

Generating Shorter Bases for Hard Random Lattices

Joël Alwen*
New York University

Chris Peikert †
SRI International

March 30, 2009

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 for generating public/secret key pairs. In these applications, a shorter basis directly corresponds to milder underlying complexity assumptions and smaller key sizes.

The contributions of this work are twofold. First, using the *Hermite normal form* as an organizing principle, we simplify and generalize an approach due to Ajtai (ICALP 1999). Second, we improve the construction and its analysis in several ways, most notably by tightening the length of the output basis essentially to the optimum value.

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

*Work 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, lattice-based *cryptographic* schemes have emerged as an intriguing alternative to more traditional ones based on, e.g., the factoring and discrete logarithm problems. Among other reasons, this is because such schemes have yet to be broken by quantum algorithms, and their security (on the average, for almost all choices of random keys) can be based solely on *worst-case* computational assumptions.

In 1996, Ajtai’s seminal work [Ajt04] in this area demonstrated a family of random lattices for which finding relatively short nonzero lattice vectors is at least as hard as approximating the well-known Shortest Vector Problem (among others) in the *worst case*. 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, MR07, Reg05, GPV08]).

Ajtai’s initial paper also showed that a hard random lattice can be generated along with *one relatively short* lattice vector, which can be useful as a secret key in cryptographic settings (though such applications seem spare; see [MV03] for the one example of which we are aware). Shortly thereafter, Goldreich, Goldwasser and Halevi [GGH97] proposed public-key cryptographic schemes (though without security proofs) in which the secret key is a *short basis* (i.e., a basis in which all of the vectors are relatively short) of some public lattice. One method proposed in [GGH97] for generating a lattice along with a short basis is first to choose the short basis vectors, and then to transform it into a “random” public basis by a sequence of lattice-preserving transformations. Unfortunately, this method does not produce lattices from the provably hard family defined in [Ajt04]. Although improvements to the GGH lattice generator and public-key cryptosystem were later proposed by Micciancio [Mic01] (following a cryptanalysis of the original scheme by Nguyen [Ngu99]), there is still no known proof that the resulting random lattices are actually hard on the average. (We should 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 [GGH97], Ajtai demonstrated an entirely different method of generating a lattice together with a short basis [Ajt99]. This generator has the important advantage that the resulting lattice is drawn, under the appropriate distribution, from the hard family defined in [Ajt04]. Interestingly, the algorithm apparently went without application until very recently, when Gentry, Peikert and Vaikuntanathan [GPV08] constructed provably secure (under worst-case assumptions) cryptographic schemes that crucially use short bases as their secret keys; see also the subsequent works [PVW08, PV08, Pei08] for other applications. At this point we should mention that technically, the main 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.

The maximal length of the generated basis vectors directly affects the security and efficiency of the application in which it is used, both in theory and in practice. More specifically, it 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 2.1 for details). Therefore, it is very desirable to generate a set that is as short as possible. Unfortunately, the result from [Ajt99] is far from optimal — the length is bounded only by $O(m^{5/2})$, versus the optimal bound of about \sqrt{m} (for commonly used parameters) — and the method appears not to have attracted much attention or improvement since its publication almost a decade ago (probably due to the lack of applications until recently).

1.1 Our Contributions

Our first contribution is to elucidate and generalize Ajtai’s algorithm [Ajt99] for generating a hard random lattice along with a relatively short full-rank set of lattice vectors. We endeavor to give a high-level, modular exposition of the method and the main concerns that motivate its structure (in the process, we also correct some minor errors in the original paper). One novelty in our approach is to design and analyze the algorithm around the concept of the *Hermite normal form* (HNF), which is a unique canonical representation for (integer) lattices. Micciancio [Mic01] has proposed using the HNF in cryptographic applications to specify a lattice in its “least revealing” representation; here we use the HNF as the central organizing tool for ensuring that the short basis corresponds to a (uniformly random) lattice from the hard family of [Ajt04].

Our second contribution is to refine the algorithm and its analysis, improving it in several ways. First and most importantly, we improve the length of its output set from $O(m^{5/2})$ to as low as $O(\sqrt{m})$, 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. Our second main improvement is to make the generator work for an *arbitrary* integer modulus q and to output a *basis* of the resulting lattice, whereas the original algorithm of [Ajt99] works only for *odd* q and produces just a full-rank set. Using an *even* modulus q happens to be important in recent cryptosystems of Peikert [Pei08] that are based on the standard worst-case shortest vector problem. Generating a basis (versus a full-rank set) seems to be less of an advantage, but it may have unanticipated uses elsewhere.

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

1.2 Relation to Ajtai’s Construction

Our construction is inspired by Ajtai’s, but differs from it in most of the details. The greatest similarity is in our use of a specially crafted unimodular matrix (called \mathbf{B} in this work) that has small entries, but whose inverse matrix \mathbf{B}^{-1} contains geometrically increasing sequences of integers. As in [Ajt99], a crucial step in our construction involves assembling other matrices with large entries via products of short vectors and \mathbf{B}^{-1} .

In terms of its main differences from [Ajt99], our construction is guided from the “top down” by the abstract block structure of the short basis, the desired distribution of its Hermite normal form, and the unimodular transformation relating the two. This approach also yields various technical simplifications and corrections. In particular, it lets us completely separate the *structural constraints* on the basis from the *randomization* of the output lattice, and it facilitates a generalization to arbitrary moduli q . (In Ajtai’s construction, the structure and randomization are tightly coupled, and q is assumed to be *odd* when arguing that the output set is full-rank.)

