

Generating Shorter Bases for Hard Random Lattices*

Joël Alwen[†]
New York University

Chris Peikert[‡]
Georgia Institute of Technology

July 10, 2010

Abstract

We revisit the problem of generating a ‘hard’ random lattice together with a basis of relatively short vectors. This problem has gained in importance lately due to new cryptographic schemes that use such a procedure to generate public/secret key pairs. In these applications, a shorter basis corresponds to milder underlying complexity assumptions and smaller key sizes.

The contributions of this work are twofold. First, we simplify and modularize an approach originally due to Ajtai (ICALP 1999). Second, we improve the construction and its analysis in several ways, most notably by making the output basis asymptotically as short as possible.

Keywords: Lattices, average-case hardness, cryptography, Hermite normal form

*An edited version of this paper is to appear in Theory of Computing Systems. A preliminary version appeared in the 26th International Symposium on Theoretical Aspects of Computer Science (STACS 2009). This is the authors’ version.

[†]Work performed while at SRI International.

[‡]Much of this work was performed while at SRI International. This material is based upon work supported by the National Science Foundation under Grants CNS-0716786 and CNS-0749931. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

A (point) *lattice* is a discrete additive subgroup of \mathbb{R}^m ; alternatively, it is the set of all integer linear combinations of some linearly independent *basis* vectors $\mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^m$. Lattices appear to be a rich source of computational hardness, and in recent years, *cryptographic* schemes based on lattices have emerged as a promising alternative to more traditional ones based on, e.g., the factoring and discrete logarithm problems. Among other reasons, this is because lattice-based schemes have yet to be broken by efficient quantum algorithms (cf. [Sho97]), and their security can often be based merely on *worst-case* computational assumptions (rather than *average-case* assumptions, which are the norm in cryptography).

In 1996, Ajtai’s seminal work [Ajt96] in this area demonstrated a particular family of lattices for which, informally speaking, finding a short nonzero lattice vector in a randomly chosen lattice from the family is at least as hard as approximating some well-studied lattice problems in the *worst case*, i.e., for *any* lattice. This family of ‘hard random lattices’ has since been used as the foundation for several important cryptographic primitives, including one-way and collision-resistant hash functions, public-key encryption, digital signatures, and identity-based encryption (see, for example, [GGH96, MR04, Reg05, GPV08]).

Ajtai’s initial work also showed how to generate a hard random lattice together with knowledge of one relatively short nonzero lattice vector. The short vector can be useful as secret information in cryptographic applications; examples include an identification scheme [MV03] and public-key cryptosystems [Reg05, GPV08]. Shortly after Ajtai’s work, Goldreich, Goldwasser and Halevi [GGH97] proposed some public-key cryptographic schemes in which the secret key is an entire *short basis* of a public lattice, i.e., a basis in which all of the vectors are relatively short in Euclidean length. Their method for generating a lattice along with a short basis was ad-hoc, and unfortunately does not produce lattices from the provably hard family defined in [Ajt96]. Although the algorithm and cryptosystem were later improved [Mic01] (following a practical cryptanalysis of the original scheme for real-world parameters [Ngu99]), there is still no known proof that the induced random lattices are actually hard on the average. Therefore, the schemes from [GGH97] lack worst-case security proofs. (We also mention that the digital signature scheme from [GGH97] has since been shown to be insecure *regardless* of the particular method used for generating lattices [NR06].)

Following the GGH proposal [GGH97], Ajtai demonstrated an entirely different method of generating a lattice together with a short basis [Ajt99]. His algorithm has the important property that the resulting lattice is drawn, under the appropriate distribution, from the hard family defined in [Ajt96]. Interestingly, the algorithm apparently went without application until recently, when Gentry, Peikert and Vaikuntanathan [GPV08] constructed several provably secure (under worst-case assumptions) cryptographic schemes that crucially use short bases as their secret keys (see also [PVW08, PV08, Pei09, CHKP10, GHV10] for representative subsequent works). At this point we note that the algorithm of [Ajt99] actually produces a *full-rank set* of short lattice vectors (not necessarily a basis), which nonetheless suffices for all the applications in question.

In the above applications, the ‘quality’ of the short basis directly affects the concrete security and efficiency of the schemes, both in theory and in practice. More precisely, the quality is measured by the maximal Euclidean length of the basis vectors, or alternatively of their *Gram-Schmidt orthogonalization* (shorter means higher quality). The quality determines the approximation factor in the underlying worst-case lattice assumptions, as well as the concrete dimensions and key sizes needed for security against real attacks (see Section 2.3 for details). Therefore, it is very desirable to generate a basis that is as short as possible. Unfortunately, the construction from [Ajt99] is far from optimal — the maximum length of the basis vectors is bounded by $m^{5/2}$, whereas the optimum is about \sqrt{m} (for commonly used parameters) — and the method seems not to have attracted much attention or improvement since its publication a decade ago (probably due to the lack of applications until recently).

1.1 Our Contributions

Our first contribution is to elucidate and modularize Ajtai’s basic approach for generating a hard random lattice along with a relatively short basis. We endeavor to give a ‘top-down’ exposition of the key aspects of the problem and the techniques used to address them (in the process, we also correct some minor errors in the original paper).

One novelty in our approach is to base the algorithm and its analysis around the concept of the *Hermite normal form* (HNF), which is an easily computable, unique canonical representation of an integer lattice. Micciancio [Mic01] has proposed using the HNF in cryptographic applications to specify a lattice in its ‘least revealing’ representation; here we use other nice properties of the HNF to bound the dimension of the output lattice and the quality of the resulting basis.

Our second contribution is to refine the algorithm and its analysis, improving it in several ways. Most importantly, we improve the length of its output basis from $m^{5/2}$ to the asymptotically optimal $O(\sqrt{m})$, where m is the dimension of the output lattice (see Section 3 for precise statements of the new bounds). For the cryptographic schemes of, e.g., [GPV08], this immediately implies security under significantly milder worst-case assumptions: we need only that lattice problems are hard to approximate to within an $\tilde{O}(n^{3/2})$ factor, rather than $\tilde{O}(n^{7/2})$ as before.

We hasten to add that [GPV08, Section 5] briefly mentions that Ajtai’s algorithm can be improved to yield an $O(m^{1+\epsilon})$ bound on the basis length, but does not provide any further details. The focus of [GPV08] is on applications of a short basis, independent of the particular generation algorithm. The present work is a full exposition of an improved generation algorithm, and is meant to support and complement the schemes of [GPV08], and any other applications requiring a short basis.

1.2 Relation to Ajtai’s Construction

Our construction is inspired by Ajtai’s [Ajt99], but differs from it substantially in both the high-level structure and most of the details. The most significant similarity is a specially crafted unimodular matrix with small entries, which is used to ‘cancel out’ the necessarily large entries of another matrix that appears in the construction.

Departing from the approach of [Ajt99], our construction is guided from the ‘top down’ by two independent aspects of the construction: the block structure of the short output basis, and the probability distribution of the output lattice. This approach helps to illuminate the essential nature of the problem, and yields several technical simplifications. In particular, it lets us completely separate the *structural constraints* on the output lattice from its *randomization* (by contrast, in [Ajt99] the structure and randomization are tightly coupled).

2 Preliminaries

For a positive integer k , let $[k]$ denote the set $\{1, \dots, k\}$; $[0]$ is the empty set. We denote the set of integers modulo an integer $q \geq 1$ by \mathbb{Z}_q , and identify it with the set of integer residues $\{0, \dots, q - 1\}$ in the natural way. The base-2 logarithm is denoted \lg .

Column vectors are named by lower-case bold letters (e.g., \mathbf{x}) and matrices by upper-case bold letters (e.g., \mathbf{X}). The i th entry of a vector \mathbf{x} is denoted x_i , and the j th column of a matrix \mathbf{X} is denoted \mathbf{x}_j . We identify a matrix \mathbf{X} with the ordered set $\{\mathbf{x}_j\}$ of its column vectors, and define $\|\mathbf{X}\| = \max_j \|\mathbf{x}_j\|$. For $\mathbf{X} \in \mathbb{R}^{n \times m}$ and $\mathbf{Y} \in \mathbb{R}^{n \times m'}$ having an equal number of rows, $[\mathbf{X}|\mathbf{Y}] \in \mathbb{R}^{n \times (m+m')}$ denotes the concatenation of the columns of \mathbf{X} followed by the columns of \mathbf{Y} . Likewise, for $\mathbf{X} \in \mathbb{R}^{n \times m}$ and $\mathbf{Y} \in \mathbb{R}^{n' \times m}$ having an equal number of columns, $[\mathbf{X}; \mathbf{Y}] \in \mathbb{R}^{(n+n') \times m}$ is the concatenation of the rows of \mathbf{X} and the rows of \mathbf{Y} .

We let \mathbf{e}_i denote the i th standard basis vector, where its dimension will be clear from context. The $d \times d$ identity matrix is denoted \mathbf{I}_d ; we omit its dimension when it is clear from context. We denote the (Euclidean) unit sphere in \mathbb{R}^m by S^{m-1} , i.e., $S^{m-1} = \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\| = 1\}$.

2.1 Matrix Decompositions

For an ordered set $\mathbf{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_m\} \subset \mathbb{R}^m$ of linearly independent vectors, the *Gram-Schmidt orthogonalization* $\tilde{\mathbf{S}}$ of \mathbf{S} is defined iteratively as follows: $\tilde{\mathbf{s}}_1 = \mathbf{s}_1$, and for $j = 2, \dots, m$, $\tilde{\mathbf{s}}_j$ is the component of \mathbf{s}_j orthogonal to $\text{span}(\mathbf{s}_1, \dots, \mathbf{s}_{j-1})$, i.e., $\tilde{\mathbf{s}}_j = \mathbf{s}_j - \sum_{i \in [j-1]} \tilde{\mathbf{s}}_i \cdot \langle \mathbf{s}_j, \tilde{\mathbf{s}}_i \rangle / \langle \tilde{\mathbf{s}}_i, \tilde{\mathbf{s}}_i \rangle$.

