perspective is how small their parameters can be made, while preserving their hardness.

We prove two main results on SIS and LWE with small parameters. For SIS, we show that the problem retains its hardness for moduli $q \geq \beta \cdot n^{\delta}$ for any constant $\delta > 0$, where β is the bound on the Euclidean norm of the solution. This improves upon prior results which required $q > \beta \cdot \sqrt{n \log n}$, and is close to optimal since the problem is trivially easy for $q \leq \beta$. For LWE, we show that it remains hard even when the errors are small (e.g., uniformly random from $\{0, 1\}$), provided that the number of samples is small enough (e.g., linear in the dimension n of the LWE secret). Prior results required the errors to have magnitude at least $\sqrt{n}$ and to come from a Gaussian-like distribution.

Keywords: Lattice cryptography, Computational hardness, SIS, LWE

1 Introduction

In modern lattice-based cryptography, two average-case computational problems serve as the foundation of almost all cryptographic schemes: Short Integer Solution (SIS), and Learning With Errors (LWE). The SIS problem dates back to Ajtai's pioneering work [1], and is defined as follows. Let n and q be integers, where n is the primary security parameter and usually $q = \text{poly}(n)$, and let $\beta > 0$. Given a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for some $m = \text{poly}(n)$, the goal is to find a nonzero integer vector $\boldsymbol{z} \in \mathbb{Z}^m$ such that $\mathbf{A}\boldsymbol{z} = \mathbf{0} \bmod q$ and $\|\boldsymbol{z}\| \leq \beta$ (where $\|\cdot\|$ denotes Euclidean norm). Observe that β should be set large enough to ensure that a solution exists (e.g., $\beta > \sqrt{n \log q}$ suffices), but

⋆ Supported by the TYPED project under the DARPA PROCEED program and the National Science Foundation under grant CNS-1117936.

⋆⋆ Supported by the National Science Foundation under CAREER Award CCF-1054495, by DARPA under agreement number FA8750-11-C-0096, and by the Alfred P. Sloan Foundation.

that $\beta \geq q$ makes the problem trivially easy to solve. Ajtai showed that for appropriate parameters, SIS enjoys a remarkable worst-case/average-case hardness property: solving it *on the average* (with any noticeable probability) is at least as hard as approximating several lattice problems on n-dimensional lattices *in the worst case*, to within $\mathrm{poly}(n)$ factors.

The LWE problem was introduced in the celebrated work of Regev [26], and has the same parameters n and q, along with a "noise rate" $\alpha \in (0,1)$. The problem (in its search form) is to find a secret vector $\boldsymbol{s} \in \mathbb{Z}_q^n$, given a "noisy" random linear system $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\boldsymbol{b} = \mathbf{A}^T\boldsymbol{s} + \boldsymbol{e} \bmod q$, where $\mathbf{A}$ is uniformly random and the entries of $\boldsymbol{e}$ are i.i.d. from a Gaussian-like distribution with standard deviation roughly αq. Regev showed that as long as $\alpha q \geq 2\sqrt{n}$, solving LWE on the average (with noticeable probability) is at least as hard as approximating lattice problems in the worst case to within $\tilde{O}(n/\alpha)$ factors using a *quantum* algorithm. Subsequently, similar results under *classical* (i.e., non quantum) hardness assumptions were proved in [23,7].

A significant line of research has been devoted to improving the tightness of worst-case/average-case connections for lattice problems. For SIS, a series of works [1,8,16,21,13] gave progressively better parameters that guarantee hardness, and smaller approximation factors for the underlying lattice problems. The state of the art (from [13], building upon techniques introduced in [21]) shows that for $q \geq \beta \cdot \omega(\sqrt{n \log n})$, finding a SIS solution with norm bounded by β is as hard as approximating worst-case lattice problems to within $\tilde{O}(\beta\sqrt{n})$ factors. (The parameter m does not play any significant role in the hardness results, and can be any polynomial in n.) For LWE, Regev's initial result remains the tightest, and the requirement that $q \geq \sqrt{n}/\alpha$ (i.e., that the errors have magnitude at least $\sqrt{n}$) is in some sense optimal: a clever algorithm due to Arora and Ge [2] solves LWE in time $2^{\tilde{O}(\alpha q)^2}$, so a proof of hardness for substantially smaller errors would imply a subexponential time (quantum) algorithm for approximate lattice problems, which would be a major breakthrough. Interestingly, the current modulus bound for LWE is in some sense better than the one for SIS by a $\tilde{\Omega}(\sqrt{n})$ factor: there are applications of LWE for $1/\alpha = \tilde{O}(1)$ and hence $q = \tilde{O}(\sqrt{n})$, whereas SIS is only useful for $\beta \geq \sqrt{n}$, and therefore requires $q \geq n$ according to the state-of-the-art reductions.

Further investigating the smallest parameters for which SIS and LWE remain provably hard is important from both a practical and theoretical perspective. On the practical side, improvements may lead to smaller cryptographic keys without compromising the theoretical security guarantees, or may provide greater confidence in more practical parameter settings that so far lack provable hardness. Also, proving the hardness of LWE for non-Gaussian error distributions (e.g., uniform over a small set) makes applications easier to implement. Theoretically, improvements may eventually shed light on related problems like Learning Parity with Noise (LPN), which can be seen as a special case of LWE for modulus $q = 2$, and which is widely used in coding-based cryptography, but which has no known proof of hardness based on worst-case complexity assumptions.

1.1 Our Results

We prove two complementary results on the hardness of SIS and LWE with small parameters. For SIS, we show that the problem retains its hardness for moduli q nearly equal to the solution bound β. For LWE, we show that it remains hard even when the errors are small (e.g., uniformly random from $\{0, 1\}$), provided that the number m of noisy equations is small enough. This qualification is necessary in light of the Arora-Ge attack [2], which for large enough m can solve LWE with binary errors in polynomial time. Details follow.

SIS with small modulus. Our first theorem says that SIS retains its hardness with a modulus as small as $q \geq \beta \cdot n^\delta$, for any $\delta > 0$. Recall that the best previous reduction [13] required $q \geq \beta \cdot \omega(\sqrt{n \log n})$, and that SIS becomes trivially easy for $q \leq \beta$, so the q obtained by our proof is close to optimal. It also essentially closes the gap between LWE and SIS, in terms of how small a useful modulus can be. More precisely, the following is a special case of our main SIS hardness theorem; see Section 2 for full details.

Theorem 1 (Corollary of Theorem 4). *Let n and $m = \mathrm{poly}(n)$ be integers, let $\beta \geq \beta_\infty \geq 1$ be reals, let $Z = \{\boldsymbol{z} \in \mathbb{Z}^m : \|\boldsymbol{z}\|_2 \leq \beta$ and $\|\boldsymbol{z}\|_\infty \leq \beta_\infty\}$, and let $q \geq \beta \cdot n^\delta$ for some constant $\delta > 0$. Then solving (on the average, with non-negligible probability)* SIS *with parameters n, m, q and solution set $Z \setminus \{\mathbf{0}\}$ is at least as hard as approximating lattice problems in the worst case on n-dimensional lattices to within $\gamma = \max\{1, \beta \cdot \beta_\infty / q\} \cdot \tilde{O}(\beta\sqrt{n})$ factors.*

Of course, the ℓ_∞ bound on the SIS solutions can be easily removed simply setting $\beta_\infty = \beta$, so that $\|\boldsymbol{z}\|_\infty \leq \|\boldsymbol{z}\|_2 \leq \beta$ automatically holds true. We include an explicit ℓ_∞ bound $\beta_\infty \leq \beta$ in order to obtain more precise hardness results, based on potentially smaller worst-case approximation factors γ. We point out that the bound β_∞ and the associated extra term $\max\{1, \beta \cdot \beta_\infty / q\}$ in the worst-case approximation factor is not present in previous results. Notice that this term can be as small as 1 (if we take $q \geq \beta \cdot \beta_\infty$, and in particular if $\beta_\infty \leq n^\delta$), and as large as β / n^δ (if $\beta_\infty = \beta$). This may be seen as the first theoretical evidence that, at least when using a small modulus q, restricting the ℓ_∞ norm of the solutions may make the SIS problem qualitatively harder than just restricting the ℓ_2 norm. There is already significant empirical evidence for this belief: the most practically efficient attacks on SIS, which use lattice basis reduction (e.g., [12,9]), only find solutions with bounded ℓ_2 norm, whereas combinatorial attacks such as [5,27] (see also [22]) or theoretical lattice attacks [10] that can guarantee an ℓ_∞ bound are much more costly in practice, and also require exponential space. Finally, we mention that setting $\beta_\infty \ll \beta$ is very natural in the usual formulations of one-way and collision-resistant hash functions based on SIS, where collisions correspond (for example) to vectors in $\{-1, 0, 1\}^m$, and therefore have ℓ_∞ bound $\beta_\infty = 1$, but ℓ_2 bound $\beta = \sqrt{m}$. Similar gaps between β_∞ and β can easily be enforced in other applications, e.g., digital signatures [13].