2 Preliminaries

For a positive integer k , let $[k]$ denote the set $\{1, \dots, k\}$. We denote the set of integers modulo q by \mathbb{Z}_q , and identify it with the set $\{0, \dots, q - 1\}$ in the natural way. Row 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

i th row of a matrix \mathbf{X} is denoted \mathbf{x}_i . We identify a matrix \mathbf{X} with the (ordered) set $\{\mathbf{x}_i\}$ of its row vectors, and define $\|\mathbf{X}\| = \max_i \|\mathbf{x}_i\|$. We let \mathbf{e}_i denote the i th standard basis vector, where its dimension will be clear from context. The symbol \mathbf{I}_d denotes the $d \times d$ identity matrix.

2.1 Lattices

Generally defined, a *lattice* Λ is a discrete additive subgroup of \mathbb{R}^m for some nonnegative integer m . In this work, every lattice will be a *full-rank integer* lattice, which is a discrete additive subgroup 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 $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_m\} \subset \mathbb{Z}^m$:

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

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 $\mathbf{B} = \mathbf{U} \cdot \mathbf{B}'$ for some unimodular matrix $\mathbf{U} \in \mathbb{Z}^{m \times m}$. The determinant of a basis matrix \mathbf{B} is exactly 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 diagonal entries (i.e., $h_{i,i} \geq 1$), and every entry above the diagonal is strictly smaller than the diagonal entry in its column (i.e., $h_{i,j} < h_{j,j}$ 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 its diagonal entries. For a lattice basis \mathbf{B} , we write $\text{HNF}(\mathbf{B})$ to denote $\text{HNF}(\mathcal{L}(\mathbf{B}))$. It follows that for $\mathbf{H} = \text{HNF}(\mathbf{B})$, there exists a (unique) unimodular matrix \mathbf{U} such that $\mathbf{B} = \mathbf{U} \cdot \mathbf{H}$. In addition, the matrices \mathbf{U} and \mathbf{H} can be computed in polynomial time given \mathbf{B} (see [MW01] and references therein).

Hard random lattices. We will be especially concerned with a certain family of lattices in \mathbb{Z}^m as first defined by Ajtai [Ajt04]. A lattice from this family is most naturally specified not by a basis, but instead by a *parity check* matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ for some positive integer n and positive integer modulus q . (We discuss the parameters m , n , and q in detail below). The associated lattice is defined as

$$\mathcal{L}^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}\mathbf{A} = \mathbf{0} \pmod{q}\} \subseteq \mathbb{Z}^m.$$

It is routine to check that $\mathcal{L}^\perp(\mathbf{A})$ contains the identity $\mathbf{0} \in \mathbb{Z}^m$ and is closed under addition, hence it is a subgroup of (and lattice in) \mathbb{Z}^m . Also observe that $q \cdot \mathbf{e}_i \in \mathcal{L}^\perp(\mathbf{A})$ for every \mathbf{A} and every $i \in [m]$, so membership in $\mathcal{L}^\perp(\mathbf{A})$ is determined solely by a vector’s entries modulo q .

We review some basic facts about this family of lattices. Let $\Lambda = \mathcal{L}^\perp(\mathbf{A})$ for some arbitrary $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$. 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 $\mathbf{x}\mathbf{A} \in \mathbb{Z}_q^n$. 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 $(\mathbf{x} - \mathbf{x}')\mathbf{A} = \mathbf{0}$ which implies $\mathbf{x} - \mathbf{x}' \in \Lambda$, i.e., $(\mathbf{x} + \Lambda) = (\mathbf{x}' + \Lambda) \in (\mathbb{Z}^m/\Lambda)$. 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 that will become clear from the statement of Proposition 2.1 below, the hardness of these lattices is most naturally parameterized by n (not m , even though m is the dimension of the lattices). Therefore, it is standard to consider the parameters $m = m(n)$ and $q = q(n)$ as functions of n . Given n and q , one of the most interesting parameter choices (which essentially minimizes the bound in (2.1)) is to let $m = c \cdot n \lg q$ for some constant $c \geq 1$. Then by (2.1), the minimum distance of $\mathcal{L}^\perp(\mathbf{A})$ for any $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is at most

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

For a *random* \mathbf{A} , a volume argument reveals that with high probability, this bound is essentially tight (up to a small constant factor). Note that for larger choices of m , the minimum distance does not increase because we can just ignore the extra rows of \mathbf{A} . As long as m does not grow extremely large, $\sqrt{n \lg q}$ remains a good estimate for the minimum distance of $\mathcal{L}^\perp(\mathbf{A})$ for random \mathbf{A} .

The following proposition, proved first by Ajtai [Ajt04] (in a quantitatively weaker form) and in its current form in [MR07, GPV08], relates the average-case and worst-case complexity of certain lattice problems.