For a matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$, a *singular value decomposition* is a factorization $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{-1}$ where $\mathbf{U} \in \mathbb{R}^{m \times m}$, $\mathbf{V} \in \mathbb{R}^{n \times n}$ are orthogonal square matrices and $\mathbf{\Sigma} \in \mathbb{R}^{m \times n}$ is diagonal with nonnegative entries. The diagonal entries of $\mathbf{\Sigma}$ called the *singular values* of \mathbf{M} , and are unique up to order. By definition, it follows that the largest (respectively, smallest) singular value of \mathbf{M} is the maximum (respectively, minimum) value of $\|\mathbf{M}\mathbf{x}\|$ over all $\mathbf{x} \in S^{n-1}$. Note also that the singular values of \mathbf{M} and \mathbf{M}^t are the same.

2.2 Probability

For two probability distributions D_1, D_2 (viewed as functions) over a finite set G , the statistical distance $\Delta(D_1, D_2)$ is defined to be $\frac{1}{2} \sum_{g \in G} |D_1(g) - D_2(g)|$. It is easy to see that statistical distance is a metric; in particular, it obeys the triangle inequality. We say that a distribution D (or a random variable having distribution D) is ϵ -uniform if its statistical distance from the uniform distribution over G is at most ϵ .

2.2.1 Hashing

Let \mathcal{X} and \mathcal{Y} be two finite domains. A family \mathcal{H} of functions mapping \mathcal{X} to \mathcal{Y} is *2-universal* if for all distinct $x, x' \in \mathcal{X}$, $\Pr_{h \leftarrow \mathcal{H}}[h(x) = h(x')] = 1/|\mathcal{Y}|$.

Lemma 2.1 (Simplified Leftover Hash Lemma [HILL99]). *Let \mathcal{H} be a family of 2-universal hash functions from a domain \mathcal{X} to range \mathcal{Y} . Then for $h \leftarrow \mathcal{H}$ and $x \leftarrow \mathcal{X}$ chosen uniformly and independently, $(h, h(x))$ is $\frac{1}{2} \sqrt{|\mathcal{Y}|/|\mathcal{X}|}$ -uniform over $\mathcal{H} \times \mathcal{Y}$.*

Let G be any finite, abelian, additive group, and let $m \geq 1$ be an integer. For $\mathbf{g} \in G^m$, define $h_{\mathbf{g}}: \{0, 1\}^m \rightarrow G$ as $h_{\mathbf{g}}(\mathbf{x}) = \sum_{i \in [m]} x_i g_i$. The family $\mathcal{H} = \{h_{\mathbf{g}}\}_{\mathbf{g} \in G^m}$ is 2-universal: for any distinct $\mathbf{x}, \mathbf{x}' \in \{0, 1\}^m$, there exists some $i \in [m]$ such that $x_i - x'_i = \pm 1$; by conditioning on any fixed values of $g_j \in G$ for $j \neq i$ and averaging over the choice of g_i , we have $\Pr_{\mathbf{g} \leftarrow G^m}[h_{\mathbf{g}}(\mathbf{x}) = h_{\mathbf{g}}(\mathbf{x}')] = 1/|G|$. Therefore by Lemma 2.1, $(\mathbf{g}, h_{\mathbf{g}}(\mathbf{x}))$ is $\frac{1}{2} \sqrt{|G|/2^m}$ -uniform over the choice of uniformly random and independent $\mathbf{g} \leftarrow G^m$ and $\mathbf{x} \leftarrow \{0, 1\}^m$.

For various reasons, it will be important for us to work with balanced (mean zero), rather than binary (zero-one), random variables. Extend $h_{\mathbf{g}}$ to have domain $\{-1, 0, 1\}^m$, and let the entries of \mathbf{x} be independent and chosen to be 0 with probability $\frac{1}{2}$, and 1 and -1 each with probability $\frac{1}{4}$. Then $(\mathbf{g}, h_{\mathbf{g}}(\mathbf{x}))$ is again $\frac{1}{2} \sqrt{|G|/2^m}$ -uniform, because \mathbf{x} may be seen as the difference between two independent uniformly random variables $\mathbf{x}', \mathbf{x}'' \leftarrow \{0, 1\}^m$, and $h_{\mathbf{g}}(\mathbf{x}) = h_{\mathbf{g}}(\mathbf{x}') - h_{\mathbf{g}}(\mathbf{x}'')$. (Note that we choose not to work with Bernoulli ± 1 random variables, because $\mathcal{H} = \{h_{\mathbf{g}}\}$ is not necessarily 2-universal on the domain $\{\pm 1\}$).

Finally, by the triangle inequality for statistical distance, we have that $(h_{\mathbf{g}}, h_{\mathbf{g}}(\mathbf{x}_1), \dots, h_{\mathbf{g}}(\mathbf{x}_k))$ is $\frac{k}{2} \sqrt{|G|/2^m}$ -uniform for independent $h_{\mathbf{g}} \leftarrow \mathcal{H}$ and $\mathbf{x}_1, \dots, \mathbf{x}_k$ chosen from either of the above distributions.

2.2.2 Subgaussian Random Variables and Matrices

We say that a random variable X is *subgaussian* with *parameter* $s > 0$ (sometimes called the *subgaussian moment*) if $\Pr[|X| > t] \leq 2 \exp(-t^2/s^2)$ for all $t \geq 0$. In particular, any bounded random variable is subgaussian. The following is a standard fact about subgaussian random variables; see, e.g., [Ver07, Lecture 5] for a proof.

Fact 2.2. *Let X_1, \dots, X_k be independent mean-zero subgaussian random variables with parameter s , and let $\mathbf{u} \in \mathbb{R}^k$ be arbitrary. Then $\sum_{i \in [k]} u_i X_i$ is subgaussian with parameter $s \cdot \|\mathbf{u}\|$.*

There is a well-developed theory for bounding the singular values of random matrices with independent entries (which need not be identically distributed). The following lemma is folklore in the area; see, e.g., [Ver07, Lecture 6] for a proof.

Lemma 2.3. *Let $\mathbf{X} \in \mathbb{R}^{m \times n}$ be a matrix whose entries are independent subgaussian random variables with parameter s . There exists a universal constant $C > 0$ such that the largest singular value of \mathbf{X} is at most $C \cdot s \cdot (\sqrt{m} + \sqrt{n})$, except with probability $2^{-\Omega(m+n)}$.*

2.3 Lattices

Generally defined, a *lattice* Λ is a discrete additive subgroup of \mathbb{R}^m . In this work, we are concerned only with *full-rank integer lattices*, which are discrete additive subgroups of \mathbb{Z}^m having finite index, i.e., the quotient group \mathbb{Z}^m/Λ is finite. The determinant of Λ , denoted $\det(\Lambda)$, is the cardinality $|\mathbb{Z}^m/\Lambda|$ of this quotient group. Geometrically, the determinant is a measure of the ‘sparsity’ of the lattice.

A lattice $\Lambda \subseteq \mathbb{Z}^m$ can also be viewed as the set of all integer linear combinations of m linearly independent *basis* vectors $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_m\} \subset \mathbb{Z}^m$:

$$\Lambda = \mathcal{L}(\mathbf{B}) = \left\{ \mathbf{B}\mathbf{c} = \sum_{i \in [m]} c_i \mathbf{b}_i : \mathbf{c} \in \mathbb{Z}^m \right\}.$$

A lattice has infinitely many bases (when $m \geq 2$), which are related to each other by unimodular transformations, i.e., \mathbf{B} and \mathbf{B}' generate the same lattice if and only if $\mathbf{B} = \mathbf{B}' \cdot \mathbf{U}$ for some unimodular $\mathbf{U} \in \mathbb{Z}^{m \times m}$. The determinant of any basis matrix \mathbf{B} coincides with the determinant of the lattice it generates, up to sign: $|\det(\mathbf{B})| = \det(\mathcal{L}(\mathbf{B}))$.

Every lattice $\Lambda \subseteq \mathbb{Z}^m$ has a *unique* canonical basis $\mathbf{H} = \text{HNF}(\Lambda) \in \mathbb{Z}^{m \times m}$ called its *Hermite normal form* (HNF). The matrix \mathbf{H} is upper triangular and has non-negative entries (i.e., $h_{i,j} \geq 0$ with equality for $i > j$), has strictly positive diagonals (i.e., $h_{i,i} \geq 1$ for every i), and every entry above the diagonal is strictly smaller than the diagonal entry in its row (i.e., $h_{i,j} < h_{i,i}$ for $i < j$). Note that because \mathbf{H} is upper triangular, its determinant is simply the product $\prod_{i \in [m]} h_{i,i} > 0$ of the diagonal entries. For a lattice basis \mathbf{B} , we write $\text{HNF}(\mathbf{B})$ to denote $\text{HNF}(\mathcal{L}(\mathbf{B}))$. Given an arbitrary basis \mathbf{B} , $\mathbf{H} = \text{HNF}(\mathbf{B})$ can be computed in polynomial time (see [MW01] and references therein).