LWE with small errors. In the case of LWE, we prove a general theorem offering a trade-off among several different parameters, including the size of the errors, the dimension and number of samples in the LWE problem, and the dimension of the underlying worst-case lattice problems. Here we mention just one instantiation for the case of prime modulus and uniformly distributed *binary* (i.e., 0-1) errors, and refer the reader to Section 3 and Theorem 6 for the more general statement and a discussion of the parameters.

Theorem 2 (Corollary of Theorem 6). *For any integers n and $m = n \cdot (1 + \Omega(1/\log n))$, and all sufficiently large polynomially bounded (prime) moduli $q \geq n^{O(1)}$,*[3] *solving* LWE *with parameters n, m, q and independent uniformly random binary errors (i.e., in $\{0,1\}$) is at least as hard as approximating lattice problems in the worst case on $\Theta(n/\log n)$-dimensional lattices within a factor $\gamma = \tilde{O}(\sqrt{n} \cdot q)$.*

We remark that our results (see Theorem 6 and discussion following it) apply to many other settings, including error vectors $\boldsymbol{e} \in X$ chosen from any (sufficiently large) subset $X \subseteq \{0,1\}^m$ of binary strings, as well as error vectors with larger entries. Interestingly, our hardness result for LWE with very small errors relies on the worst-case hardness of lattice problems in dimension $n' = O(n/\log n)$, which is smaller than (but still quasi-linear in) the dimension n of the LWE problem; however, this is needed only when considering very small error vectors. Theorem 6 also shows that if $\boldsymbol{e}$ is chosen uniformly at random with entries bounded by n^ϵ (which is still much smaller than $\sqrt{n}$), then the dimension of the underlying worst-case lattice problems (and the number $m - n$ of extra samples, beyond the LWE dimension n) can be linear in n.

The restriction that the number of LWE samples $m = O(n)$ be linear in the dimension of the secret can also be relaxed slightly. But some restriction is necessary, because LWE with small errors can be solved in polynomial time when given an arbitrarily large polynomial number of samples. We focus on linear $m = O(n)$ because this is enough for most (but not all) applications in lattice cryptography, including identity-based encryption and fully homomorphic encryption, when the parameters are set appropriately. (The one exception that we know of is the security proof for pseudorandom functions [3].)

We remark that state-of-the-art reductions from worst-case lattice problems [21,13,26,7] are not tight enough to provide useful estimates on the concrete security of lattice cryptography, and they are best interpreted as qualitative results showing that there is no structural flaw in cryptographic constructions/instantiations. Still, these reductions are very valuable because even small

[3] Making the asymptotic notation explicit, the theorem asserts that for any constant $c_1 > 0$ there is a constant $c_2 > 0$ such that if $m = n \cdot (1 + c_1/\log n)$ and $q \geq n^{c_2}$, then LWE with binary errors is hard. Notice that this dependency of the modulus q on the number of samples m is necessary for the theorem to be nontrivial. In fact, the LWE function with binary errors maps $\log q^n + m$ input bits to $\log q^m$ output bits. When $m = n \cdot (1 + c_1/\log n)$ and $q = n^{c_2}$ the LWE function stretches the input by $\log q^m - (\log q^n + m) = (c_1 c_2 - 1 - o(1))n$ bits, and for the theorem to be useful/nontrivial (and give, e.g., a pseudorandom generator) one needs $c_1 c_2 > 1$.

changes in parameters can easily lead to new avenues of attack, like the polynomial time algorithm of [2] to LWE with binary errors. Likewise, our work also provides results that are primarily qualitative, showing that SIS and LWE are asymptotically secure even when $q \approx \sqrt{n}$ and the errors are small (provided the number of samples is suitably restricted). The evaluation of the concrete level of security/efficiency offered by SIS and LWE for specific small parameter values still requires careful cryptanalysis, and consideration of the best known attacks. (See, for example, [22,15,9].)

1.2 Techniques and Comparison to Related Work

Our results for SIS and LWE are technically disjoint, and all they have in common is the goal of proving hardness results for smaller values of the parameters. So, we describe our technical contributions in the analysis of these two problems separately.

SIS with small modulus. For SIS, as a warm-up, we first give a proof for a special case of the problem where the input is restricted to vectors of a special form (e.g., binary vectors). For this restricted version of SIS, we are able to give a self-reduction (from SIS to SIS) which reduces the size of the modulus. So, we can rely on previous worst-case to average-case reductions for SIS as "black boxes," resulting in an extremely simple proof. However, this simple self-reduction has some drawbacks. Beside the undesirable restriction on the SIS inputs, our reduction is rather loose with respect to the underlying worst-case lattice approximation problem: in order to establish the hardness of SIS with small moduli q (and restricted inputs), one needs to assume the worst-case hardness of lattice problems for rather large polynomial approximation factors. (By contrast, previous hardness results for larger moduli [21,13] only assumed hardness for quasi-linear approximation factors.) We address both drawbacks by giving a direct reduction from worst-case lattice problems to SIS with small modulus. This is our main SIS result, and it combines ideas from previous work [21,13] with two new technical ingredients:

- All previous SIS hardness proofs [1,8,16,21,13] solved worst-case lattice problems by iteratively finding (sets of linearly independent) lattice vectors of shorter and shorter length. Our first new technical ingredient (inspired by the pioneering work of Regev [26] on LWE) is the use a different intermediate problem: instead of finding progressively shorter lattice vectors, we consider the problem of sampling lattice vectors according to Gaussian-like distributions of progressively smaller widths. To the best of our knowledge, this is the first use of Gaussian lattice sampling as an intermediate worst-case problem in the study of SIS, and it appears necessary to lower the SIS modulus below n. We mention that Gaussian lattice sampling has been used before to reduce the modulus in hardness reductions for SIS [13], but still within the framework of iteratively finding short vectors (used in [13] to generate fresh Gaussian samples for the reduction), which results in larger moduli $q > n$.

– The use of Gaussian lattice sampling as an intermediate problem within the SIS hardness proof yields linear combinations of several discrete Gaussian samples with adversarially chosen coefficients. Our second technical ingredient, used to analyze these linear combinations, is a new convolution theorem for discrete Gaussians (Theorem 3), which strengthens similar ones previously proved in [24,6]. Here again, the strength of our new convolution theorem appears necessary to obtain hardness results for SIS with modulus smaller than n.

Our new convolution theorem may be of independent interest, and might find applications in the analysis of other lattice algorithms.

LWE with small errors. We now move to our results on LWE. For this problem, the best provably hard parameters to date were those obtained in the original paper of Regev [26], which employed Gaussian errors, and required them to have (expected) magnitude at least $\sqrt{n}$. These results were believed to be optimal due to a clever algorithm of Arora and Ge [2], which solves LWE in subexponential time when the errors are asymptotically smaller than $\sqrt{n}$. The possibility of circumventing this barrier by limiting the number of LWE samples was first suggested by Micciancio and Mol [18], who gave "sample preserving" search-to-decision reductions for LWE, and asked if LWE with small uniform errors could be proved hard when the number of available samples is sufficiently small. Our results provide a first answer to this question, and employ concepts and techniques from the work of Peikert and Waters [25] (see also [4]) on *lossy* (trapdoor) functions. In brief, a lossy function family is an indistinguishable pair of function families $\mathcal{F}, \mathcal{L}$ such that functions in $\mathcal{F}$ are injective and those in $\mathcal{L}$ are lossy, in the sense that they map their common domain to much smaller sets, and therefore lose information about the input. As shown in [25], from the indistinguishability of $\mathcal{F}$ and $\mathcal{L}$, it follows that the function families $\mathcal{F}$ and $\mathcal{L}$ are both one-way.

In the full version of this paper [20] we present a generalized framework for the study of lossy function families, which does not require the functions to have trapdoors, and applies to arbitrary (not necessarily uniform) input distributions. While the techniques we use are all standard, and our definitions are minor generalizations of the ones given in [25], we believe that our framework provides a conceptual simplification of previous work, relating the relatively new notion of lossy functions to the classic security definitions of second-preimage resistance and uninvertibility.

The lossy function framework is used to prove the hardness of LWE with small uniform errors and (necessarily) a small number of samples. Specifically, we use the standard LWE problem (with large Gaussian errors) to set up a lossy function family $\mathcal{F}, \mathcal{L}$. (Similar families with trapdoors were constructed in [25,4], but not for the parameterizations required to obtain interesting hardness results for LWE.) The indistinguishability of $\mathcal{F}$ and $\mathcal{L}$ follows directly from the hardness of the underlying LWE problem. The new hardness result for LWE (with small errors) is equivalent to the one-wayness of $\mathcal{F}$, and is proved by a relatively

standard analysis of the second-preimage resistance and uninvertibility of certain subset-sum functions associated to $\mathcal{L}$.