Proposition 2.1. *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 $\mathbf{x} \in \mathcal{L}^\perp(\mathbf{A})$ having length at most β for uniformly random $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ (with nonnegligible probability over the choice of \mathbf{A} and the randomness of the algorithm) is at least as hard as solving (with overwhelming probability) the approximate shortest vector problem GapSVP (and others) on n -dimensional lattices to within a $\gamma(n) = \beta \cdot \tilde{O}(\sqrt{n})$ factor in the worst case.*

Note that Proposition 2.1 is meaningful only when β is at least the minimum distance of a random $\mathcal{L}^\perp(\mathbf{A})$, otherwise no nonzero vector $\mathbf{x} \in \mathcal{L}^\perp(\mathbf{A})$ of length at most β is likely to exist. For $q = \text{poly}(n)$ and m as described above, we can therefore take β to be as small as $O(\sqrt{n \lg n})$, which yields a problem that is hard on the average assuming the worst-case hardness of approximating GapSVP (and other problems) to within an $\tilde{O}(n)$ factor.

In certain cryptographic applications, however, an adversary that breaks the scheme is guaranteed only to produce lattice vectors that are much longer than the shortest vector in the lattice, so one needs to assume average-case hardness for larger values of β . For example, the secret key in the digital signature schemes of [GPV08] is a basis of $\mathcal{L}^\perp(\mathbf{A})$ having some length L , and its signatures are vectors of length $\approx L\sqrt{m}$. It is shown in [GPV08] that an adversary that is capable of forging a signature is also capable of finding a nonzero lattice vector of length $\beta \approx L\sqrt{m}$ in $\mathcal{L}^\perp(\mathbf{A})$, which by Proposition 2.1 (for our choice of m) is as hard as approximating GapSVP in the worst case to within $L \cdot \tilde{O}(n)$ factors. Therefore, a shorter secret basis immediately induces a weaker underlying hardness assumption.

Note also that Proposition 2.1 requires the modulus q to exceed β by a significant amount (otherwise the trivial vector $q \cdot \mathbf{e}_1$ would be a valid solution), and that m grows with $\lg q$. Therefore, a polynomial factor improvement in the length L of the basis also yields a constant factor improvement in the dimension m and magnitude q of entries in the parity check matrix \mathbf{A} (i.e., the public key).

2.2 Probability

We denote the uniform probability distribution over a finite set G by $U(G)$. 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 $\frac{1}{2} \sum_{g \in G} |D_1(g) - D_2(g)|$.

Lemma 2.2 (Leftover Hash Lemma (Simplified) [HILL99]). *Let \mathcal{H} be a family of 2-universal hash functions from a domain \mathcal{X} to range \mathcal{Y} . and let X be a random variable over \mathcal{X} . Then for $h \leftarrow \mathcal{H}$ and $X \leftarrow \mathcal{X}$ chosen independently and uniformly, $(h, h(X))$ is $\frac{1}{2} \sqrt{|\mathcal{Y}| / |\mathcal{X}|}$ -uniform over $\mathcal{H} \times \mathcal{Y}$.*

3 Construction

Our goal is to generate a (nearly) uniform parity check matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, together with a basis $\mathbf{S} \in \mathbb{Z}^{m \times m}$ of $\mathcal{L}^\perp(\mathbf{A})$ whose vectors are relatively short. Our approach consists of two steps. First we investigate the structure of the Hermite normal form of $\mathcal{L}^\perp(\mathbf{A})$, for a given (random) \mathbf{A} . Then we describe how to generate \mathbf{S} so that its HNF has the appropriate structure and distribution, so as to induce a (nearly) uniform parity check matrix \mathbf{A} .

We give two constructions that are, in general, incomparable. The first construction, described in Theorem 3.1 below, works for a small dimension $m = O(n \log q)$, but the resulting basis length is $\tilde{O}(m)$, which is not optimal. The second construction, described in Theorem 3.2, provides a basis of essentially *optimal* length $O(\sqrt{n \log q})$, but at the cost of a somewhat larger dimension $m = O(n \log^2 q)$. More generally, Theorem 3.2 can actually be parameterized by a base r to yield various trade-offs between the basis length and dimension m ; in general, we can obtain a basis of length $\Theta(r \cdot \sqrt{n \log q})$ with a dimension $m = \Theta(n \log q \log_r q)$.

Most applications use a polynomial modulus $q = \text{poly}(n)$, so the extra $\log q = O(\log n)$ factor (or $\log_r q = O(1/\delta)$ factor, when $r = n^\delta$) in the dimension m in Theorem 3.2 is of little consequence for the resulting key sizes and underlying hardness assumptions, at least asymptotically. However, certain applications (like the GapSVP-based cryptosystems of [Pei08]) in some cases rely on an exponentially large $q \approx 2^n$, in which case the extra $\log q$ factor increases the key size significantly.

Theorem 3.1. *There is a probabilistic polynomial-time algorithm that, on input a positive integer n (in unary), positive integer $q \geq 2$ (in binary), and a $\text{poly}(n)$ -bounded positive integer $m \geq 3(1 + \delta)n \lg q$ for some $\delta > 0$, outputs a pair $(\mathbf{A} \in \mathbb{Z}_q^{m \times n}, \mathbf{S} \in \mathbb{Z}^{m \times m})$ such that:*

- \mathbf{A} is $(m \cdot q^{-\delta n/2})$ -uniform over $\mathbb{Z}_q^{m \times n}$,
- \mathbf{S} is a basis of $\mathcal{L}^\perp(\mathbf{A})$, and
- For any $\omega(\sqrt{\log n})$ function, $\|\mathbf{S}\| \leq m \cdot \omega(\sqrt{\log n})$ with all but $n^{-\omega(1)}$ probability.