2.4 Hard Random Lattices

We will be especially concerned with a certain family of lattices in \mathbb{Z}^m as defined by Ajtai [Ajt96]. A lattice from this family is most naturally specified not by a basis, but instead by a *parity check* matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$

for some positive integer n and positive integer modulus q . (We discuss the parameters n , q , and m in detail below; see also the survey [MR09]). The lattice associated with \mathbf{A} is defined as

$$\Lambda^\perp(\mathbf{A}) = \left\{ \mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\mathbf{x} = \sum_{j \in [m]} x_j \cdot \mathbf{a}_j = \mathbf{0} \in \mathbb{Z}_q^n \right\} \subseteq \mathbb{Z}^m.$$

It is routine to check that $\Lambda^\perp(\mathbf{A})$ contains the identity $\mathbf{0} \in \mathbb{Z}^m$ and is closed under negation and addition, hence it is a subgroup of (and lattice in) \mathbb{Z}^m . Also observe that $\Lambda^\perp(\mathbf{A})$ is ‘ q -ary,’ that is, $q \cdot \mathbb{Z}^m \subseteq \Lambda^\perp(\mathbf{A})$ for every \mathbf{A} , so membership in $\Lambda^\perp(\mathbf{A})$ is determined solely by an integer vector’s entries modulo q .

2.4.1 Hermite Normal Form

Let $\mathbf{H} \in \mathbb{Z}^{m \times m}$ be the Hermite normal form of a lattice $\Lambda = \Lambda^\perp(\mathbf{A})$ for some arbitrary parity check matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. Given \mathbf{A} , the matrix \mathbf{H} may be computed efficiently (e.g., by first computing a basis of Λ). In one of our constructions, we use the fact that every diagonal entry of \mathbf{H} is at most q , which we now prove.

We can determine \mathbf{H} as follows. Starting with the first column $\mathbf{h}_1 = h_{1,1} \cdot \mathbf{e}_1 \in \Lambda$, it must be the case that

$$\mathbf{A} \cdot \mathbf{h}_1 = h_{1,1} \cdot \mathbf{a}_1 = \mathbf{0} \in \mathbb{Z}_q^n.$$

Let $k \leq q$ be the smallest positive integer solution to $k \cdot \mathbf{a}_1 = \mathbf{0} \in \mathbb{Z}_q^n$. Then $k \cdot \mathbf{e}_1 \in \Lambda$, so we must be able to write $k \cdot \mathbf{e}_1 = \mathbf{H}\mathbf{z}$ for some $\mathbf{z} \in \mathbb{Z}^m$. Now because every diagonal $h_{i,i} > 0$ and \mathbf{H} is upper triangular, it must be the case that $z_i = 0$ for all $i > 1$. This implies that $z_1 \cdot h_{1,1} = k$, and because $0 < k \leq h_{1,1}$, we must have $z_1 = 1$ and thus $h_{1,1} = k \leq q$.

More generally, suppose that $\mathbf{h}_1, \dots, \mathbf{h}_{j-1}$ are determined for some $j \in [m]$. Then by similar reasoning as above, $\mathbf{h}_j \in \mathbb{Z}^m$ is given by the unique solution to the equation

$$h_{j,j} \cdot \mathbf{a}_j + \sum_{i \in [j-1]} h_{i,j} \cdot \mathbf{a}_i = \mathbf{0} \in \mathbb{Z}_q^n$$

in which $h_{j,j} > 0$ is minimized and $0 \leq h_{i,j} < h_{i,i} \leq q$ for every $i < j$. In particular, $q \cdot \mathbf{e}_j$ is a solution to the above relation, hence $h_{j,j} \leq q$. We conclude by induction that every diagonal entry of \mathbf{H} is at most q .

2.4.2 Geometric Facts

Let $\Lambda = \Lambda^\perp(\mathbf{A})$ for some arbitrary $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. First, we have $\det(\Lambda) \leq q^n$, by the following argument: let $\phi : (\mathbb{Z}^m / \Lambda) \rightarrow \mathbb{Z}_q^n$ be the homomorphism mapping the residue class $(\mathbf{x} + \Lambda)$ to $\mathbf{A}\mathbf{x} \in \mathbb{Z}_q^n$. Then ϕ is injective, because if $\phi(\mathbf{x} + \Lambda) = \phi(\mathbf{x}' + \Lambda)$ for some $\mathbf{x}, \mathbf{x}' \in \mathbb{Z}^m$, we have $\mathbf{A}(\mathbf{x} - \mathbf{x}') = \mathbf{0}$ which implies $\mathbf{x} - \mathbf{x}' \in \Lambda$, i.e., $\mathbf{x} = \mathbf{x}' \pmod{\Lambda}$. Therefore, there are at most $|\mathbb{Z}_q^n| = q^n$ residue classes in \mathbb{Z}^m / Λ . Minkowski’s first inequality states that the minimum distance of Λ (i.e., the length of a shortest nonzero lattice vector) is at most

$$\sqrt{m} \cdot \det(\Lambda)^{1/m} \leq \sqrt{m} \cdot q^{n/m}. \quad (2.1)$$

For reasons related to Proposition 2.4 below, the family of lattices under discussion is most naturally parameterized by n (even though m is the lattice dimension), and the parameters $q = q(n)$ and $m = m(n)$ are viewed as functions of n . Given n and $q = q(n)$, a typical choice of the parameter m , which essentially minimizes the bound in (2.1), is $m = c \cdot n \lg q$ for some constant $c > 0$. Then by (2.1), the minimum distance of $\Lambda^\perp(\mathbf{A})$ for any $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is at most

$$\sqrt{m} \cdot q^{n/m} = \sqrt{m} \cdot q^{1/(c \lg q)} = \sqrt{m} \cdot 2^{1/c} = O(\sqrt{n \log q}).$$

For the above parameters, a simple counting argument shows that the above bound on $\lambda_1(\Lambda^\perp(\mathbf{A}))$ is asymptotically tight, with high probability over the uniformly random choice of \mathbf{A} . For simplicity, suppose that q is prime. (With a bit more care, the argument can be extended to composite q as well.) Then for any fixed nonzero $\mathbf{z} \in \mathbb{Z}^m$, the probability over the choice of \mathbf{A} that $\mathbf{z} \in \Lambda^\perp(\mathbf{A})$, i.e., that $\mathbf{A}\mathbf{z} = \mathbf{0} \in \mathbb{Z}_q^n$, is exactly q^{-n} . Then as long as

$$N_{\alpha,m} := |\{\mathbf{z} \in \mathbb{Z}^m : \|\mathbf{z}\| \leq \sqrt{\alpha m}\}| \leq q^{n/2}$$

for some constant $\alpha > 0$, a union bound implies that $\lambda_1(\Lambda^\perp(\mathbf{A})) \geq \sqrt{\alpha m} = \Omega(\sqrt{n \log q})$ except with probability $q^{-n/2}$.

To bound $N_{\alpha,m}$, we use a result of Mazo and Odlyzko [MO90]. An immediate consequence of [MO90, Lemma 1] is that for any constant $\delta > 1$, there exists a constant $\alpha > 0$ such that $N_{\alpha,m} \leq \delta^m$. The desired bound holds by choosing $\delta = q^{n/2m} = 2^{1/2c} > 1$. For larger choices of $m = c(n) \cdot n \lg q$ where $c(n) = \omega(1)$, a more refined analysis using the Mazo-Odlyzko bound shows that the minimum distance remains bounded from below by $\Omega(\sqrt{n \log q} / \log c(n))$, and from above by $O(\sqrt{n \log q})$ because we can simply ignore the extra columns of \mathbf{A} .

2.4.3 Average-Case Hardness

The following proposition, proved first by Ajtai [Ajt96] (in a quantitatively weaker form) and in its current form in [MR04, GPV08], relates the average-case and worst-case complexity of certain lattice problems.

Proposition 2.4 (Informal). *For any $m = m(n)$, $\beta = \beta(n) = \text{poly}(n)$ and any $q = q(n) \geq \beta \cdot \omega(\sqrt{n \log n})$, finding a nonzero $\mathbf{x} \in \Lambda^\perp(\mathbf{A})$ having length at most β for uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ (with at least $1/\text{poly}(n)$ probability over the choice of \mathbf{A} and the randomness of the algorithm) is at least as hard as approximating several lattice problems on n -dimensional lattices to within a $\gamma(n) = \beta \cdot \tilde{O}(\sqrt{n})$ factor in the worst case.*

Note that Proposition 2.4 is meaningful only when β is at least the typical minimum distance of $\Lambda^\perp(\mathbf{A})$ for uniformly random \mathbf{A} . For $m = c \cdot n \lg q$ as described above, we can therefore take β to be as small as $O(\sqrt{n \lg n})$, which yields a hard-on-average problem assuming the worst-case hardness of approximating lattice problems to within an $\tilde{O}(n)$ factor.

In certain cryptographic applications, however, an adversary that breaks a cryptographic scheme is guaranteed only to produce a lattice vector whose length is substantially more than the minimum distance, so one needs average-case hardness for larger values of β . For example, the secret key in the digital signature schemes of [GPV08] is a basis of $\Lambda^\perp(\mathbf{A})$ having some length L , and signatures are vectors of length $\approx L\sqrt{m}$. It is shown that a signature forger may be used to find a nonzero lattice vector of length $\beta \approx L\sqrt{m}$ in $\Lambda^\perp(\mathbf{A})$, which by Proposition 2.4 (for our choice of m) is as hard as approximating lattice problems in the worst case to within $L \cdot \tilde{O}(n)$ factors. Therefore, using a shorter secret basis in the signature scheme has the immediate advantage of a weaker underlying hardness assumption.