Our results, as well as previous work based on lossyness arguments, relies to some extent on the entropy in the secret vector. For simplicity, we specialize our analysis to the uniform input distribution (over arbitrary sets of short vectors), which makes counting arguments and entrody arguments essentially the same. Our results do generalize without much difficulty to other non-uniform distributions having sufficient min-entropy (unpredictability) over input sets with small diameter.

Comparison to related work. In a recent and independent work Döttling and Müller-Quade [11] also used a lossyness argument to prove new hardness results for LWE. Beside the use of a similar high level lossyness argument, the technical details of the proof are quite different, and the results are different as well. Just like our work, [11] proves hardness for uniformly distributed errors, and requires the number of m of samples to be fixed in advance. However, [11] requires the noise bound to be bigger than $\sqrt{n}$ (in fact, at least $m\sqrt{n}$, where m is the number of samples,) while in our work the errors can be smaller than $\sqrt{n}$, or even binary. On the other hand, when the magnitude of the errors is large $\sqrt{n}m \gg \sqrt{n}$, [11] allows the number of samples $m = n^{O(1)}$ to be an arbitrary large polynomial, while here we require it to be linear $m = \Theta(n)$. Another (more technical) difference between the two results is that our proof is based on a fairly general counting argument that allows error distributions that are uniform over arbitrary sets (of short vectors), while [11] only applies to uniform distributions over more structured sets, e.g., all vectors within a regularly shaped convex region of space.

1.3 Notation and Background

We briefly recall the (mostly standard) notation and background used in this paper, and refer the reader to the full version [20] for a more detailed account. We use standard asymptotic notation $O, \tilde{O}, \Omega, o, \omega$, and write ω_n as an abbreviation for $\omega(\sqrt{\log n})$. We write $[\mathcal{X}]$ to denote the support of a probability distribution $\mathcal{X}$, and $\mathcal{U}(X)$ for the uniform distribution over a set X. We assume familiarity with the notion of *computational indistinguishability*, and standard notions of security for function families, like *collision resistance*, *one-wayness*, *uninvertibility*, *second-preimage resistance*, and *pseudorandomness*. A *lossy function family* (slightly generalizing the concept of lossy trapdoor functions introduced in [25]) is a pair of computationally indistinguishable probability distributions $\mathcal{L}, \mathcal{F}$ over (descriptions of) functions $F \subseteq X \to Y$, such that $\mathcal{L}$ is an uninvertible function family, and $\mathcal{F}$ is a second preimage resistant function family, both with respect to some input distribution $\mathcal{X}$ over the domain X. It easily follows from the definition that both $\mathcal{F}$ and $\mathcal{L}$ are one-way function families with respect to input $\mathcal{X}$. A function family $\mathcal{L}$ is uninvertible with respect to the uniform distribution $\mathcal{X} = \mathcal{U}(X)$ if (and only if) the expectation $\mathbb{E}_{f \leftarrow \mathcal{L}}[|f(X)|]/|X|$ is negligible. Moreover, if $\mathcal{F} : X \to Y$ is uninvertible (with respect to input distribution $\mathcal{X}$) and

$\mathcal{G} : Y \to Z$ is an arbitrary function family, then the composition $\mathcal{G} \circ \mathcal{F} : X \to Z$ (defined in the obvious way, as the result of sampling two functions from $\mathcal{G}$ and $\mathcal{F}$ independently at random, and taking their function composition) is also uninvertible on input $\mathcal{X}$.

An n-dimensional (full rank) *lattice* is the set Λ of integer combinations of n linearly independent (basis) vectors $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n \in \mathbb{R}^n$. The (decision version of the) Shortest Vector Problem, GapSVP_γ asks to approximate within a factor γ the (Euclidean) length of the shortest nonzero vector in the lattice generated by an input basis $\mathbf{B} = [\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n]$. The Shortest Independent Vectors Problem SIVP_γ asks to find n linearly independent lattice vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$ of length $\max_i \|\boldsymbol{v}_i\| \leq \gamma\lambda_n$ within a factor γ from the optimum value λ_n. The *Gram-Schmidt* minimum of a lattice Λ is $\tilde{bl}(\Lambda) = \min_{\mathbf{B}} \|\widetilde{\mathbf{B}}\|$, where $\widetilde{\boldsymbol{b}}_i$ is projection of $\boldsymbol{b}_i$ orthogonal to $\mathrm{span}(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_{i-1})$, and $\mathbf{B} = [\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n]$ ranges over all possible bases of Λ.

The *Gaussian function* $\rho_s(\boldsymbol{x}) = \exp(-\pi\|\boldsymbol{x}\|^2/s^2)$ defines a *discrete Gaussian distribution* $D_{\Lambda+\boldsymbol{c},s}$ over a lattice coset $\Lambda + \boldsymbol{c}$, which samples each element $\boldsymbol{x} \in \Lambda + \boldsymbol{c}$ with probability $\rho_s(\boldsymbol{x})/\rho_s(\Lambda + \boldsymbol{c})$. For any (typically negligible) $\epsilon > 0$, the *smoothing parameter* $\eta_\epsilon(\Lambda)$ [21] is the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon$, where $\Lambda^* = \{\boldsymbol{x} : \forall \boldsymbol{y} \in \Lambda, \langle \boldsymbol{x}, \boldsymbol{y} \rangle \in \mathbb{Z}\}$ is the *dual lattice*. The tensor product of two lattices with bases $\mathbf{B}, \mathbf{B}'$ is the lattice with a basis $\mathbf{B} \otimes \mathbf{B}'$ obtained replacing each entry $b_{i,j}$ of $\mathbf{B}$ with the block $b_{i,j} \cdot \mathbf{B}'$, and it satisfies $\eta(\Lambda_1 \otimes \Lambda_2) \leq \tilde{bl}(\Lambda_1) \cdot \eta(\Lambda_2)$. Sampling $D_{\Lambda,\sigma}$ for some $\sigma \geq 2\eta(\Lambda)$ allows to solve SIVP_γ within a factor $\gamma = \tilde{O}(\sigma\sqrt{n})$.

The *Short Integer Solution* function family $\mathrm{SIS}(m, n, q, X)$ is the set of all functions $f_{\mathbf{A}}(\boldsymbol{x}) = \mathbf{A}\boldsymbol{x} \bmod q$ indexed by $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with domain $X \subseteq \mathbb{Z}^m$ and range $Y = \mathbb{Z}_q^n$. The *Learning With Errors* function family $\mathrm{LWE}(m, n, q, X)$, is the set of all functions $g_{\mathbf{A}}(\boldsymbol{s}, \boldsymbol{x}) = \mathbf{A}^T\boldsymbol{s} + \boldsymbol{x} \bmod q$ indexed by $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with domain $\mathbb{Z}_q^n \times X$ and range $Y = \mathbb{Z}_q^m$. We sometimes write $\mathrm{SIS}(m, n, q, \beta)$ for some real $\beta > 0$, to refer to the problem of finding a nonzero vector $\boldsymbol{z} \in \mathbb{Z}^m$ of length at most $\|\boldsymbol{z}\| \leq \beta$ such that $f_{\mathbf{A}}(\boldsymbol{z}) = \mathbf{0}$. The SIS input is usually chosen according to either the uniform distribution $\mathcal{U}(X)$ over the set $X = \{0, \ldots, s-1\}^m$ or $X = \{-s, \ldots, 0, \ldots, s\}^m$, or the discrete Gaussian distribution $D^m_{\mathbb{Z},s}$. For the LWE problem, the input is usually chosen according to distribution $\mathcal{U}(\mathbb{Z}_q^n) \times \mathcal{X}$, where $\mathcal{X}$ is one of the SIS input distributions. This makes the $\mathrm{LWE}(m, n, q)$ and $\mathrm{SIS}(m, m-n, q)$ function families equivalent via a lattice duality argument [17,18]]. The SIS and LWE functions can be shown to be collision resistant, one-way, uninvertible, pseudorandom, etc., for appropriate parameters and input distributions, based on the assumption that SIVP_γ and/or GapSVP_γ are (quantum) hard in the worst case [26,23,7].

2 Hardness of SIS with Small Modulus

As a warm-up, we first give a simplified proof that the $\mathrm{SIS}(m, n, q, \beta)$ function family is collision resistant for moduli q as small as $n^{1/2+\delta}$, by a reduction between SIS problems with different parameters. Previous hardness results

based on worst-case lattice assumptions require the modulus q to be at least $\beta \cdot \omega(\sqrt{n \log n})$ [13, Theorem 9.2], and $\beta \geq \sqrt{n \log q}$ is needed to guarantee that a nontrivial solution exists. For such parameters, SIS is collision resistant assuming the hardness of approximating worst-case lattice problems to within $\approx \beta\sqrt{n}$ factors.