Theorem 3.2. *There is a probabilistic polynomial-time algorithm that, on input the parameters n , q , and m as above with $m \geq 2n \lg^2 q$, outputs a pair (\mathbf{A}, \mathbf{S}) as above, where*

- $\|\mathbf{S}\| \leq 5\sqrt{n \lg q}$ for every $i \in [m]$.

The remainder of this section is devoted to proving the theorems.

3.1 Parity Check and Hermite Normal Form

As a warm-up to motivate the construction, we first consider how a given parity check matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ relates to the Hermite normal form of the lattice $\mathcal{L}^\perp(\mathbf{A})$. One may imagine that the rows $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{Z}_q^n$ of \mathbf{A} are uniformly random and independent, though most of the discussion below applies to arbitrary \mathbf{A} .

We determine the HNF matrix $\mathbf{H} \in \mathbb{Z}^{m \times m}$ of the lattice $\Lambda = \mathcal{L}^\perp(\mathbf{A})$ inductively from the bottom up. Starting with the m th row $\mathbf{h}_m = (0, \dots, 0, h_{m,m}) = h_{m,m} \cdot \mathbf{e}_m \in \mathbb{Z}^m$, it must be the case that

$$\mathbf{h}_m \cdot \mathbf{A} = h_{m,m} \cdot \mathbf{a}_m = \mathbf{0} \in \mathbb{Z}_q^n,$$

because every row of \mathbf{H} must be in Λ . Let $k \leq q$ be the smallest positive integer solution to $k \cdot \mathbf{a}_m = \mathbf{0} \in \mathbb{Z}_q^n$. Then $k \cdot \mathbf{e}_m \in \Lambda$, so we must be able to write $k \cdot \mathbf{e}_m = \sum_{i \in [m]} z_i \mathbf{h}_i$ for some integers z_i . Now because $h_{i,i} > 0$ for every $i \in [m]$, it must therefore be the case that $z_i = 0$ for all $i < m$, which implies $h_{m,m} = k$.

Observe that when \mathbf{a}_m is uniformly random, we typically have $h_{m,m} = q$, but other values of $h_{m,m}$ are also possible. For example, if q is even and every entry of \mathbf{a}_m also happens to be even, then we would have $h_{m,m} \leq q/2$.

More generally, suppose that we have determined $\mathbf{h}_{i+1}, \dots, \mathbf{h}_m$ for some $1 \leq i < m$. Then by similar reasoning, $\mathbf{h}_i \in \mathbb{Z}^m$ is given by the unique solution to the equation

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

in which $h_{i,i} > 0$ is minimized and $0 \leq h_{i,j} < h_{j,j} \leq q$ for every $j > i$. To illustrate further, let $M_{i+1} \subseteq \mathbb{Z}_q^n$ be the subgroup of \mathbb{Z}_q^n generated by (all integer linear combinations of) $\mathbf{a}_{i+1}, \dots, \mathbf{a}_m$. Then if $\mathbf{a}_i \in M$, we have $\mathbf{a}_i = \sum_{j=i+1}^m z_j \mathbf{a}_j$ for some integers z_j , so $h_{i,i} = 1$, $h_{i,j} = -z_j \bmod h_{j,j}$, and $M_i = M_{i+1}$. On the other hand, if $\mathbf{a}_i \notin M$, then we have $1 < h_{i,i} \leq q$ and $M_i \supsetneq M_{i+1}$. Note that once $M_i = \mathbb{Z}_q^n$, we have $h_{i',i'} = 1$ and $h_{i',j'} = 0$ for every $i' < j' < i$.

Now suppose that \mathbf{A} is uniformly random, and that $d = (1 + \delta)n \lg q \leq m$ for some positive constant $\delta > 0$. Let $m' = m - d$, and break \mathbf{A} into two matrices $\mathbf{A}_1 \in \mathbb{Z}_q^{m' \times n}$ and $\mathbf{A}_2 \in \mathbb{Z}_q^{d \times n}$, where \mathbf{A}_1 consists of the first m' rows of \mathbf{A} and \mathbf{A}_2 consists of the remaining d . It can be shown (e.g., using the leftover hash lemma) that the rows of \mathbf{A}_2 generate the *entire* group \mathbb{Z}_q^n with overwhelming probability over the choice of \mathbf{A}_2 . So almost all lattices $\mathcal{L}^\perp(\mathbf{A})$ have an HNF of the form

$$\mathbf{H} = \left[\begin{array}{c|c} \mathbf{I}_{m'} & \mathbf{H}_1 \\ \hline \mathbf{0} & \mathbf{H}_2 \end{array} \right], \quad (3.1)$$

where $\mathbf{H}_2 \in \mathbb{Z}_q^{d \times d}$ is the Hermite normal form of the lattice $\mathcal{L}^\perp(\mathbf{A}_2) \subset \mathbb{Z}^d$, which has determinant q^n . Note that there is a bijection between \mathbb{Z}^d modulo \mathbf{H}_2 (formally, the group $\mathbb{Z}^d / (\mathbb{Z}^d \cdot \mathbf{H}_2)$) and \mathbb{Z}_q^n , given by $\phi(\mathbf{h}) = \mathbf{h} \cdot \mathbf{A}_2 \in \mathbb{Z}_q^n$. Note also that because $\mathbf{H} \cdot \mathbf{A} = \mathbf{0} \in \mathbb{Z}_q^{m' \times n}$, we have $\mathbf{A}_1 = -\mathbf{H}_1 \cdot \mathbf{A}_2 \in \mathbb{Z}_q^{m' \times d}$. Therefore, the rows of \mathbf{A}_1 are uniformly random if and only if the rows of \mathbf{H}_1 are uniformly random modulo \mathbf{H}_2 . Our construction (described below) produces a basis \mathbf{S} of short vectors whose HNF has the above form and a nearly identical probability distribution.