Note also that Proposition 2.4 requires the modulus q to exceed β (otherwise $q \cdot \mathbf{e}_1$ would trivially be a valid solution), and that m grows with $\lg q$. Therefore, a polynomial factor improvement in the length L also yields a constant factor improvement in the dimension m and modulus q , which translates to a constant factor improvement in the size of the public key \mathbf{A} (all other variables remaining the same).

3 Constructions

We give two algorithms for constructing a hard random lattice together with a relatively short basis. Strictly speaking, our two constructions are incomparable. The first is relatively simple and gives a guaranteed bound on the basis quality, but is slightly suboptimal in either the lattice dimension or basis length. Our second construction is more involved, but it is simultaneously optimal (up to constant factors) in both the lattice dimension and another useful measure of quality.

Theorem 3.1. *Let $\delta > 0$ be any fixed constant. There is a probabilistic polynomial-time algorithm that, on input positive integers n (in unary), $q, r \geq 2$ (in binary), and $m \geq (1 + \delta)(1 + \lceil \lg_r q \rceil) \cdot n \lg q$ (in unary), outputs $(\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{S} \in \mathbb{Z}^{m \times m})$ such that:*

- \mathbf{A} is $(m \cdot q^{-\delta n/2})$ -uniform over $\mathbb{Z}_q^{n \times m}$,
- \mathbf{S} is a basis of $\Lambda^\perp(\mathbf{A})$, and
- $\|\mathbf{S}\| \leq 2r\sqrt{m}$.

Setting $r = 2$ in the above theorem, the algorithm generates a basis of length $O(\sqrt{m}) = O(\sqrt{n \log^2 q})$ for a random lattice having dimension $m = O(n \log^2 q)$. These quantities are larger than our ultimate goal by $O(\sqrt{\log q})$ and $O(\log q)$ factors, respectively. Alternatively, if $q = \text{poly}(n)$, we may set $r = n^\epsilon$ for some small constant $\epsilon > 0$, which implies $\log_r q = O(1)$. In this case, the algorithm generates a basis of only slightly suboptimal length $O(n^\epsilon \cdot \sqrt{n \log q})$ for a random lattice having dimension $m = O(n \log q)$.

Our next construction *simultaneously* optimizes the lattice dimension and basis quality, when the quality is measured according to the *Gram-Schmidt orthogonalization* of the basis. As explained in the introduction, this measure of quality is appropriate for all known applications. The somewhat large constant factor in the lower bound for m is a consequence of the theorem's generality, and can be improved in specific cases, such as when q is a prime.

Theorem 3.2. *Let $\delta > 0$ be any fixed constant. There is a probabilistic polynomial-time algorithm that, on input positive integers n (in unary), $q \geq 2$ (in binary), and $m \geq (5 + 3\delta) \cdot n \lg q$ (in unary), outputs $(\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{S} \in \mathbb{Z}^{m \times m})$ such that:*

- \mathbf{A} is $(m \cdot q^{-\delta n/2})$ -uniform over $\mathbb{Z}_q^{n \times m}$,
- $\|\mathbf{S}\| = O(n \log q)$ with probability $1 - 2^{-\Omega(n)}$, and
- $\|\tilde{\mathbf{S}}\| = O(\sqrt{n \log q})$ with probability $1 - 2^{-\Omega(n)}$.

3.1 Common Approach

Here we describe the common framework, specified in Algorithm 1, that underlies the two concrete constructions from Theorems 3.1 and 3.2. (The details of each construction are given below in Sections 3.2 and 3.3, respectively.)

Let $m = m_1 + m_2$ for some sufficiently large dimensions m_1, m_2 . The algorithm is given a uniformly random matrix $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times m_1}$ as input, and extends \mathbf{A}_1 to $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2] \in \mathbb{Z}_q^{n \times m}$ by generating $\mathbf{A}_2 \in \mathbb{Z}_q^{n \times m_2}$ together with some short basis $\mathbf{S} \in \mathbb{Z}^{m \times m}$ of $\Lambda^\perp(\mathbf{A})$.

Algorithm 1 Framework for constructing $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and basis \mathbf{S} of $\Lambda^\perp(\mathbf{A})$.

Input: $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times m_1}$ and dimension m_2 (in unary).

Output: $\mathbf{A}_2 \in \mathbb{Z}_q^{n \times m_2}$ and basis \mathbf{S} of $\Lambda^\perp(\mathbf{A})$, where $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2] \in \mathbb{Z}_q^{n \times m}$ for $m = m_1 + m_2$.

- 1: Generate component matrices $\mathbf{U} \in \mathbb{Z}^{m_2 \times m_2}$; $\mathbf{G}, \mathbf{R} \in \mathbb{Z}^{m_1 \times m_2}$; $\mathbf{P} \in \mathbb{Z}^{m_2 \times m_1}$; and $\mathbf{C} \in \mathbb{Z}^{m_1 \times m_1}$ such that \mathbf{U} is nonsingular and $(\mathbf{G}\mathbf{P} + \mathbf{C}) \subset \Lambda^\perp(\mathbf{A}_1)$, e.g., as described in Section 3.2 or 3.3.
 - 2: Let $\mathbf{A}_2 = -\mathbf{A}_1 \cdot (\mathbf{R} + \mathbf{G}) \in \mathbb{Z}_q^{n \times m_2}$.
 - 3: Let $\mathbf{S} = \begin{pmatrix} (\mathbf{G} + \mathbf{R})\mathbf{U} & \mathbf{R}\mathbf{P} - \mathbf{C} \\ \mathbf{U} & \mathbf{P} \end{pmatrix} \in \mathbb{Z}^{m \times m}$.
 - 4: **return** \mathbf{A}_2 and \mathbf{S} .
-

$$n \left\{ \left[\begin{array}{c|c} \mathbf{A}_1 & \mathbf{A}_2 \\ \hline \end{array} \right] \begin{array}{c} \left[\begin{array}{c|c} (\mathbf{G} + \mathbf{R})\mathbf{U} & \mathbf{R}\mathbf{P} - \mathbf{C} \\ \hline \mathbf{U} & \mathbf{P} \end{array} \right] \\ \left. \begin{array}{l} \underbrace{\hspace{1.5cm}}_{m_2} \quad \underbrace{\hspace{1.5cm}}_{m_1} \end{array} \right\} \begin{array}{l} \left. \begin{array}{l} \hspace{1.5cm} \end{array} \right\} m_1 \\ \left. \begin{array}{l} \hspace{1.5cm} \end{array} \right\} m_2 \end{array} \right\} = \mathbf{0} \in \mathbb{Z}_q^{n \times m}.$$

Figure 1: Block structure of the equation $\mathbf{AS} = \mathbf{0} \in \mathbb{Z}_q^{n \times m}$.

The output matrix \mathbf{S} has a block structure as shown in Figure 1, which uses four main component matrices \mathbf{U} , \mathbf{G} , \mathbf{P} , and \mathbf{R} that are provided by an instantiation of the framework. (The fifth matrix \mathbf{C} is inessential to the basic construction, and is included only for some extra flexibility later on; for now we may take $\mathbf{C} = \mathbf{0}$.) The components are named according to their essential properties:

- \mathbf{U} is nonsingular (invertible over the reals) and typically unimodular;
- \mathbf{G} typically has entries that grow geometrically (from left to right);
- \mathbf{P} ‘picks out’ certain columns of \mathbf{G} via the matrix product \mathbf{GP} ;
- \mathbf{R} is a random, typically ‘short’ matrix with an appropriate distribution (e.g., random $0, \pm 1$ entries).

In both of our constructions, all of the components except \mathbf{R} are constructed deterministically (depending on the input \mathbf{A}_1), and the desired near-uniform distribution of $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2]$ follows from the uniformity of \mathbf{A}_1 and the random choice of \mathbf{R} (via the leftover hash lemma). The utility of \mathbf{S} ’s particular block structure will become clear as we see how it allows for satisfying the various constraints on the component matrices.

First, consider the requirement that $\mathbf{S} \subset \Lambda^\perp(\mathbf{A})$, i.e., $\mathbf{AS} = \mathbf{0} \in \mathbb{Z}_q^{n \times m}$. We need to satisfy

$$\mathbf{A}_1 \cdot (\mathbf{G} + \mathbf{R})\mathbf{U} + \mathbf{A}_2 \cdot \mathbf{U} = \mathbf{0} \in \mathbb{Z}_q^{n \times m_2} \quad (3.1)$$

$$\mathbf{A}_1 \cdot (\mathbf{RP} - \mathbf{C}) + \mathbf{A}_2 \cdot \mathbf{P} = \mathbf{0} \in \mathbb{Z}_q^{n \times m_1}. \quad (3.2)$$

We can immediately satisfy Equation (3.1) by letting

$$\mathbf{A}_2 = -\mathbf{A}_1 \cdot (\mathbf{G} + \mathbf{R}) \in \mathbb{Z}_q^{n \times m_2}. \quad (3.3)$$

(Indeed, if \mathbf{U} is unimodular then this choice of \mathbf{A}_2 is necessary, because \mathbf{U} can be cancelled out of Equation (3.1).) Note that for uniformly random \mathbf{A}_1 and a suitable random choice of \mathbf{R} (independent of \mathbf{G}), the matrix $[\mathbf{A}_1 | \mathbf{A}_1 \mathbf{R}]$ will be close to uniformly random by the leftover hash lemma, hence so will the parity-check matrix $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2] = [\mathbf{A}_1 | -\mathbf{A}_1(\mathbf{G} + \mathbf{R})]$.