The intuition behind our proof for smaller moduli is easily explained. We reduce SIS with modulus q^c and solution bound β^c (for any constant integer $c \geq 1$) to SIS with modulus q and bound β. Then as long as $(q/\beta)^c \geq \omega(\sqrt{n \log n})$, the former problem enjoys worst-case hardness, hence so does the latter. Thus we can take $q = \beta \cdot n^\delta$ for any constant $\delta > 0$, and $c > 1/(2\delta)$. Notice, however, that the underlying approximation factor for worst-case lattice problems is $\approx \beta^c\sqrt{n} \geq n^{1/2+1/(4\delta)}$, which, while still polynomial, degrades severely as δ approaches 0. In the next subsection we give a direct reduction from worst-case lattice problems to SIS with a small modulus, which does not have this drawback.

The above discussion is formalized in the following proposition. For technical reasons, we prove that $\mathrm{SIS}(m, n, q, X)$ is collision resistant assuming that the domain X has the property that all SIS solutions $\boldsymbol{z} \in (X - X) \setminus \{\mathbf{0}\}$ satisfy $\gcd(\boldsymbol{z}, q) = 1$. This restriction is satisfied in many (but not all) common settings, e.g., when $q > \beta$ is prime, or when $X \subseteq \{0, 1\}^m$ is a set of binary vectors. For simplicity, we describe the reduction assuming an SIS oracle that finds collisions with overwhelming probability. The reduction can be easily adapted to oracles that solve SIS with only nonnegligible probability using standard repetition/amplification techniques.

Proposition 1. *Let n, q, m, β and $X \subseteq \mathbb{Z}^m$ be such that $\gcd(\boldsymbol{x} - \boldsymbol{x}', q) = 1$ and $\|\boldsymbol{x} - \boldsymbol{x}'\| \leq \beta$ for any distinct $\boldsymbol{x}, \boldsymbol{x}' \in X$. For any positive integer c, there is a deterministic reduction from collision-finding for $\mathrm{SIS}(m^c, n, q^c, \beta^c)$ to collision-finding for $\mathrm{SIS}(m, n, q, X)$ (in both cases, with overwhelming advantage). The reduction runs in time polynomial in its input size, and makes fewer than m^c calls to its oracle.*

Proof. Let $\mathcal{A}$ be an efficient algorithm that finds a collision for $\mathrm{SIS}(m, n, q, X)$ with overwhelming advantage. We use it to find a nonzero solution for $\mathrm{SIS}(m^c, n, q^c, \beta^c)$. Let $\mathbf{A} \in \mathbb{Z}_{q^c}^{n \times m^c}$ be an input SIS instance. Partition the columns of $\mathbf{A}$ into m^{c-1} blocks $\mathbf{A}_i \in \mathbb{Z}_{q^c}^{n \times m}$, and for each one, invoke $\mathcal{A}$ to find a collision modulo q, i.e., a pair of distinct vectors $\boldsymbol{x}_i, \boldsymbol{x}'_i \in X$ such that $\mathbf{A}_i \boldsymbol{z}_i = \mathbf{0} \bmod q$, where $\boldsymbol{z}_i = \boldsymbol{x}_i - \boldsymbol{x}'_i$ and $\|\boldsymbol{z}_i\| \leq \beta$.

For each i, since $\gcd(\boldsymbol{z}_i, q) = 1$ and $\mathbf{A}_i \boldsymbol{z}_i = \mathbf{0} \bmod q$, the vector $\boldsymbol{a}'_i = (\mathbf{A}_i \boldsymbol{z}_i)/q \in \mathbb{Z}_{q^{c-1}}^n$ is uniformly random, even after conditioning on $\boldsymbol{z}_i$ and $\mathbf{A}_i \bmod q$. So, the matrix $\mathbf{A}' \in \mathbb{Z}_{q^{c-1}}^{n \times m^{c-1}}$ made up of all these columns is uniformly random. By induction on c, using $\mathcal{A}$ we can find a nonzero solution $\boldsymbol{z}' \in \mathbb{Z}^{m^{c-1}}$ such that $\mathbf{A}'\boldsymbol{z}' = \mathbf{0} \bmod q^{c-1}$ and $\|\boldsymbol{z}'\| \leq \beta^{c-1}$. Then it is easy to verify that a nonzero solution for the original instance $\mathbf{A}$ is given by $\boldsymbol{z} = (z'_1 \cdot \boldsymbol{z}_1, \ldots, z'_{m^{c-1}} \cdot \boldsymbol{z}_{m^{c-1}}) \in \mathbb{Z}^{m^c}$, and that $\|\boldsymbol{z}\| \leq \|\boldsymbol{z}'\| \cdot \max_i \|\boldsymbol{z}_i\| \leq \beta^c$. Finally, the total number of calls to $\mathcal{A}$ is $\sum_{i=0}^{c-1} m^i < m^c$, as claimed. □

Direct Reduction As mentioned above, the large worst-case approximation factor associated with the use of Proposition 1 is undesirable, as is (to a lesser extent) the restriction that $\gcd(X-X, q) = 1$. To eliminate these drawbacks, we next give a direct proof that SIS is collision resistant for small q, based on the assumed hardness of worst-case lattice problems. The underlying approximation factor for these problems can be as small as $\tilde{O}(\beta\sqrt{n})$, which matches the best known factors obtained by previous proofs (which require a larger modulus q). Our new proof combines ideas from [21,13] and Proposition 1, as well as a new convolution theorem for discrete Gaussians which strengthens similar ones previously proved in [24,6].

Our proof of the convolution theorem is substantially different and, we believe, technically simpler than the prior ones. In particular, it handles the sum of many Gaussian samples all at once, whereas previous proofs used induction from a base case of two samples. With the inductive approach, it is technically complex to verify that all the intermediate Gaussian parameters (which involve harmonic means) satisfy the hypotheses. Moreover, the intermediate parameters can depend on the order in which the samples are added in the induction, leading to unnecessarily strong hypotheses on the original parameters. Due to space constraints, the proof of the convolution theorem is given in the full version [20].

Theorem 3. *Let Λ be an n-dimensional lattice, $\boldsymbol{z} \in \mathbb{Z}^m$ a nonzero integer vector, $s_i \geq \sqrt{2}\|\boldsymbol{z}\|_\infty \cdot \eta(\Lambda)$, and $\Lambda + \boldsymbol{c}_i$ arbitrary cosets of Λ for $i = 1, \ldots, m$. Let $\boldsymbol{y}_i$ be independent vectors with distributions $D_{\Lambda+\boldsymbol{c}_i, s_i}$, respectively. Then the distribution of $\boldsymbol{y} = \sum_i z_i \boldsymbol{y}_i$ is statistically close to $D_{Y,s}$, where $Y = \gcd(\boldsymbol{z})\Lambda + \boldsymbol{c}$, $\boldsymbol{c} = \sum_i z_i \boldsymbol{c}_i$, and $s = \sqrt{\sum_i (z_i s_i)^2}$. In particular, if $\gcd(\boldsymbol{z}) = 1$ and $\sum_i z_i \boldsymbol{c}_i \in \Lambda$, then $\boldsymbol{y}$ is distributed statistically close to $D_{\Lambda,s}$.*

The convolution theorem implies the following simple but useful lemma, which shows how to convert samples having a broad range of parameters into ones having parameters in a desired narrow range.

Lemma 1. *There is an efficient algorithm which, given a basis $\mathbf{B}$ of some lattice Λ, some $R \geq \sqrt{2}$, and access to samples $(\boldsymbol{y}_i, s_i)$ where each $s_i \in [\sqrt{2}, R] \cdot \eta(\Lambda)$ is* arbitrary *and each $\boldsymbol{y}_i$ has distribution D_{Λ,s_i}, with overwhelming probability outputs a pair $(\boldsymbol{y}, s)$ such that $s \in [R, \sqrt{2}R] \cdot \eta(\Lambda)$ and $\boldsymbol{y}$ has distribution statistically close to $D_{\Lambda,s}$.*

Proof. Let $\omega_n = \omega(\sqrt{\log n})$ satisfy $\omega_n \leq \sqrt{n}$. The algorithm draws $2n^2$ input samples, and works as follows: if at least n^2 of the samples have parameters $s_i \leq R \cdot \eta(\Lambda)/(\sqrt{n} \cdot \omega_n)$, then with overwhelming probability they all have lengths bounded by $R \cdot \eta(\Lambda)/\omega_n$ and they include n linearly independent vectors. Using such vectors we can construct a basis $\mathbf{S}$ such that $\|\widetilde{\mathbf{S}}\| \leq R \cdot \eta(\Lambda)/\omega_n$, and with the sampling algorithm of [13, Theorem 4.1] we can generate samples having parameter $R \cdot \eta(\Lambda)$. Otherwise, at least n^2 of the samples $(\boldsymbol{y}_i, s_i)$ have parameters $s_i \geq \max\{R/n, \sqrt{2}\} \cdot \eta(\Lambda)$. Then by summing an appropriate subset of those $\boldsymbol{y}_i$, by the convolution theorem we can obtain a sample having parameter in the desired range. □

The next lemma is the heart of our reduction. The novel part, corresponding to the properties described in the second item, is a way of using a collision-finding oracle to reduce the Gaussian width of samples drawn from a lattice. The first item corresponds to the guarantees provided by previous reductions.