3.2 The Block Structure

To guarantee that the HNF matrix \mathbf{H} of our constructed basis \mathbf{S} has the desired structure and distribution, we design \mathbf{S} together with the unimodular matrix \mathbf{U} relating it to \mathbf{H} . We first set up the basic block structures of \mathbf{S} and \mathbf{U} according to the principal equation $\mathbf{S} = \mathbf{U} \cdot \mathbf{H}$. We then make a few simplifying choices and extract a few constraints on the blocks, and specify the blocks so as to satisfy these constraints.

Our construction first chooses $\mathbf{A}_2 \in \mathbb{Z}_q^{d \times n}$ uniformly at random and computes the HNF \mathbf{H}_2 of the induced lattice $\mathcal{L}^\perp(\mathbf{A}_2) \subseteq \mathbb{Z}^d$. Recall that \mathbf{H}_2 is nonsingular and $|\det(\mathbf{H}_2)| \leq q^n$ (note that it will usually be the case that \mathbf{A}_2 generates all of \mathbb{Z}_q^n and $|\det(\mathbf{H}_2)| = q^n$, though we do not need this fact explicitly.) Following the form of \mathbf{H} in (3.1), we obtain the following block structure on \mathbf{S} and \mathbf{U} , where we have named the blocks of \mathbf{S} for convenience.

$$\mathbf{S} = \begin{bmatrix} \mathbf{B} & \mathbf{D} \\ \mathbf{P} & \mathbf{V} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{B} & \mathbf{U}_{1,2} \\ \mathbf{P} & \mathbf{U}_{2,2} \end{bmatrix}}_{\mathbf{U}} \times \underbrace{\begin{bmatrix} \mathbf{I}_{m'} & \mathbf{H}_1 \\ \mathbf{0} & \mathbf{H}_2 \end{bmatrix}}_{\mathbf{H}} \quad (3.2)$$

Strictly speaking, our construction of \mathbf{S} and \mathbf{U} does not correspond to an \mathbf{H} that is in full normal form; specifically, some entries of \mathbf{H}_1 might exceed their corresponding diagonal entries in \mathbf{H}_2 . This is not a problem, because the rows of \mathbf{H}_1 can always be reduced modulo \mathbf{H}_2 via additional unimodular operations. But it is not even necessary to compute this reduced form of \mathbf{H}_1 in our algorithm; instead, it suffices to output \mathbf{S} , \mathbf{A}_2 , and $\mathbf{A}_1 = -\mathbf{H}_1 \cdot \mathbf{A}_2 \in \mathbb{Z}_q^{m' \times n}$, and to show that the joint distribution of $(\mathbf{A}_1, \mathbf{A}_2)$ is nearly uniform.

One of the most sensitive conditions to satisfy is to make \mathbf{U} unimodular. Because we only care about \mathbf{H}_1 modulo \mathbf{H}_2 , the particular choices of the rightmost blocks $\mathbf{U}_{1,2}$ and $\mathbf{U}_{2,2}$ are not of much consequence. For convenience, we make \mathbf{U} *block lower-triangular*, setting $\mathbf{U}_{1,2} = \mathbf{0}$ and $\mathbf{U}_{2,2} = -\mathbf{I}_d$, which implies that \mathbf{B} must be unimodular. Substituting these choices, we obtain the following constraints.

$$\mathbf{H}_1 = \mathbf{B}^{-1} \cdot \mathbf{D} \quad (3.3)$$

$$\mathbf{V} + \mathbf{H}_2 = \mathbf{P} \cdot \mathbf{H}_1 = \mathbf{P} \cdot \mathbf{B}^{-1} \cdot \mathbf{D} \quad (3.4)$$

Note that the left-hand sides of the above equations have large entries, while we need all the blocks of \mathbf{S} to have small entries. The \mathbf{B}^{-1} term will therefore bear the sole responsibility for generating large entries. Note also the common term $\mathbf{H}_1 = \mathbf{B}^{-1} \cdot \mathbf{D}$ that appears in both equations, which causes tension between the two constraints: while we need \mathbf{H}_1 to be nearly uniform modulo \mathbf{H}_2 , we also need to be able to construct \mathbf{P} with small entries so that $\mathbf{P} \cdot \mathbf{H}_1$ closely approximates the matrix \mathbf{H}_2 that is imposed upon us.