Next, substituting Equation (3.3) into Equation (3.2) and rearranging, we obtain the constraint

$$\mathbf{A}_1 \cdot (\mathbf{GP} + \mathbf{C}) = \mathbf{0} \in \mathbb{Z}_q^{n \times m_1}. \quad (3.4)$$

That is, we need $(\mathbf{GP} + \mathbf{C}) \subset \Lambda^\perp(\mathbf{A}_1)$. Lemma 3.3 below shows that in order for \mathbf{S} to be nonsingular, $\mathbf{GP} + \mathbf{C}$ may be any basis or full-rank subset of $\Lambda^\perp(\mathbf{A}_1)$. We will typically use the Hermite normal form basis $\text{HNF}(\Lambda^\perp(\mathbf{A}_1))$, due to its nice properties (specifically, efficient computability and bounded entries).

Lemma 3.3 (Correctness of Algorithm 1). *Adopt the notation and hypotheses of Algorithm 1. Then if $\mathbf{GP} + \mathbf{C} \subset \Lambda^\perp(\mathbf{A}_1)$, we have $\mathbf{S} \subset \Lambda^\perp(\mathbf{A})$. Moreover, \mathbf{S} is a basis (respectively, full-rank subset) of $\Lambda^\perp(\mathbf{A})$ if and only if $\mathbf{GP} + \mathbf{C}$ is a basis (resp., full-rank subset) of $\Lambda^\perp(\mathbf{A}_1)$.*

Proof. By the above discussion, we have $\mathbf{S} \subset \Lambda^\perp(\mathbf{A})$ if $\mathbf{GP} + \mathbf{C} \subset \Lambda^\perp(\mathbf{A}_1)$. Now because \mathbf{U} is unimodular, it is invertible. Using the formula $|\det \begin{pmatrix} \mathbf{W} & \mathbf{X} \\ \mathbf{Y} & \mathbf{Z} \end{pmatrix}| = |\det(\mathbf{X} - \mathbf{WY}^{-1}\mathbf{Z})|$ (for square invertible \mathbf{Y}) for the determinant of a block matrix, we have

$$|\det(\mathbf{S})| = |\det((\mathbf{RP} - \mathbf{C}) - (\mathbf{G} + \mathbf{R})\mathbf{U} \cdot \mathbf{U}^{-1} \cdot \mathbf{P})| = |\det(\mathbf{GP} + \mathbf{C})|.$$

Therefore, $\mathbf{GP} + \mathbf{C}$ is full-rank (nonsingular) if and only if \mathbf{S} is full-rank. To see when \mathbf{S} is a *basis* of $\Lambda^\perp(\mathbf{A})$, observe that the additive subgroup $\mathbb{G} \subseteq \mathbb{Z}_q^n$ generated by the columns of \mathbf{A}_1 is exactly the subgroup generated

by the columns of $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2]$, because the columns of $\mathbf{A}_2 = -\mathbf{A}_1(\mathbf{G} + \mathbf{R})$ are in \mathbb{G} by construction. Therefore,

$$\det(\Lambda^\perp(\mathbf{A})) = |\mathbb{G}| = \det(\Lambda^\perp(\mathbf{A}_1)).$$

Thus $|\det(\mathbf{S})| = \det(\Lambda^\perp(\mathbf{A}))$ — i.e., \mathbf{S} is a basis of $\Lambda^\perp(\mathbf{A})$ — exactly when $|\det(\mathbf{GP} + \mathbf{C})| = \det(\Lambda^\perp(\mathbf{A}_1))$ — i.e., $\mathbf{GP} + \mathbf{C}$ is a basis of $\Lambda^\perp(\mathbf{A}_1)$. \square

The remaining main constraint is that \mathbf{S} must be relatively short. This presents a dilemma: clearly \mathbf{P} must be short, but we need the columns of \mathbf{GP} to be nontrivial vectors in $\Lambda^\perp(\mathbf{A}_1)$, and it is hard to find short nonzero vectors in this lattice. (Here we are assuming for simplicity that $\mathbf{C} = \mathbf{0}$; in any case, \mathbf{C} needs to be short because \mathbf{R} is short as well.) Therefore, at least some of the columns of \mathbf{G} should be ‘long.’ At the same time, \mathbf{GU} must be short because it appears in \mathbf{S} as part of the block $(\mathbf{G} + \mathbf{R})\mathbf{U}$, and because both \mathbf{U} and \mathbf{R} are short.

The dilemma may be resolved by a judicious choice of the \mathbf{G} and \mathbf{U} matrices. We construct \mathbf{G} so that its columns grow geometrically to include long vectors that are themselves in $\Lambda^\perp(\mathbf{A}_1)$, or that have known small combinations belonging to $\Lambda^\perp(\mathbf{A}_1)$. This makes it easy to construct a short \mathbf{P} so that $\mathbf{GP} \subset \Lambda^\perp(\mathbf{A}_1)$. We also construct a short nonsingular matrix \mathbf{U} so that \mathbf{GU} is short. This is possible because the small entries of \mathbf{U} can cancel adjacent columns of \mathbf{G} to always yield short vectors. For example, the entries in the j th column of \mathbf{G} can be 2^j , while \mathbf{U} can simply have 1s along the diagonal and -2 s above the diagonal.

The remainder of the paper is dedicated to concrete instantiations of Algorithm 1, and to analyzing the quality of \mathbf{S} for the particular constructions.

3.2 First Construction

We begin with a relatively simple instantiation of Algorithm 1. Its properties are summarized in the following lemma, of which Theorem 3.1 is an immediate corollary.

Lemma 3.4. *Let $\delta > 0$ be any fixed constant. There is a probabilistic polynomial-time algorithm that, given uniformly random $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times m_1}$ for any $m_1 \geq d = (1 + \delta)n \lg q$, an integer $r \geq 2$, and any integer $m_2 \geq m_1 \cdot \ell$ (in unary) where $\ell = \lceil \log_r q \rceil$, outputs matrices $\mathbf{U}, \mathbf{G}, \mathbf{R}, \mathbf{P}$, and $\mathbf{C} = \mathbf{I}$ as required by Step 1 of Algorithm 1 such that:*

- $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2]$ is $(m_2 \cdot q^{-\delta n/2})$ -uniform, where \mathbf{A}_2 is as in Step 2 of Algorithm 1.
- $\|\mathbf{S}\| \leq 2r\sqrt{m_1 + 1}$, where \mathbf{S} is as in Step 3 of Algorithm 1.

The remainder of this subsection consists of the proof of Lemma 3.4.

3.2.1 Construction

Given \mathbf{A}_1 , let $\mathbf{H} \in \mathbb{Z}^{m_1 \times m_1}$ be the Hermite normal form of $\Lambda^\perp(\mathbf{A}_1)$. The basic idea of the construction is that \mathbf{G} itself contains the m_1 columns of $\mathbf{H}' = \mathbf{H} - \mathbf{I}$ (among many others), and \mathbf{P} simply selects those columns to yield $\mathbf{GP} = \mathbf{H}'$. To ensure a short unimodular \mathbf{U} such that \mathbf{GU} is also short, we include additional columns in \mathbf{G} that increase geometrically (with base r) to the desired columns of \mathbf{H}' ; this is the reason for the extra $\ell = \log_r q$ factor in the dimension m_2 .

Definition of G. Write

$$\mathbf{G} = \left[\mathbf{G}^{(1)} \mid \dots \mid \mathbf{G}^{(m_1)} \mid \mathbf{0} \right] \in \mathbb{Z}^{m_1 \times m_2}$$

as a block matrix consisting of m_1 blocks $\mathbf{G}^{(i)}$ having ℓ columns each, and a final zero block consisting of the remaining $m_2 - m_1 \cdot \ell$ columns (if any). As per our usual notation, $\mathbf{g}_j^{(i)}$ and \mathbf{h}'_j denote the j th columns of $\mathbf{G}^{(i)}$ and \mathbf{H}' , respectively. For each $i \in [m_1]$, $\mathbf{G}^{(i)}$ is defined as follows: let $\mathbf{g}_\ell^{(i)} = \mathbf{h}'_i$, and for each $j = \ell - 1, \dots, 1$, let

$$\mathbf{g}_j^{(i)} = \lfloor \mathbf{g}_{j+1}^{(i)} / r \rfloor = \lfloor \mathbf{h}'_i / r^{\ell-j} \rfloor,$$

where the division and floor operations are coordinate-wise.

Note that because all the entries of \mathbf{h}'_i are less than $q \leq r^\ell$, all the entries of $\mathbf{g}_1^{(i)}$ are in the range $[0, r - 1]$.

Definition of P. For each $j \in [m_1]$, let $\mathbf{p}_j = \mathbf{e}_{j\ell} \in \mathbb{Z}^{m_2}$, the $(j\ell)$ th standard basis vector. Observe that the i th column of \mathbf{P} simply selects the rightmost column of $\mathbf{G}^{(i)}$, yielding $\mathbf{G}\mathbf{P} = \mathbf{H}'$, as desired. Clearly, $\|\mathbf{p}_j\|^2 = 1$ for all $j \in [m_1]$.

Definition of U. Define the unimodular upper-triangular matrix $\mathbf{T}_\ell \in \mathbb{Z}^{\ell \times \ell}$ to have diagonal entries equal to 1 (i.e., $t_{i,i} = 1$ for every $i \in [\ell]$), upper diagonal entries equal to $-r$ (i.e., $t_{i,i+1} = -r$ for every $i \in [\ell - 1]$), and zero entries elsewhere. Define $\mathbf{U} \in \mathbb{Z}^{m_2 \times m_2}$ to be the block-diagonal matrix

$$\mathbf{U} = \text{diag}(\mathbf{T}_\ell, \dots, \mathbf{T}_\ell, \mathbf{I})$$

consisting of m_1 blocks \mathbf{T}_ℓ , followed by the square identity matrix of dimension $m_2 - m_1 \cdot \ell$.