Lemma 2. *Let m, n be integers, $S = \{\boldsymbol{z} \in \mathbb{Z}^m \setminus \{\mathbf{0}\} \mid \|\boldsymbol{z}\| \leq \beta \wedge \|\boldsymbol{z}\|_\infty \leq \beta_\infty\}$ for some real $\beta \geq \beta_\infty > 0$, and q an integer modulus with at most $\mathrm{poly}(n)$ integer divisors less than β_∞. There is a probabilistic polynomial time reduction that, on input any basis $\mathbf{B}$ of a lattice Λ and sufficiently many samples $(\boldsymbol{y}_i, s_i)$, where each $s_i \geq \sqrt{2}q \cdot \eta(\Lambda)$ may be arbitrary and each $\boldsymbol{y}_i$ has distribution D_{Λ,s_i}, and given access to an $\mathrm{SIS}(m, n, q, S)$ oracle (that finds a solution $\boldsymbol{z} \in S$ with nonnegligible probability) outputs (with overwhelming probability) a sample $(\boldsymbol{y}, s)$ where $\min s_i/q \leq s \leq (\beta/q) \cdot \max s_i$, and $\boldsymbol{y} \in \Lambda$ is such that:*

- $\mathbb{E}[\|\boldsymbol{y}\|] \leq (\beta\sqrt{n}/q) \cdot \max s_i$, *and for any fixed subspace $H \subset \mathbb{R}^n$ of dimension at most $n-1$, we have $\Pr[\boldsymbol{y} \notin H] \geq 1/10$.*
- *Moreover, if each $s_i \geq \sqrt{2}\beta_\infty q \cdot \eta(\Lambda)$, then the distribution of $\boldsymbol{y}$ is statistically close to $D_{\Lambda,s}$*

Proof. Let $\mathcal{A}$ be the SIS oracle. Without loss of generality, we can assume that whenever $\mathcal{A}$ outputs a valid solution $\boldsymbol{z} \in S$, we have that $\gcd(\boldsymbol{z})$ divides q. This is because for any integer vector $\boldsymbol{z}$, if $\mathbf{A}\boldsymbol{z} = \mathbf{0} \bmod q$ then also $\mathbf{A}((g/d)\boldsymbol{z}) = \mathbf{0} \bmod q$, where $d = \gcd(\boldsymbol{z})$ and $g = \gcd(d, q)$. Moreover, $(g/d)\boldsymbol{z} \in S$ holds true and $\gcd((g/d)\boldsymbol{z}) = \gcd(\boldsymbol{z}, q)$ divides q. Let d be such that $\mathcal{A}$ outputs, with non-negligible probability, a valid solution $\boldsymbol{z}$ satisfying $\gcd(\boldsymbol{z}) = d$. Such a d exists because $\gcd(\boldsymbol{z})$ is bounded by β_∞ and divides q, so by assumption there are only polynomially many possible values of d. Let $q' = q/d$, which is an integer. By increasing m and using standard amplification techniques, we can make the probability that $\mathcal{A}$ outputs such a solution (satisfying $\boldsymbol{z} \in S$, $\mathbf{A}\boldsymbol{z} = \mathbf{0} \pmod q$ and $\gcd(\boldsymbol{z}) = d$) exponentially close to 1.

Let $(\boldsymbol{y}_i, s_i)$ for $i = 1, \ldots, m$ be input samples, where $\boldsymbol{y}_i$ has distribution D_{Λ,s_i}. Write each $\boldsymbol{y}_i$ as $\boldsymbol{y}_i = \mathbf{B}\boldsymbol{a}_i \bmod q'\Lambda$ for $\boldsymbol{a}_i \in \mathbb{Z}_{q'}^n$. Since $s_i \geq q'\eta(\Lambda)$ the distribution of $\boldsymbol{a}_i$ is statistically close to uniform over $\mathbb{Z}_{q'}^n$. Let $\mathbf{A} = [\boldsymbol{a}_1 \mid \cdots \mid \boldsymbol{a}_m] \in \mathbb{Z}_q^{n \times m}$, and choose $\mathbf{A'} \in \mathbb{Z}_d^{n \times m}$ uniformly at random. Since $\mathbf{A}$ is statistically close to uniform over $\mathbb{Z}_{q'}^{n \times m}$, the matrix $\mathbf{A} + q'\mathbf{A'}$ is statistically close to uniform over $\mathbb{Z}_q^{n \times m}$. Call the oracle $\mathcal{A}$ on input $\mathbf{A} + q'\mathbf{A'}$, and obtain (with overwhelming probability) a nonzero $\boldsymbol{z} \in S$ with $\gcd(\boldsymbol{z}) = d$, $\|\boldsymbol{z}\| \leq \beta$, $\|\boldsymbol{z}\|_\infty \leq \beta_\infty$ and $(\mathbf{A} + q'\mathbf{A'})\boldsymbol{z} = \mathbf{0} \bmod q$. Notice that $q'\mathbf{A'}\boldsymbol{z} = q\mathbf{A'}(\boldsymbol{z}/d) = \mathbf{0} \bmod q$ because $(\boldsymbol{z}/d)$ is an integer vector. Therefore $\mathbf{A}\boldsymbol{z} = \mathbf{0} \bmod q$. Finally, the reduction outputs $(\boldsymbol{y}, s)$, where $\boldsymbol{y} = \sum_i z_i \boldsymbol{y}_i/q$ and $s = \sqrt{\sum_i (s_i z_i)^2}/q$. Notice that $z_i \boldsymbol{y}_i \in q\Lambda + \mathbf{B}(z_i \boldsymbol{a}_i)$ because $\gcd(\boldsymbol{z}) = d$, so $\boldsymbol{y} \in \Lambda$.

Notice that s satisfies the stated bounds because $\boldsymbol{z}$ is a nonzero integer vector. We next analyze the distribution of $\boldsymbol{y}$. For any fixed $\boldsymbol{a}_i$, the conditional distribution of each $\boldsymbol{y}_i$ is $D_{q'\Lambda + \mathbf{B}\boldsymbol{a}_i, s_i}$, where $s_i \geq \sqrt{2}\eta(q'\Lambda)$. The claim on $\mathbb{E}[\|\boldsymbol{y}\|]$ then follows from [21, Lemma 2.11 and Lemma 4.3] and Hölder's inequality. The claim on the probability that $\boldsymbol{y} \notin H$ was initially shown in the preliminary version of [21]; see also [26, Lemma 3.15].

Now assume that $s_i \geq \sqrt{2}\beta_\infty q \cdot \eta(\Lambda) \geq \sqrt{2}\|\boldsymbol{z}\|_\infty \cdot \eta(q'\Lambda)$ for all i. By Theorem 3 the distribution of $\boldsymbol{y}$ is statistically close to $D_{Y/q,s}$ where $Y = \gcd(\boldsymbol{z}) \cdot q'\Lambda + \mathbf{B}(\mathbf{A}\boldsymbol{z})$. Using $\mathbf{A}\boldsymbol{z} = \mathbf{0} \bmod q$ and $\gcd(\boldsymbol{z}) = d$, we get $Y = q\Lambda$. Therefore $\boldsymbol{y}$ has distribution statistically close to $D_{\Lambda,s}$, as claimed. □

Building on Lemma 2, our next lemma shows that for any $q \geq \beta \cdot n^{\Omega(1)}$, an SIS oracle can be used to obtain Gaussian samples of width close to $2\beta\beta_\infty \cdot \eta(\Lambda)$.

Lemma 3. *Let m, n, q, S as in Lemma 2, and also assume $q/\beta \geq n^\delta$ for some constant $\delta > 0$. There is an efficient reduction that, on input any basis $\mathbf{B}$ of an n-dimensional lattice Λ, an upper bound $\eta \geq \eta(\Lambda)$, and given access to an $\mathrm{SIS}(m, n, q, S)$ oracle (which finds a solution $\boldsymbol{z} \in S$ with nonnegligible probability), outputs (with overwhelming probability) a sample $(\boldsymbol{y}, s)$ where $\sqrt{2}\beta_\infty \cdot \eta \leq s \leq 2\beta_\infty\beta \cdot \eta$ and $\boldsymbol{y}$ has distribution statistically close to $D_{\Lambda,s}$.*

Proof. By applying the LLL basis reduction algorithm [14] to the basis $\mathbf{B}$, we can assume without loss of generality that $\|\widetilde{\mathbf{B}}\| \leq 2^n \cdot \eta(\Lambda)$. The main procedure, described below, produces samples having parameters in the range $[1, q] \cdot \sqrt{2}\beta_\infty \cdot \eta$. On these samples we run the procedure from Lemma 1 (with $R = \sqrt{2}\beta_\infty q \cdot \eta$) to obtain samples having parameters in the range $[\sqrt{2}, 2] \cdot \beta_\infty q \cdot \eta$. Finally, we invoke the reduction from Lemma 2 on those samples to obtain a sample satisfying the conditions in the Lemma statement.