To resolve this tension, we write \mathbf{H}_1 as the sum of two matrices, a random matrix \mathbf{R} and a deterministic “structured” matrix \mathbf{G} :

$$\mathbf{H}_1 = \mathbf{B}^{-1} \cdot \mathbf{D} = \mathbf{G} + \mathbf{R}.$$

- Each row r_i of \mathbf{R} is an independent, uniformly random vector in $\{0, 1\}^d$ with random sign. We show using the leftover hash lemma that $\mathbf{R} \cdot \mathbf{A}_2$ is nearly uniform in $\mathbb{Z}_q^{m' \times n}$, hence so is $\mathbf{A}_1 = -\mathbf{H}_1 \cdot \mathbf{A}_2$.
- The matrix \mathbf{G} is designed so that small integer combinations of its rows may be assembled to produce (a matrix close to) \mathbf{H}_2 ; more specifically,

$$\mathbf{P} \cdot \mathbf{G} = \mathbf{H}'_2 = \mathbf{H}_2 - \mathbf{I}_d$$

for some \mathbf{P} having small entries (we subtract \mathbf{I}_d from \mathbf{H}_2 simply for convenience, to put the diagonals of \mathbf{H}'_2 in the range $\{0, \dots, q-1\}$). Furthermore, \mathbf{G} and \mathbf{B} are designed together to make $\mathbf{B} \cdot \mathbf{G}$ have small entries, so that

$$\mathbf{D} = \mathbf{B} \cdot \mathbf{H}_1 = \mathbf{B} \cdot \mathbf{G} + \mathbf{B} \cdot \mathbf{R}$$

has small entries as well.

We then let $\mathbf{V} = \mathbf{P} \cdot \mathbf{R} - \mathbf{I}_d$; observe that \mathbf{V} has small entries because \mathbf{P} , \mathbf{R} , and \mathbf{I}_d do, and that (3.4) is satisfied because

$$\mathbf{P} \cdot \mathbf{H}_1 = \mathbf{P} \cdot (\mathbf{G} + \mathbf{R}) = \mathbf{H}_2 + \mathbf{V}.$$

3.3 Building the Blocks

Here we list the principal constraints on the as-yet undefined matrices \mathbf{B} , \mathbf{P} , and \mathbf{G} from the above discussion, and show how to satisfy those constraints.

1. Matrix \mathbf{B} must be unimodular and have small entries.
2. The product $\mathbf{W} = \mathbf{B} \cdot \mathbf{G}$ must have small entries.
3. We must satisfy $\mathbf{P} \cdot \mathbf{G} = \mathbf{H}'_2 = \mathbf{H}_2 - \mathbf{I}_d$ for some \mathbf{P} with small entries.

Below we give two constructions, corresponding to Theorems 3.1 and 3.2, respectively. In both constructions, we assemble \mathbf{B} from copies of a certain component matrix $\mathbf{T}_k \in \mathbb{Z}^{k \times k}$, which is defined to be the $k \times k$ lower-triangular matrix with 1s along the diagonal, -2 s directly below the diagonal, and 0s elsewhere, i.e., $t_{i,i} = 1$ for $i \in [k]$ and $t_{i+1,i} = -2$ for $i \in [k-1]$. It may be verified that \mathbf{T}_k is lower triangular and unimodular. Moreover, its inverse \mathbf{T}_k^{-1} has a very useful form: its (i, j) th entry is 2^{i-j} for every $i \geq j$, and zero elsewhere.

3.3.1 Construction for Theorem 3.1

Define $m' = m - d \geq 2d$. The basic idea is to construct $\mathbf{G} = \mathbf{B}^{-1} \cdot \mathbf{W} \in \mathbb{Z}^{m' \times d}$ so that it contains enough power-of-2 multiples of each of the standard basis vectors in \mathbb{Z}^d ; this is done by assembling \mathbf{B} from copies of \mathbf{T}_k and letting \mathbf{W} have small entries, thus satisfying constraint 2. Then any vector in \mathbb{Z}^d with bounded entries (specifically, every row of \mathbf{H}'_2) can be expressed as a binary combination of the rows of \mathbf{G} , thus satisfying constraint 3.

We now proceed in more detail. Recall that we are given $\mathbf{H}_2 \in \mathbb{Z}^{d \times d}$; say its diagonal entries (from top to bottom) are r_1, \dots, r_d , and recall that their product is (at most) q^n . Let $\ell_j = \lceil \lg r_j \rceil \leq 1 + \lg r_j$, and define the partial sums $s_0 = 0$, $s_j = s_{j-1} + \ell_j$ for $j \in [d]$, and define the total sum $s = s_d \leq d + n \lg q \leq m'$.

Define $\mathbf{B} \in \mathbb{Z}^{m' \times m'}$ to be the block diagonal matrix

$$\mathbf{B} = \text{diag}(\mathbf{T}_{\ell_1}, \dots, \mathbf{T}_{\ell_d}, \mathbf{I}_{m'-s}),$$

i.e., the direct sum of \mathbf{T}_{ℓ_j} for $j \in [d]$, plus an identity matrix of the appropriate remaining dimension. Observe that \mathbf{B} is lower triangular and unimodular, and that $\mathbf{B}^{-1} = \text{diag}(\mathbf{T}_{\ell_1}^{-1}, \dots, \mathbf{T}_{\ell_d}^{-1}, \mathbf{I}_{m'-s})$.