Note that \mathbf{U} is unimodular and that $\|\mathbf{u}_j\|^2 \leq r^2 + 1$ for all j . Also observe that

$$\mathbf{G}\mathbf{U} = \left[\mathbf{G}^{(1)} \cdot \mathbf{T}_\ell \mid \dots \mid \mathbf{G}^{(m_1)} \cdot \mathbf{T}_\ell \mid \mathbf{0} \right].$$

We claim that all the entries of each block $\mathbf{F}^{(i)} = \mathbf{G}^{(i)} \cdot \mathbf{T}_\ell$ are integers in the range $[0, r - 1]$, and thus $\|\mathbf{f}_j^{(i)}\|^2 \leq m_1 \cdot (r - 1)^2$. First observe that the claim is true for $\mathbf{f}_1^{(i)} = \mathbf{g}_1^{(i)}$, as explained above. Moreover, for each $j \in [\ell - 1]$ we have

$$\mathbf{f}_{j+1}^{(i)} = \mathbf{g}_{j+1}^{(i)} - r \cdot \mathbf{g}_j^{(i)} = \mathbf{g}_{j+1}^{(i)} - r \cdot \lfloor \mathbf{g}_{j+1}^{(i)} / r \rfloor,$$

which establishes the claim.

Definition of R. Each entry in the top $d = (1 + \delta)n \lg q$ rows of \mathbf{R} is an independent $\{0, \pm 1\}$ -valued random variable that is 0 with probability $\frac{1}{2}$, 1 with probability $\frac{1}{4}$, and -1 with probability $\frac{1}{4}$. The remaining entries are all 0.

Observe that $\|\mathbf{r}_j\|^2 \leq d$ for all j . Also, by Lemma 2.1 and the discussion following it (with $G = \mathbb{Z}_q^n$), we have that $\mathbf{A} = [\mathbf{A}_1 \mid -\mathbf{A}_1(\mathbf{G} + \mathbf{R})]$ is $(m_2 \cdot q^{-\delta n/2})$ -uniform over $\mathbb{Z}_q^{n \times m}$, as claimed. (Note that it is also suitable to use uniform and independent 0-1 random variables in the top d rows of \mathbf{R} .)

3.2.2 Quality of \mathbf{S}

We now analyze the length of the basis matrix \mathbf{S} . By the triangle inequality and Pythagorean theorem,

$$\|\mathbf{S}\|^2 \leq \max \{ (\|\mathbf{GU}\| + \|\mathbf{RU}\|)^2 + \|\mathbf{U}\|^2, \|\mathbf{RP} - \mathbf{I}\|^2 + \|\mathbf{P}\|^2 \}.$$

We have

$$\|\mathbf{P}\|^2 = 1 \quad \text{and} \quad \|\mathbf{RP} - \mathbf{I}\|^2 \leq 4d < 4r^2 m_1,$$

because each entry of $\mathbf{RP} - \mathbf{I}$ has magnitude at most 2. Therefore, $\|\mathbf{RP} - \mathbf{I}\|^2 + \|\mathbf{P}\|^2 < 4r^2(m_1 + 1)$.

Next, we have

$$\|\mathbf{GU}\|^2 \leq m_1(r-1)^2 \quad \text{and} \quad \|\mathbf{RU}\|^2 \leq d(r+1)^2 \leq m_1(r+1)^2,$$

because every entry in the top d rows of \mathbf{RU} has magnitude at most $r+1$ (and the other entries are zero). Thus $(\|\mathbf{GU}\| + \|\mathbf{RU}\|)^2 \leq 4r^2 m_1$, and because $\|\mathbf{U}\|^2 \leq r^2 + 1 < 4r^2$, the claim follows.

3.3 Second Construction

Theorem 3.2 is an immediate corollary of the following lemma.

Lemma 3.5. *Let $\delta > 0$ be any fixed constant. There is a universal constant $C > 0$ and a probabilistic polynomial-time algorithm that, given uniformly random $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times m_1}$ for any $m_1 \geq d = (1 + \delta)n \lg q$, and any integer $m_2 \geq (4 + 2\delta)n \lg q$ (in unary), outputs matrices $\mathbf{U}, \mathbf{G}, \mathbf{R}, \mathbf{P}$ and $\mathbf{C} = \mathbf{I}$ as required by Step 1 of Algorithm 1 such that:*

- $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2]$ is $(m_2 \cdot q^{-\delta n/2})$ -uniform, where \mathbf{A}_2 is as in Step 2 of Algorithm 1.
- $\|\mathbf{S}\| \leq Cn \lg q$ with probability $1 - 2^{-\Omega(n)}$ over the choice of \mathbf{R} , where \mathbf{S} is as in Step 3 of Algorithm 1.
- $\|\tilde{\mathbf{S}}\| \leq 1 + C\sqrt{d} = O(\sqrt{n \log q})$ with probability $1 - 2^{-\Omega(n)}$ over the choice of \mathbf{R} .

We have not attempted to optimize the exact constant C appearing in the above bounds, but it is not exceedingly large (at most 20, certainly). The remainder of this subsection is devoted to proving the lemma.

3.3.1 Construction

Given \mathbf{A}_1 , let $\mathbf{H} \in \mathbb{Z}^{m_1 \times m_1}$ be the Hermite normal form of $\Lambda^\perp(\mathbf{A}_1)$. The basic idea behind the construction is to ensure that the columns of \mathbf{G} include sufficiently many power-of-2 multiples of each standard basis vector $\mathbf{e}_i \in \mathbb{Z}^{m_1}$. This allows us to express each vector in $\mathbf{H}' = \mathbf{H} - \mathbf{I}$ simply as a binary combination of such vectors. (The $-\mathbf{I}$ term is included to make every entry in the i th row of \mathbf{H}' strictly smaller than $h_{i,i}$, which yields a tighter bound on m_2 .) To obtain a good bound on the length of the Gram-Schmidt orthogonalization $\tilde{\mathbf{S}}$, we additionally ensure that certain rows of \mathbf{G} are mutually orthogonal and sufficiently long. This ensures that adding the random matrix \mathbf{R} to \mathbf{G} does not ‘distort the shape’ of \mathbf{G} by much, which is important in the analysis of the orthogonalization.

Recall that every diagonal entry $h_{i,i}$ of the Hermite normal form $\mathbf{H} \in \mathbb{Z}^{m_1 \times m_1}$ is at least 1, that

$$\prod_{i \in [m_1]} h_{i,i} = \det(\mathbf{H}) = \det(\Lambda^\perp(\mathbf{A}_1)) \leq q^n,$$

and that $0 \leq h_{i,j} < h_{i,i}$ for every $j \neq i$. Therefore, every column \mathbf{h}'_j of $\mathbf{H}' = \mathbf{H} - \mathbf{I}$ belongs to the Cartesian product

$$\prod_{i \in [m_1]} [0, \dots, h_{i,i} - 1] \subset \mathbb{Z}^{m_1},$$

which has size $\prod_{i \in [m_1]} h_{i,i} \leq q^n$.

Definition of \mathbf{G} . Write

$$\mathbf{G} = \left[\mathbf{G}^{(1)} \mid \dots \mid \mathbf{G}^{(m_1)} \mid \mathbf{M} \mid \mathbf{0} \right] \in \mathbb{Z}^{m_1 \times m_2}$$

as a block matrix of m_1 blocks $\mathbf{G}^{(i)}$ having various widths, followed by a special block \mathbf{M} , followed by a zero block of any remaining columns. For each $i \in [m_1]$, block $\mathbf{G}^{(i)}$ has width $w_i = \lceil \lg h_{i,i} \rceil < 1 + \lg h_{i,i}$, and its j th column is $\mathbf{g}_j^{(i)} = 2^{j-1} \cdot \mathbf{e}_i \in \mathbb{Z}^{m_1}$. Note that if $h_{i,i} = 1$, block $\mathbf{G}^{(i)}$ actually has width 0, and that there are at most $n \lg q$ values of i for which $h_{i,i} > 1$. Taking all blocks $\mathbf{G}^{(i)}$ together, the total number of columns is therefore

$$\sum_{i \in [m_1]} w_i \leq n \lg q + \sum_{i \in [m_1]} \lg h_{i,i} \leq 2n \lg q.$$

(In the special case that q is prime, there are at most n values of $h_{i,i}$ that are greater than 1, which are all q , so the total number of columns in this case is at most $n \lceil \lg q \rceil$.)

The block \mathbf{M} is a special component needed only for the analysis of $\|\tilde{\mathbf{S}}\|$; the bound on the length $\|\mathbf{S}\|$ from Lemma 3.5 holds even if we leave out \mathbf{M} (which allows for a smaller value of m_2). The block \mathbf{M} has width w , where w is the largest power of 2 in the range $[d, m_2 - 2n \lg q]$. Note that $m_2 - 2n \lg q \geq 2d$, so a power of 2 always exists in the given range, and that $w \geq m_2/2 - n \lg q \geq m_2/4$.

Block \mathbf{M} is zero in all but its first d rows, which are distinct rows of a square *Hadamard* matrix of dimension w , times a suitably large constant $C' > 0$. Recall that a Hadamard matrix is a square ± 1 matrix whose rows are mutually orthogonal; a Hadamard matrix in any dimension 2^k may be constructed in time $\text{poly}(2^k)$ using Sylvester's recursive formula $\mathbf{H}_{2^k} = \begin{pmatrix} \mathbf{H}_{2^{k-1}} & \mathbf{H}_{2^{k-1}} \\ \mathbf{H}_{2^{k-1}} & -\mathbf{H}_{2^{k-1}} \end{pmatrix}$, with base case $\mathbf{H}_1 = [1]$.