The main procedure works in a sequence of phases $i = 0, 1, 2, \ldots$. In phase i, the input is a basis $\mathbf{B}_i$ of Λ, where initially $\mathbf{B}_0 = \mathbf{B}$. The basis $\mathbf{B}_i$ is used in the discrete Gaussian sampling algorithm of [13, Theorem 4.1] to produce samples $(\boldsymbol{y}, s_i)$, where $s_i = \max\{\|\widetilde{\mathbf{B}}_i\| \cdot \omega_n, \sqrt{2}\beta_\infty\eta\} \geq \sqrt{2}\beta_\infty\eta$ and $\boldsymbol{y}_i$ has distribution statistically close to D_{Λ,s_i}. Phase i either manages to produce a sample $(\boldsymbol{y}, s)$ with s in the desired range $[1, q] \cdot \sqrt{2}\beta_\infty\eta$, or it produces a new basis $\mathbf{B}_{i+1}$ for which $\|\widetilde{\mathbf{B}}_{i+1}\| \leq \|\widetilde{\mathbf{B}}_i\|/2$, which is the input to the next phase. The number of phases before termination is polynomial in n, by hypothesis on $\mathbf{B}$.

If $\|\widetilde{\mathbf{B}}_i\| \cdot \omega_n \leq \sqrt{2}q\beta_\infty\eta$, then this already gives samples with $s_i \in [1, q]\sqrt{2}\beta_\infty\eta$ in the desired range, and we can terminate the main phase. So, we may assume that $s_i = \|\widetilde{\mathbf{B}}_i\| \cdot \omega_n \geq \sqrt{2}q\beta_\infty\eta$. Each phase i proceeds in some constant $c \geq 1/\delta$ number of sub-phases $j = 1, 2, \ldots, c$, where the inputs to the first sub-phase are the samples $(\boldsymbol{y}, s_i)$ generated as described above. We recall that these samples satisfy $s_i \geq \sqrt{2}q\beta_\infty\eta$. The same will be true for the samples passed as input to all other subsequent subphases. So, each subphase receives as input samples $(\boldsymbol{y}, s)$ satisfying all the hypotheses of Lemma 2, and we can run the reduction from that lemma to generate new samples $(\boldsymbol{y}', s')$ having parameters s' bounded from above by $s_i \cdot (\beta/q)^j$, and from below by $\sqrt{2}\beta_\infty\eta$. If any of the produces samples satisfies $s' \leq q\sqrt{2}\beta_\infty\eta$, then we can terminate the main procedure with $(\boldsymbol{y}', s')$ as output. Otherwise, all samples produced during the subphase satisfy $s' > q\sqrt{2}\beta_\infty\eta$, and they can be passed as input to the next sub-phase. Notice that the total runtime of all the sub-phases is $\mathrm{poly}(n)^c$, because each invocation of the reduction from Lemma 2 relies on $\mathrm{poly}(n)$ invocations of the reduction in the previous sub-phase; this is why we need to limit the number of sub-phases to a constant c.

If phase i ends up running all its sub-phases without ever finding a sample with $s' \in [1,q]\sqrt{2}\beta_\infty\eta$, then it has produced samples whose parameters are bounded by $(\beta/q)^c \le s_i \le s_i/\sqrt{n}$. It uses n^2 of these samples, which with overwhelming probability have lengths all bounded by $s_i/\sqrt{n}$, and include n linearly independent vectors. It transforms those vectors into a basis $\mathbf{B}_{i+1}$ with $\|\widetilde{\mathbf{B}}_{i+1}\| \le s_i/\sqrt{n} \le \|\widetilde{B}\|_i\omega_n/\sqrt{n} \le \|\widetilde{\mathbf{B}}_i\|/2$, as input to the next phase. □

We can now prove our main theorem, reducing worst-case lattice problems with $\max\{1, \beta\beta_\infty/q\} \cdot \tilde{O}(\beta\sqrt{n})$ approximation factors to SIS, when $q \ge \beta \cdot n^{\Omega(1)}$.

Theorem 4. *Let m, n be integers, $S = \{\boldsymbol{z} \in \mathbb{Z}^m \setminus \{\mathbf{0}\} \mid \|\boldsymbol{z}\| \le \beta \wedge \|\boldsymbol{z}\|_\infty \le \beta_\infty\}$ for some real $\beta \ge \beta_\infty > 0$, and $q \ge \beta \cdot n^{\Omega(1)}$ be an integer modulus with at most $\mathrm{poly}(n)$ integer divisors less than β_∞. For some $\gamma = \max\{1, \beta\beta_\infty/q\} \cdot O(\beta\sqrt{n})$, there is an efficient reduction from $\mathrm{SIVP}_\gamma^\eta$ (and hence also from standard $\mathrm{SIVP}_{\gamma\cdot\omega_n}$) on n-dimensional lattices to solving the collision-finding problem $\mathrm{SIS}(m, n, q, S)$ with non-negligible advantage.*

Proof. Given an input basis $\mathbf{B}$ of a lattice Λ, we can apply the LLL algorithm to obtain a 2^n-approximation to $\eta(\Lambda)$, and by scaling we can assume that $\eta(\Lambda) \in [1, 2^n]$. For $i = 1, \ldots, n$, we run the procedure described below for each candidate upper bound $\eta_i = 2^i$ on $\eta(\Lambda)$. Each call to the procedure either fails, or returns a set of linearly independent vectors in Λ whose lengths are all bounded by $(\gamma/2) \cdot \eta_i$. We return the first such obtained set (i.e., for the minimal value of i). As we show below, as long as $\eta_i \ge \eta(\Lambda)$ the procedure returns a set of vectors with overwhelming probability. Since exactly one $\eta_i \in [1, 2) \cdot \eta(\Lambda)$, our reduction solves $\mathrm{SIVP}_\gamma^\eta$ with overwhelming probability, as claimed.

The procedure invokes the reduction from Lemma 3 with $\eta = \eta_i$ to obtain samples with parameters in the range $[\sqrt{2}\beta_\infty, \sqrt{2}\beta\beta_\infty] \cdot \eta$. On these samples we run the procedure from Lemma 1 with $R = \max\{\sqrt{2}q, \sqrt{2}\beta\beta_\infty\}$ to obtain samples having parameters in the range $[R, \sqrt{2}R] \cdot \eta$. On such samples we repeatedly run (using independent samples each time) the reduction from Lemma 2. After enough runs, we obtain with overwhelming probability a set of linearly independent lattice vectors all having lengths at most $(\gamma/2) \cdot \eta$, as required. □

3 Hardness of LWE with Small Uniform Errors

In this section we prove the hardness of inverting the LWE function even when the error vectors have very small entries, provided the number of samples is sufficiently small. We proceed similarly to [25,4], by using the LWE assumption (for discrete Gaussian error) to construct a lossy family of functions with respect to a uniform distribution over small inputs. However, the parameterization we obtain is different from those in [25,4], allowing us to obtain *pseudorandomness* of LWE under *very small* (e.g., binary) inputs, for a number of LWE samples that exceeds the LWE dimension.

Our results and proofs are more naturally formulated using the SIS function family. So, we will first study the problem in terms of SIS, and then reformulate

the results in terms of LWE by lattice duality [17,18]. All statements in this section are proved by relatively simple counting arguments, and the reader is referred to the full version [20] for details. We recall that the main difference between this section and Section 2, is that here we consider parameters for which the resulting functions are essentially injective, or more formally, statistically second-preimage resistant. The following lemma gives sufficient conditions that ensure this property.