Now define \mathbf{W} so that $\mathbf{w}_{s_{j-1}+1} = \mathbf{e}_j \in \mathbb{Z}^d$ for each $j \in [d]$, and $\mathbf{w}_i = \mathbf{0}$ elsewhere. Recalling that $\mathbf{G} = \mathbf{B}^{-1} \cdot \mathbf{W}$, one can then check that for each $j \in [d]$ and each $k \in [\ell_j]$, we have

$$\mathbf{g}_{s_{j-1}+k} = 2^{k-1} \cdot \mathbf{e}_j \in \mathbb{Z}^d$$

(and $\mathbf{g}_i = \mathbf{0}$ for $s < i \leq m'$).

Because \mathbf{G} has such a useful form, satisfying constraint 3 (i.e., making $\mathbf{P} \cdot \mathbf{G} = \mathbf{H}_2 - \mathbf{I}_d$) is straightforward. For each $j \in [d]$, every entry of the j th column of \mathbf{H}'_2 is in $\{0, \dots, r_j - 1\}$, by construction of \mathbf{H}_2 . Therefore, each row of \mathbf{H}'_2 can be represented as a binary combination of rows $\mathbf{g}_1, \dots, \mathbf{g}_s$ of \mathbf{G} . These binary combinations are specified in the natural way via the d rows of \mathbf{P} , and we have satisfied constraint 3 where each entry of \mathbf{P} has magnitude at most 1.

3.3.2 Construction for Theorem 3.2

Define $m' = m - d \geq d \cdot \lceil \lg q \rceil$. The basic idea is to construct $\mathbf{G} = \mathbf{B}^{-1} \cdot \mathbf{W} \in \mathbb{Z}^{m' \times d}$ so that \mathbf{G} itself contains the rows of \mathbf{H}'_2 , which can then be trivially selected by very short rows \mathbf{p}_i having length 1 (rather than almost \sqrt{m} as above). To do this, we let \mathbf{B} be made up of copies of \mathbf{T}_k much like above, and let \mathbf{W} encode the binary representation of each row of \mathbf{H}'_2 . Note that \mathbf{H}'_2 has d rows with entries that can be as large as $q - 1$, so we can represent it in binary using $d \cdot \lceil \lg q \rceil \leq m'$ rows. (More generally, using the base- r analog of \mathbf{T}_k instead of base 2, we can represent \mathbf{H}'_2 using $d \cdot \log_r q$ rows, at the expense of using vectors \mathbf{b}_i having length $O(r)$.)

Define $\ell = \lceil \lg(q - 1) \rceil$ and define $\mathbf{B} \in \mathbb{Z}^{m' \times m'}$ be the block diagonal matrix

$$\mathbf{B} = \text{diag}(\mathbf{T}_\ell, \dots, \mathbf{T}_\ell, \mathbf{I}_{m'-d\ell})$$

(where the above expression includes d copies of \mathbf{T}_ℓ). Observe that \mathbf{B} is lower triangular and unimodular, and that $\mathbf{B}^{-1} = \text{diag}(\mathbf{T}_\ell^{-1}, \dots, \mathbf{T}_\ell^{-1}, \mathbf{I}_{m'-d\ell})$.

We now define \mathbf{W} . Let $\mathbf{h}'_j \in \mathbb{Z}^d$ denote the j th row of \mathbf{H}'_2 , and observe that every entry of \mathbf{h}'_j is nonnegative and at most $q - 1$, so it can be written in binary using ℓ bits. Therefore \mathbf{h}'_j can be seen as the ℓ th row of $\mathbf{T}_\ell^{-1} \cdot \mathbf{W}_j$ for a binary matrix $\mathbf{W}_j \in \{0, 1\}^{\ell \times d}$, where the rows of \mathbf{W}_j consist of the coordinate-wise bits of \mathbf{h}'_j from most significant down to least significant. Finally, let $\mathbf{W} \in \mathbb{Z}^{m' \times d}$ be the vertical block matrix consisting of \mathbf{W}_1 through \mathbf{W}_d , followed by the zero matrix of dimension $(m' - d \cdot \ell) \times d$. Then for $\mathbf{G} = \mathbf{B}^{-1} \cdot \mathbf{W}$, it is apparent from the above discussion that row $\mathbf{g}_{j \cdot \ell} = \mathbf{h}'_j$ for each $j \in [d]$. The corresponding rows of \mathbf{P} are $\mathbf{p}_j = \mathbf{e}_{j \cdot \ell} \in \mathbb{Z}^{m'}$ for $j \in [d]$.

3.4 Analysis

We now prove that the above constructions satisfy the claims in Theorems 3.1 and 3.2, respectively. We have already shown by construction that \mathbf{S} is a basis of $\mathcal{L}^\perp(\mathbf{A})$. It remains to show that the distribution of \mathbf{A} is statistically close to uniform over $\mathbb{Z}_q^{m \times n}$, and that the rows of \mathbf{S} are all relatively short (in both constructions).

3.4.1 Distribution of \mathbf{A}

Recall that in both constructions, \mathbf{A} is of the form

$$(\mathbf{A}_1 = -\mathbf{H}_1 \cdot \mathbf{A}_2, \mathbf{A}_2) = (-(\mathbf{G} + \mathbf{R}) \cdot \mathbf{A}_2, \mathbf{A}_2) \in \mathbb{Z}_q^{m \times n},$$

where $\mathbf{A}_2 \in \mathbb{Z}_q^{d \times n}$ is uniform, \mathbf{G} is deterministic, and each row of \mathbf{R} is independent and uniform from $\{0, 1\}^d$ (with random sign).