Definition of \mathbf{P} . Mirroring the structure of \mathbf{G} , we write

$$\mathbf{P} = \left[\mathbf{P}^{(1)}; \dots; \mathbf{P}^{(m_1)}; \mathbf{0}; \mathbf{0} \right] \in \mathbb{Z}^{m_2 \times m_1}$$

as a vertical block matrix where each block $\mathbf{P}^{(i)} \in \mathbb{Z}^{w_i \times m_1}$.

For each $i, j \in [m_1]$, the j th column $\mathbf{p}_j^{(i)}$ of $\mathbf{P}^{(i)}$ contains the binary representation of $h'_{i,j} \in [0, \dots, h_{i,i} - 1]$, which has length at most w_i . Specifically, $\mathbf{P}^{(i)}$ contains entries $p_{k,j}^{(i)} \in \{0, 1\}$ such that

$$h'_{i,j} = \sum_{k \in [w_i]} p_{k,j}^{(i)} \cdot 2^{k-1}.$$

Note that $\|\mathbf{p}_j\|^2 \leq \sum_{i \in [m_1]} w_i \leq 2n \lg q$.

By definition of $\mathbf{G}^{(i)}$, we have $\mathbf{G}^{(i)} \mathbf{p}_j^{(i)} = \mathbf{e}_i \cdot \sum_{k \in [m_1]} p_{k,j}^{(i)} \cdot 2^{k-1} = \mathbf{e}_i \cdot h'_{i,j}$, hence

$$\mathbf{G} \mathbf{P} = \sum_{i \in [m_1]} \mathbf{G}^{(i)} \mathbf{P}^{(i)} = \mathbf{H}',$$

as desired.

Definition of \mathbf{U} . Let $\mathbf{T}_w \in \mathbb{Z}^{w \times w}$ be the upper-triangular unimodular matrix with 1s along the diagonal and -2 s along the upper diagonal, i.e., $t_{i,i} = 1$ for $i \in [w]$ and $t_{i,i+1} = -2$ for $i \in [w-1]$ (all other entries are zero). By definition of $\mathbf{G}^{(i)}$, observe that $\mathbf{F}^{(i)} = \mathbf{G}^{(i)} \cdot \mathbf{T}_{w_i} \in \mathbb{Z}^{m_1 \times w_i}$ is simply \mathbf{e}_i in its first column and zero elsewhere. Then letting \mathbf{U} be the block diagonal matrix

$$\mathbf{U} = \text{diag}(\mathbf{T}_{w_1}, \dots, \mathbf{T}_{w_{m_1}}, \mathbf{I}) \in \mathbb{Z}^{m_2 \times m_2},$$

we see that \mathbf{U} is unimodular and very short, i.e., $\|\mathbf{U}\|^2 \leq 5$, and that

$$\mathbf{G}\mathbf{U} = \left[\mathbf{F}^{(1)} \mid \dots \mid \mathbf{F}^{(m_1)} \mid \mathbf{M} \mid \mathbf{0} \right]$$

is also short, i.e., $\|\mathbf{G}\mathbf{U}\| \leq C'\sqrt{d}$.

Definition of \mathbf{R} . Each entry in the top $d = (1 + \delta)n \lg q$ rows of \mathbf{R} is an independent $\{0, \pm 1\}$ -valued random variable that is 0 with probability $\frac{1}{2}$, 1 with probability $\frac{1}{4}$, and -1 with probability $\frac{1}{4}$. The remaining entries are all 0.

Observe that $\|\mathbf{r}_j\|^2 \leq d$ for all j . Also, by Lemma 2.1 and the discussion following it (with $G = \mathbb{Z}_q^n$), we have that $\mathbf{A} = [\mathbf{A}_1 \mid -\mathbf{A}_1(\mathbf{G} + \mathbf{R})]$ is $(m_2 \cdot q^{-\delta n/2})$ -uniform over $\mathbb{Z}_q^{n \times m}$, as claimed.

3.3.2 Quality of \mathbf{S}

We now analyze $\|\mathbf{S}\|$ and $\|\tilde{\mathbf{S}}\|$. For both analyses, we partition \mathbf{S} into two sets of vectors,

$$\mathbf{S}_1 = \{\mathbf{s}_j\}_{j \in [m_2]} = [(\mathbf{G} + \mathbf{R})\mathbf{U}; \mathbf{U}] \quad \text{and} \quad \mathbf{S}_2 = \{\mathbf{s}_j\}_{j > m_2} = [\mathbf{R}\mathbf{P} - \mathbf{I}; \mathbf{P}].$$

Length of basis vectors. We have

$$\|\mathbf{S}\| = \max\{\|\mathbf{S}_1\|, \|\mathbf{S}_2\|\}.$$

By the Pythagorean theorem and the triangle inequality,

$$\|\mathbf{S}_1\|^2 \leq \|\mathbf{G}\mathbf{U} + \mathbf{R}\mathbf{U}\|^2 + \|\mathbf{U}\|^2 \leq (C'\sqrt{d} + 3\sqrt{d})^2 + 5 \leq (C\sqrt{d} + 1)^2, \quad (3.5)$$

for some large enough constant $C > 0$.

For $\|\mathbf{S}_2\|$, observe that \mathbf{R} is zero on all but a $d \times m_2$ submatrix whose entries are independent subgaussian random variables with some constant parameter $C'' > 0$. Therefore by Fact 2.2, for every fixed \mathbf{p}_j , the first d entries of $\mathbf{R}\mathbf{p}_j \in \mathbb{R}^{m_1}$ are independent subgaussian variables with parameter $C'' \cdot \|\mathbf{p}_j\| = O(\sqrt{n \log q})$. By Lemma 2.3, the largest singular value of $\mathbf{R}\mathbf{p}_j$, and hence the length $\|\mathbf{R}\mathbf{p}_j\|$, is at most $O(\sqrt{dn \log q}) = O(n \log q)$ except with probability $2^{-\Omega(n)}$. By the union bound and triangle inequality, we conclude that $\|\mathbf{S}_2\| = O(n \log q)$ except with probability $2^{-\Omega(n)}$, as desired.

Length of Gram-Schmidt vectors. First we review some preliminary facts that are needed in the analysis. Let $\mathbf{X} \in \mathbb{R}^{m \times \ell}$ be any set of $\ell \leq m$ linearly independent vectors. Then $\pi_{\mathbf{X}} := \mathbf{X} \cdot (\mathbf{X}^t \mathbf{X})^{-1} \cdot \mathbf{X}^t \in \mathbb{R}^{m \times m}$ is the projection matrix of the orthogonal linear projection from \mathbb{R}^m to $\text{span}(\mathbf{X}) \subseteq \mathbb{R}^m$. (Note that the Gram matrix $\mathbf{X}^t \mathbf{X}$ is invertible because the vectors in \mathbf{X} are linearly independent.) This fact may be verified by observing that any $\mathbf{v} \in \text{span}(\mathbf{X})$ may be written as $\mathbf{v} = \mathbf{X}\mathbf{c}$ for some $\mathbf{c} \in \mathbb{R}^\ell$, hence

$$\pi_{\mathbf{X}} \cdot \mathbf{v} = \mathbf{X} \cdot (\mathbf{X}^t \mathbf{X})^{-1} \cdot \mathbf{X}^t \mathbf{X} \cdot \mathbf{c} = \mathbf{X}\mathbf{c} = \mathbf{v};$$

moreover, for any $\mathbf{v} \in \text{span}^\perp(\mathbf{X})$ we have $\mathbf{X}^t \mathbf{v} = \mathbf{0}$ and hence $\pi_{\mathbf{X}} \cdot \mathbf{v} = \mathbf{0}$. Also note that for any $\mathbf{v} \in \mathbb{R}^m$,

$$\|\pi_{\mathbf{X}} \cdot \mathbf{v}\|^2 = \langle \pi_{\mathbf{X}} \cdot \mathbf{v}, \pi_{\mathbf{X}} \cdot \mathbf{v} \rangle = \langle \mathbf{v}, \pi_{\mathbf{X}} \cdot \mathbf{v} \rangle = \mathbf{v}^t \cdot \pi_{\mathbf{X}} \cdot \mathbf{v} = (\mathbf{X}^t \mathbf{v})^t \cdot (\mathbf{X}^t \mathbf{X})^{-1} \cdot (\mathbf{X}^t \mathbf{v}), \quad (3.6)$$

because $\mathbf{v} - \pi_{\mathbf{X}} \cdot \mathbf{v}$ is orthogonal to $\pi_{\mathbf{X}} \cdot \mathbf{v}$.