Lemma 4. *For any integers m, k, q, s and set $X \subseteq [s]^m$, the function family $\mathrm{SIS}(m, k, q)$ is (statistically) ϵ-second preimage resistant with respect to the uniform input distribution $\mathcal{U}(X)$ for $\epsilon = |X| \cdot (s'/q)^k$, where s' is the largest factor of q smaller than s.*

Proof. Let $\boldsymbol{x} \leftarrow \mathcal{U}(X)$ and $\mathbf{A} \leftarrow \mathrm{SIS}(m, k, q)$ be chosen at random. We want to evaluate the probability that there exists an $\boldsymbol{x}' \in X \setminus \{\boldsymbol{x}\}$ such that $\mathbf{A}\boldsymbol{x} = \mathbf{A}\boldsymbol{x}' \pmod q$, or, equivalently, $\mathbf{A}(\boldsymbol{x} - \boldsymbol{x}') = \mathbf{0} \pmod q$. Fix any two distinct vectors $\boldsymbol{x}, \boldsymbol{x}' \in X$ and let $\boldsymbol{z} = \boldsymbol{x} - \boldsymbol{x}'$. The vector $\mathbf{A}\boldsymbol{z} \pmod q$ is distributed uniformly at random in $(d\mathbb{Z}/q\mathbb{Z})^k$, where $d = \gcd(q, z_1, \ldots, z_m)$. All coordinates of $\boldsymbol{z}$ are in the range $z_i \in \{-(s-1), \ldots, (s-1)\}$, and at least one of them is nonzero. Therefore, d is at most s' and $|d\mathbb{Z}_q^k| = (q/d)^k \geq (q/s')^k$. By union bound (over $\boldsymbol{x}' \in X \setminus \{\boldsymbol{x}\}$) for any $\boldsymbol{x}$, the probability that there is a second preimage $\boldsymbol{x}'$ is at most $(|X| - 1)(s'/q)^k$. □

We remark that, as shown in Section 2, even for parameter settings that do not fall within the range specified in Lemma 4, $\mathrm{SIS}(m, k, q)$ is collision resistant, and therefore also (computationally) second-preimage-resistant. This is all that is needed in the rest of this section. However, when $\mathrm{SIS}(m, k, q)$ is not statistically second-preimage resistant, the one-wayness proof that follows (see Theorem 5) is not very interesting: typically, in such settings, $\mathrm{SIS}(m, k, q)$ is also statistically uninvertible, and the one-wayness of $\mathrm{SIS}(m, k, q)$ trivially follows. So, below we focus on parameter settings covered by Lemma 4. We prove the one-wayness of $\mathcal{F} = \mathrm{SIS}(m, k, q, X)$ with respect to the uniform input distribution $\mathcal{X} = \mathcal{U}(X)$ by building a lossy function family $(\mathcal{L}, \mathcal{F}, \mathcal{X})$ where $\mathcal{L}$ is an auxiliary function family that we will prove to be uninvertible and computationally indistinguishable from $\mathcal{F}$. Due to space limitations, we only provide a sketch of the construction and analysis of the lossy function family, followed by formal statement of our main result, and description of some notable instantiations. The reader is referred to the full version [20] for more details.

Outline The idea behind our construction and proof is easily explained. The functions in our family $\mathcal{F}$ are defined by uniformly chosen random matrices $\mathbf{F} \in \mathbb{Z}_q^{k \times m}$. We define the auxiliary function family $\mathcal{L}$ by choosing the first $\ell < m$ colums of $\mathbf{L} = [\mathbf{A} \mid \ldots] \in \mathbb{Z}_q^{k \times m}$ uniformly at random, and setting the remaining columns to $\mathbf{AY}$ where $\mathbf{Y} \in \mathbb{Z}_q^{\ell \times m}$ is a random matrix with discrete gaussian entries of small width $\sigma > \sqrt{n}$. It follows from previous worst-case to average-case reductions [26,23,7], search-to-decision reductions and standard lattice duality

results (e.g., see [18],) that the matrix $\mathbf{L} = [\mathbf{A} \mid \mathbf{AY}]$ is pseudorandom. So, $\mathcal{F}$ and $\mathcal{L}$ are computationally indistinguishable. We already know from Lemma 4 that $\mathcal{F}$ is second preimage resistant. In order to conclude that $(\mathcal{F}, \mathcal{L}, \mathcal{X})$ is a lossy function family (and therefore, one-way, see Section 1.3,) it only remains to prove that $\mathcal{L}$ is uninvertible with respect to input distribution $\mathcal{X}$. To this end, express $\mathbf{L}$ as a product $\mathbf{L} = \mathbf{A}[\mathbf{I} \mid \mathbf{Y}]$. Notice that $\mathbf{L}$ is the composition of (linear) functions $[\mathbf{I} \mid \mathbf{Y}]$ and $\mathbf{A}$. The first function $[\mathbf{I} \mid \mathbf{Y}]$ is uninvertible with respect to input distribution $\mathcal{X}$ because its matrix has small entries and necessarily maps small input vectors X to a small set of short output vectors. The composition of an uninvertible function with an arbitrary function remains uninvertible. Therefore also $\mathbf{L} = [\mathbf{A} \mid \mathbf{AY}]$ is uninvertible. The rest of the proof is standard, and follows the general lossy function approach of [25]: the function $\mathbf{F}$ is uninvertible because any efficient inverter for $\mathbf{F}$ would allow to distinguish $\mathbf{F}$ from the uninvertible function $\mathbf{L}$, and contradict the pseudorandomness of $\mathbf{L}$. Since $\mathbf{F}$ is both uninvertible and second-preimage resistant, it is also one-way. (See Section 1.3 and full version [20].) Finally, using the search-to-decision reduction of [18] we conclude that $\mathcal{F}$ is not only one-way, but pseudorandom.

Main result We begin by defining our uninvertible function family consisting of matrices with small entries.

Definition 1. *For any probability distribution $\mathcal{Y}$ over $\mathbb{Z}^\ell$ and integer $m \geq \ell$, let $\mathcal{I}(m, \ell, \mathcal{Y})$ be the probability distribution over linear functions $[\mathbf{I} \mid \mathbf{Y}] \colon \mathbb{Z}^m \to \mathbb{Z}^\ell$ where $\mathbf{I}$ is the $\ell \times \ell$ identity matrix, and $\mathbf{Y} \in \mathbb{Z}^{\ell \times (m-\ell)}$ is obtained choosing each column of $\mathbf{Y}$ independently at random from $\mathcal{Y}$.*

The next lemma shows that, for appropriate parameter values, this is indeed an uninvertible function family.

Lemma 5. *Let $\mathcal{Y} = D^\ell_{\mathbb{Z},\sigma}$ be the discrete Gaussian distribution with parameter $\sigma > 0$, and let $X \subseteq \{-s, \ldots, s\}^m$. Then $\mathcal{I}(m, \ell, \mathcal{Y})$ is ϵ-uninvertible with respect to $\mathcal{U}(X)$, for $\epsilon = O(\sigma m s/\sqrt{\ell})^\ell/|X| + 2^{-\Omega(m)}$.*

Proof. It is enough to bound the expected size of $f(X)$ when $f \leftarrow \mathcal{I}(m, \ell, \mathcal{Y})$ is chosen at random. Remember that $f = [\mathbf{I} \mid \mathbf{Y}]$ where $\mathbf{Y} \leftarrow D^{\ell \times (m-\ell)}_{\mathbb{Z},\sigma}$. Since the entries of $\mathbf{Y} \in \mathbb{R}^{\ell \times (m-\ell)}$ are independent zero-mean subgaussians with parameter σ, by a standard bound from the theory of random matrices, the largest singular value $s_1(\mathbf{Y}) = \max_{\mathbf{0} \neq \boldsymbol{x} \in \mathbb{R}^m} \|\mathbf{Y}\boldsymbol{x}\|/\|\boldsymbol{x}\|$ of $\mathbf{Y}$ is at most $\sigma \cdot O(\sqrt{\ell} + \sqrt{m-\ell}) = \sigma \cdot O(\sqrt{m})$, except with probability $2^{-\Omega(m)}$. We now bound the ℓ_2 norm of all vectors in the image $f(X)$. Let $\boldsymbol{u} = (\boldsymbol{u}_1, \boldsymbol{u}_2) \in X$, with $\boldsymbol{u}_1 \in \mathbb{Z}^\ell$ and $\boldsymbol{u}_2 \in \mathbb{Z}^{m-\ell}$. Then $\|f(\boldsymbol{u})\| \leq \|\boldsymbol{u}_1 + \mathbf{Y}\boldsymbol{u}_2\| \leq \|\boldsymbol{u}_1\| + \|\mathbf{Y}\boldsymbol{u}_2\| \leq \left(\sqrt{\ell} + s_1(\mathbf{Y})\sqrt{m-\ell}\right) s \leq \left(\sqrt{\ell} + \sigma \cdot O(\sqrt{m})\sqrt{m-\ell}\right) s = O(\sigma m s)$. The number of integer points in the ℓ-dimensional zero-centered ball of radius $R = O(\sigma m s)$ can be bounded by a simple volume argument, as $|f(X)| \leq (R + \sqrt{\ell}/2)^n V_\ell = O(\sigma m s/\sqrt{\ell})^\ell$, where $V_\ell = \pi^{\ell/2}/(\ell/2)!$ is the volume of the ℓ-dimensional unit ball. Dividing by the size of X and accounting for the rare event that $s_1(\mathbf{Y})$ is not bounded as above, we get that $\mathcal{I}(m, \ell, \mathcal{Y})$ is ϵ-uninvertible for $\epsilon = O(\sigma m s/\sqrt{\ell})^\ell/|X| + 2^{-\Omega(m)}$. □

We can now prove the one-wayness of the SIS function family by defining and analyzing an appropriate lossy function family. The parameters below are set up to expose the connection with LWE, via lattice duality: $\mathrm{SIS}(m, m-n, q)$ corresponds to LWE in n dimensions (given m samples), whose one-wayness we are proving, while $\mathrm{SIS}(\ell = m-n+k, m-n, q)$ corresponds to LWE in $k \leq n$ dimensions, whose pseudorandomness we are assuming.