We claim that $\{h_{\mathbf{A}_2} : h_{\mathbf{A}_2}(\mathbf{r}) = \mathbf{r} \mathbf{A}_2\}$ is a family of 2-universal hash functions from domain $\{0, 1\}^d$ to range \mathbb{Z}_q^n . First, note that $\mathbf{r} \mathbf{A}_2 = \mathbf{r}' \mathbf{A}_2$ if and only if $(\mathbf{r} - \mathbf{r}') \mathbf{A}_2 = \mathbf{0}$, and $\mathbf{0} \neq \mathbf{r} - \mathbf{r}' \in \{0, \pm 1\}$ for any distinct $\mathbf{r}, \mathbf{r}' \in \{0, 1\}^d$. Fix such \mathbf{r}, \mathbf{r}' , and suppose that they differ in their i th entry. Finally, observe that

$$\Pr_{\mathbf{A}_2}[(\mathbf{r} - \mathbf{r}') \mathbf{A}_2 = \mathbf{0}] = q^{-n} = 1/|\mathbb{Z}_q^n|,$$

by averaging over any fixed choice of all but the i th row of \mathbf{A}_2 .

Now because $d = (1 + \delta)n \lg q$ for some constant $\delta > 0$, Lemma 2.2 (the leftover hash lemma) and the triangle inequality imply that $(\mathbf{R} \cdot \mathbf{A}_2, \mathbf{A}_2)$ is $(m \cdot q^{-\delta n/2})$ -uniform over $\mathbb{Z}_q^{m \times n}$, as desired.

3.4.2 Length of \mathbf{S}

We need to analyze the lengths of the rows of \mathbf{B} , \mathbf{P} , \mathbf{D} , and \mathbf{V} , where

$$\begin{aligned}\mathbf{D} &= \mathbf{BG} + \mathbf{BR} \\ \mathbf{V} &= \mathbf{PR} - \mathbf{I}_d\end{aligned}$$

- In both constructions, $\mathbf{BG} = \mathbf{W}$ is a binary matrix (or in the base- r generalization of Theorem 3.2, an r -ary matrix). Thus $\|\mathbf{BG}\| \leq \sqrt{d}$ (more generally, $(r-1)\sqrt{d}$).
- We have $\|\mathbf{R}\| \leq \sqrt{d}$ by construction, and the ℓ_1 norm (i.e., the sum of the absolute values of each entry) of each \mathbf{b}_i is at most 3 (more generally, at most $r+1$). So by the triangle inequality, we have $\|\mathbf{BR}\| \leq 3\sqrt{d}$ (more generally, $(r+1)\sqrt{d}$).
- Note that $\|\mathbf{V}\| \leq \|\mathbf{PR}\| + 1$ by the triangle inequality.

It remains to analyze $\|\mathbf{PR}\|$ for the two constructions. In the construction for Theorem 3.2, each \mathbf{p}_i has just a single 1 entry (and 0s elsewhere), so $\|\mathbf{PR}\| \leq \sqrt{d}$. Putting all the blocks of \mathbf{S} together, we conclude that in the construction for Theorem 3.2, $\|\mathbf{S}\| \leq (2r+1)\sqrt{d}$, as desired.

We now analyze the construction for Theorem 3.1. Let s be the random variable corresponding to any entry of \mathbf{PR} . Because \mathbf{P} is a fixed binary matrix, s is the sum of at most m independent random variables $r_{i,j}$ that individually have expectation 0 (because the sign of each \mathbf{r}_i is random) and magnitude at most 1. By the Hoeffding bound, we have $|s| \leq t \cdot \sqrt{m}$ except with probability at most $\exp(-\Omega(t^2))$. Setting $t = \omega(\sqrt{\log n})$ and taking a union bound over all $\text{poly}(n)$ entries of \mathbf{PR} , we conclude that $\|\mathbf{PR}\| \leq t \cdot \sqrt{m \cdot d} \leq t \cdot m$, except with probability $n^{-\omega(1)}$, as desired.

References

- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In *ICALP*, pages 1–9, 1999.
- [Ajt04] Miklós Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in *STOC* 1996.
- [GGH96] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(42), 1996.
- [GGH97] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Public-key cryptosystems from lattice reduction problems. In *CRYPTO*, pages 112–131, 1997.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [Mic01] Daniele Micciancio. Improving lattice based cryptosystems using the Hermite normal form. In *CaLC*, pages 126–145, 2001.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in *FOCS* 2004.

- [MV03] Daniele Micciancio and Salil P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *CRYPTO*, pages 282–298, 2003.
- [MW01] Daniele Micciancio and Bogdan Warinschi. A linear space algorithm for computing the Hermite normal form. In *ISSAC*, pages 231–236, 2001.
- [Ngu99] Phong Q. Nguyen. Cryptanalysis of the Goldreich-Goldwasser-Halevi cryptosystem from Crypto '97. In *CRYPTO*, pages 288–304, 1999.
- [NR06] Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. In *EUROCRYPT*, pages 271–288, 2006.
- [Pei08] Chris Peikert. Public key cryptosystems from the worst-case shortest vector problem. In submission, 2008.
- [PV08] Chris Peikert and Vinod Vaikuntanathan. Noninteractive statistical zero-knowledge proofs for lattice problems. In *CRYPTO*, pages 536–553, 2008.
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571, 2008.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.