In particular, we define $\mathbf{X} \in \mathbb{R}^{m \times m_1}$ as

$$\mathbf{X} = [-\mathbf{I} | \mathbf{G} + \mathbf{R}]^t,$$

and observe that the columns of \mathbf{X} are linearly independent and form a basis of $\text{span}^\perp(\mathbf{S}_1)$, because

$$\dim \text{span}(\mathbf{X}) = m_1 = m - \dim \text{span}(\mathbf{S}_1) \quad \text{and} \quad \mathbf{X}^t \cdot \mathbf{S}_1 = -(\mathbf{G} + \mathbf{R})\mathbf{U} + (\mathbf{G} + \mathbf{R})\mathbf{U} = \mathbf{0}.$$

We now analyze $\|\tilde{\mathbf{S}}\|$. Observe that

$$\|\tilde{\mathbf{S}}\| = \max_{j \in [m]} \|\tilde{\mathbf{s}}_j\| \leq \max \{\|\mathbf{S}_1\|, \|\pi_{\mathbf{X}} \cdot \mathbf{S}_2\|\}, \quad (3.7)$$

because $\|\tilde{\mathbf{s}}_j\| \leq \|\mathbf{s}_j\|$ for all $j \in [m_2]$, and $\tilde{\mathbf{s}}_j$ is the orthogonal projection of \mathbf{s}_j onto a linear subspace of $\text{span}(\mathbf{X})$ for all $j > m_2$. Equation (3.5) has already established that $\|\mathbf{S}_1\| \leq C\sqrt{d} + 1$.

Bounding $\|\pi_{\mathbf{X}} \cdot \mathbf{S}_2\|$ is more involved. We start by setting up some additional notation that will make the analysis more convenient. Define

$$\hat{\mathbf{G}} = [-\mathbf{I} | \mathbf{G}], \quad \hat{\mathbf{R}} = [\mathbf{0} | \mathbf{R}] \in \mathbb{Z}^{m_1 \times m}, \quad \hat{\mathbf{P}} = [\mathbf{0}; \mathbf{P}] \in \mathbb{Z}^{m \times m_1}, \quad \hat{\mathbf{S}}_2 = \mathbf{S}_2 + [\mathbf{I}; \mathbf{0}] = [\mathbf{R}; \mathbf{I}] \cdot \mathbf{P}.$$

We have

$$\|\pi_{\mathbf{X}} \cdot \mathbf{S}_2\| \leq \|\pi_{\mathbf{X}} \cdot \hat{\mathbf{S}}_2\| + \|\pi_{\mathbf{X}} \cdot [\mathbf{I}; \mathbf{0}]\| \leq \|\pi_{\mathbf{X}} \cdot \hat{\mathbf{S}}_2\| + 1,$$

by the triangle inequality and the fact that $\pi_{\mathbf{X}}$ represents an orthogonal projection onto a subspace of \mathbb{R}^m . Therefore, it is enough to bound $\|\pi_{\mathbf{X}} \cdot \hat{\mathbf{S}}_2\|$. To do so, we analyze the two main components of the right-hand side of Equation (3.6). We have

$$\begin{aligned} \mathbf{X}^t \cdot \hat{\mathbf{S}}_2 &= [-\mathbf{I} | \mathbf{G} + \mathbf{R}] \cdot [\mathbf{R}; \mathbf{I}] \cdot \mathbf{P} = \mathbf{G} \cdot \mathbf{P} = \hat{\mathbf{G}} \cdot \hat{\mathbf{P}}, \\ \mathbf{X}^t \mathbf{X} &= (\hat{\mathbf{G}} + \hat{\mathbf{R}})(\hat{\mathbf{G}} + \hat{\mathbf{R}})^t. \end{aligned}$$

We therefore want to analyze the properties of the positive semidefinite matrix

$$\mathbf{Z} = \hat{\mathbf{G}}^t \cdot \left((\hat{\mathbf{G}} + \hat{\mathbf{R}})(\hat{\mathbf{G}} + \hat{\mathbf{R}})^t \right)^{-1} \cdot \hat{\mathbf{G}}. \quad (3.8)$$

Note that the rows of $\hat{\mathbf{G}}$ are orthogonal by construction (because the rows of \mathbf{G} are), that all its rows have length at least 1, and that its first d rows have length at least $C'\sqrt{w} \geq C'\sqrt{m_2}/2$ by the properties of the block \mathbf{M} . Therefore, we may factor $\hat{\mathbf{G}}$ as

$$\hat{\mathbf{G}} = \mathbf{D} \cdot \mathbf{V}$$

where the rows of $\mathbf{V} \in \mathbb{R}^{m_1 \times m}$ are orthonormal (i.e., $\mathbf{V}\mathbf{V}^t = \mathbf{I}$), and $\mathbf{D} \in \mathbb{R}^{m_1 \times m_1}$ is a nonsingular square diagonal matrix whose first d diagonal entries are all at least $C'\sqrt{m_2}/2$. Bringing \mathbf{D} into the inverted central term of Equation (3.8) from both sides, we therefore have

$$\mathbf{Z} = \mathbf{V}^t \cdot \left((\mathbf{V} + \mathbf{D}^{-1}\hat{\mathbf{R}})(\mathbf{V} + \mathbf{D}^{-1}\hat{\mathbf{R}})^t \right)^{-1} \cdot \mathbf{V}.$$

Below, we show that the singular values of $\mathbf{Y} = \mathbf{V} + \mathbf{D}^{-1}\hat{\mathbf{R}}$ are all at least $\frac{1}{2}$, with very high probability. Given this, it follows that the singular values of \mathbf{Z} , which are also its eigenvalues because \mathbf{Z} is positive semidefinite, are all at most 4. Now \mathbf{Z} may be factored as $\mathbf{Z} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}$ for some orthogonal matrix \mathbf{Q} and diagonal matrix $\mathbf{\Lambda}$ whose diagonal entries are the eigenvalues of \mathbf{Z} . From this we have

$$\|\pi_{\mathbf{x}} \cdot \hat{\mathbf{S}}_2\|^2 = \max_{j \in [m_1]} \|\hat{\mathbf{p}}_j^t \cdot \mathbf{Z} \cdot \hat{\mathbf{p}}_j\|^2 \leq \max_{j \in [m_1]} (4 \cdot \|\hat{\mathbf{p}}_j\|^2) \leq 8n \lg q < (3\sqrt{d})^2.$$

It remains to bound the singular values of $\mathbf{Y} = \mathbf{V} + \mathbf{D}^{-1}\hat{\mathbf{R}}$ from below by $\frac{1}{2}$. To do so, it suffices to bound the singular values of $\mathbf{D}^{-1}\hat{\mathbf{R}}$ from above by $\frac{1}{2}$, because by the triangle inequality and the fact that the rows of \mathbf{V} are orthonormal, the smallest singular value of \mathbf{Y} is

$$\min_{\mathbf{x} \in S^{m-1}} \|\mathbf{V}^t \mathbf{x} + (\mathbf{D}^{-1}\hat{\mathbf{R}})^t \mathbf{x}\| \geq 1 - \max_{\mathbf{x} \in S^{m-1}} \|(\mathbf{D}^{-1}\hat{\mathbf{R}})^t \mathbf{x}\| \geq \frac{1}{2}.$$

By definition of $\hat{\mathbf{R}}$ and the properties of \mathbf{D} , the matrix $\mathbf{D}^{-1}\hat{\mathbf{R}}$ is zero on all but a $d \times m_2$ submatrix whose entries are independent subgaussian random variables of parameter $1/(C''\sqrt{m_2})$, where $C'' > 0$ is some constant multiple of C' . Lemma 2.3 implies that with probability $1 - 2^{-\Omega(d)}$, the singular values of $\mathbf{D}^{-1}\hat{\mathbf{R}}$ are all at most

$$\frac{C(\sqrt{d} + \sqrt{m_2})}{C''\sqrt{m_2}} \leq \frac{1}{2}$$

(for sufficiently large constant C''), and the proof is complete.

Acknowledgments

We thank Daniele Micciancio and the anonymous referees for helpful comments on the presentation.

References

- [Ajt96] M. Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.
- [Ajt99] M. Ajtai. Generating hard instances of the short basis problem. In *ICALP*, pages 1–9. 1999.
- [CHKP10] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552. 2010.
- [GGH96] O. Goldreich, S. Goldwasser, and S. Halevi. Collision-free hashing from lattice problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(42), 1996.
- [GGH97] O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *CRYPTO*, pages 112–131. 1997.
- [GHV10] C. Gentry, S. Halevi, and V. Vaikuntanathan. A simple BGN-type cryptosystem from LWE. In *EUROCRYPT*, pages 506–522. 2010.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. 2008.

- [HILL99] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [Mic01] D. Micciancio. Improving lattice based cryptosystems using the Hermite normal form. In *CaLC*, pages 126–145. 2001.
- [MO90] J. E. Mazo and A. M. Odlyzko. Lattice points in high-dimensional spheres. *Monatshefte für Mathematik*, 110(1):47–61, March 1990.
- [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.
- [MR09] D. Micciancio and O. Regev. Lattice-based cryptography. In *Post Quantum Cryptography*, pages 147–191. Springer, February 2009.
- [MV03] D. Micciancio and S. P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *CRYPTO*, pages 282–298. 2003.
- [MW01] D. Micciancio and B. Warinschi. A linear space algorithm for computing the Hermite normal form. In *ISSAC*, pages 231–236. 2001.
- [Ngu99] P. Q. Nguyen. Cryptanalysis of the Goldreich-Goldwasser-Halevi cryptosystem from Crypto ’97. In *CRYPTO*, pages 288–304. 1999.
- [NR06] P. Q. Nguyen and O. Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *J. Cryptology*, 22(2):139–160, 2009. Preliminary version in Eurocrypt 2006.
- [Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. 2009.
- [PV08] C. Peikert and V. Vaikuntanathan. Noninteractive statistical zero-knowledge proofs for lattice problems. In *CRYPTO*, pages 536–553. 2008.
- [PVW08] C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571. 2008.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009. Preliminary version in STOC 2005.
- [Sho97] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.
- [Ver07] R. Vershynin. Lecture notes on non-asymptotic theory of random matrices, 2007. Available at <http://www-personal.umich.edu/~romanv/teaching/2006-07/280/>, last accessed 17 Feb 2010.