Theorem 5. *Let q be a modulus and let $\mathcal{X}, \mathcal{Y}$ be two distributions over $\mathbb{Z}^m$ and $\mathbb{Z}^\ell$ respectively, where $\ell = m-n+k$ for some $0 < k \leq n \leq m$, such that*

1. *$\mathcal{I}(m, \ell, \mathcal{Y})$ is uninvertible with respect to input distribution $\mathcal{X}$,*
2. *$\mathrm{SIS}(\ell, m-n, q)$ is pseudorandom with respect to input distribution $\mathcal{Y}$, and*
3. *$\mathrm{SIS}(m, m-n, q)$ is second-preimage resistant with respect to input distribution $\mathcal{X}$.*

Then $\mathcal{F} = \mathrm{SIS}(m, m-n, q)$ is one-way with respect to input distribution $\mathcal{X}$. In particular, if $\mathrm{SIS}(\ell, m-n, q)$ is pseudorandom with respect to the discrete Gaussian distribution $\mathcal{Y} = D^\ell_{\mathbb{Z},\sigma}$, then $\mathrm{SIS}(m, m-n, q)$ is $(2\epsilon + 2^{-\Omega(m)})$-one-way with respect to the uniform input distribution $\mathcal{X} = \mathcal{U}(X)$ over any set $X \subseteq \{-s, \ldots, s\}^m$ satisfying $(C'\sigma m s/\sqrt{\ell})^\ell/\epsilon \leq |X| \leq \epsilon \cdot (q/s')^{m-n}$, where s' is the largest divisor of q that is smaller than or equal to $2s$, and C' is the universal constant hidden by the $O(\cdot)$ notation from Lemma 5.

In order to conclude that the LWE function is pseudorandom (under worst-case lattice assumptions) for uniformly random small errors, we combine Theorem 5 with previous hardness results [26,23,7] and search-to-decision reductions [18,19]. For simplicity, we focus on the important case of a prime modulus q. Nearly identical results for composite moduli (e.g., those divisible by only small primes) are also easily obtained.

Theorem 6. *Let $0 < k \leq n \leq m - \omega(\log k) \leq k^{O(1)}$, $\ell = m-n+k$, $s \geq (Cm)^{\ell/(n-k)}$ for a large enough universal constant C, and q be a prime such that $\max\{3\sqrt{k}, (4s)^{m/(m-n)}\} \leq q \leq k^{O(1)}$. For any set $X \subseteq \{-s, \ldots, s\}^m$ of size $|X| \geq s^m$, the $\mathrm{SIS}(m, m-n, q)$ (equivalently, $\mathrm{LWE}(m, n, q)$) function family is one-way (and pseudorandom) with respect to the uniform input distribution $\mathcal{X} = \mathcal{U}(X)$, under the assumption that SIVP_γ and GapSVP_γ are (quantum) hard to approximate, in the worst case, on k-dimensional lattices to within a factor $\gamma = \tilde{O}(\sqrt{k} \cdot q)$.*

Example parameters We conclude the section with a few notable instantiations of the last theorem. We remark that the condition $|X| \geq s^m$ in Theorem 6 is not essential, and the same proof yields similar results also with weaker lower bounds on $|X|$. To obtain pseudorandomness for binary errors, we need $s = 2$ and $X = \{0,1\}^m$. For this value of s, the condition $s \geq (Cm)^{\ell/(n-k)}$ can be equivalently be rewritten as $m \leq (n-k) \cdot (1 + 1/\log_2(Cm))$, which can be satisfied by taking $k = n/(C' \log_2 n)$ and $m = n(1 + 1/(c \log_2 n))$ for any desired $c > 1$ and a sufficiently large constant $C' > 1/(1-1/c)$. For these values, the

modulus should satisfy $q \geq 8^{m/(m-n)} = 8n^{3c} = k^{O(1)}$, and can be set to any sufficiently large prime $p = k^{O(1)}$.[4]

Notice that for binary errors, both the worst-case lattice dimension k and the number $m - n$ of "extra" LWE samples (i.e., the number of samples beyond the LWE dimension n) are sublinear in the LWE dimension n: we have $k = \Theta(n/\log n)$ and $m - n = O(n/\log n)$. This corresponds to both a stronger worst-case security assumption, and a less useful LWE problem. By using larger errors, say, bounded by $s = n^\epsilon$ for some constant $\epsilon > 0$, it is possible to make both the worst-case lattice dimension k and number of extra samples $m - n$ into (small) linear functions of the LWE dimension n, which may be sufficient for some cryptographic applications of LWE. Specifically, for any constant $\epsilon < 1$, one may set $k = (\epsilon/3)n$ and $m = (1 + \epsilon/3)n$, which are easily verified to satisfy all the hypotheses of Theorem 6 when $q = k^{O(1)}$ is sufficiently large. These parameters correspond to $(\epsilon/3)n = \Omega(n)$ extra samples (beyond the LWE dimension n), and to the worst-case hardness of lattice problems in dimension $(\epsilon/3)n = \Omega(n)$. Notice that for $\epsilon < 1/2$, this version of LWE has much smaller errors than allowed by previous LWE hardness proofs, and it would be subject to subexponential-time attacks [2] if the number of samples were not restricted. Our result shows that if the number of samples is limited to $(1 + \epsilon/3)n$, then LWE maintains its provable security properties and conjectured exponential-time hardness in the dimension n.

One last instantiation allows for a linear number of samples $m = c \cdot n$ for any desired constant $c \geq 1$, which is enough for most applications of LWE in lattice cryptography. In this case we can choose (say) $k = n/2$, and it suffices to set the other parameters so that $s \geq (Cm)^{2c-1}$ and $q \geq (4s)^{c/(c-1)} \geq 4^{c/(c-1)} \cdot (Ccn)^{2c+1+1/(c-1)} = k^{O(1)}$. (We can also obtain better lower bounds on s and q by letting k be a smaller constant fraction of n.) This proves the hardness of LWE with uniform noise of polynomial magnitude $s = n^{O(1)}$, and any linear number of samples $m = O(n)$. Note that for $m = cn$, any instantiation of the parameters requires the magnitude s of the errors to be at least n^{c-1}. For $c > 3/2$, this is more noise than is typically used in the standard LWE problem, which allows errors of magnitude as small as $O(\sqrt{n})$, but requires them to be independent and follow a Gaussian-like distribution. The novelty in this last instantiation of Theorem 6 is that it allows for a much wider class of error distributions, including the uniform distribution, and distributions where different components of the error vector are correlated.

References

1. M. Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.

[4] Here we have not tried to optimize the value of q, and smaller values of the modulus are certainly possible: a close inspection of the proof of Theorem 6 reveals that for binary errors, the condition $q \geq 8n^{3c}$ can be replaced by $q \geq n^{c'}$ for any constant $c' > c$.

2. S. Arora and R. Ge. New algorithms for learning in presence of errors. In *ICALP (1)*, pages 403–415, 2011.
3. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, pages 719–737, 2012.
4. M. Bellare, E. Kiltz, C. Peikert, and B. Waters. Identity-based (lossy) trapdoor functions and applications. In *EUROCRYPT*, pages 228–245, 2012.
5. A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.
6. D. Boneh and D. M. Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In *PKC*, pages 1–16, 2011.
7. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584, 2013.
8. J.-Y. Cai and A. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *FOCS*, pages 468–477, 1997.
9. Y. Chen and P. Q. Nguyen. BKZ 2.0: Better lattice security estimates. In *ASIACRYPT*, pages 1–20, 2011.
10. D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *FOCS*, pages 580–589, 2011.
11. N. Döttling and J. Müller-Quade. Lossy codes and a new variant of the learning-with-errors problem. In *EUROCRYPT*, pages 18–34, 2013.
12. N. Gama and P. Q. Nguyen. Predicting lattice reduction. In *EUROCRYPT*, pages 31–51, 2008.
13. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
14. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.
15. R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In *CT-RSA*, pages 319–339, 2011.
16. D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM J. Comput.*, 34(1):118–169, 2004.
17. D. Micciancio. Duality in lattice cryptography. In *PKC*, 2010. Invited talk.
18. D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *CRYPTO*, pages 465–484, 2011.
19. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718, 2012.
20. D. Micciancio and C. Peikert. Hardness of sis and lwe with small parameters. *IACR Cryptology ePrint Archive*, 2013:69, 2013.
21. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Prelim. version in FOCS 2004.
22. D. Micciancio and O. Regev. Lattice-based cryptography. In *Post Quantum Cryptography*, pages 147–191. Springer, February 2009.
23. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342, 2009.
24. C. Peikert. An efficient and parallel Gaussian sampler for lattices. In *CRYPTO*, pages 80–97, 2010.
25. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196, 2008.
26. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.
27. D. Wagner. A generalized birthday problem. In *CRYPTO*, pages 288–303, 2002.