An Improved Worst-Case to Average-Case Connection for Lattice Problems

(extended abstract)

Jin-Yi Cai * Ajay P. Nerurkar †

Abstract

We improve a connection of the worst-case complexity and the average-case complexity of some well-known lattice problems. This fascinating connection was first discovered by Ajtai [1] in 1996. We improve the exponent of this connection from 8 to $3.5 + \epsilon$.

*Department of Computer Science, State University of New York at Buffalo, Buffalo, NY 14260. Research supported in part by NSF grants CCR-9319393 and CCR-9634665, and an Alfred P. Sloan Fellowship. Email: `cai@cs.buffalo.edu`

†Department of Computer Science, State University of New York at Buffalo, Buffalo, NY 14260. Research supported in part by NSF grants CCR-9319393 and CCR-9634665. Email: `apn@cs.buffalo.edu`

1 Introduction

A lattice L is a discrete additive subgroup of $\mathbf{R}^n$. There are many fascinating problems concerning lattices, both from a structural and from an algorithmic point of view [12, 20, 11, 13]. The study of lattice problems can be traced back to Gauss, Dirichlet and Hermite, among others [8, 6, 14]. The subject was first conceived as a bridge between geometry and Diophantine approximation and the theory of quadratic forms. The field *Geometry of Numbers* was christened by Minkowski when he proved his fundamental theorems on shortest vectors and successive minima. In recent years, there is enormous interest in the algorithmic aspects of the theory, especially in connection with basis reduction [18, 23], algorithmic Diophantine approximation and combinatorial optimization [11], integer programming [19], volume estimation for convex bodies [7, 21, 15] and, cryptography [1, 2, 10, 17].

There is an inherent beauty in many problems in the theory of *Geometry of Numbers*. Moreover, major algorithmic progress in the field, such as Lovász's basis reduction algorithm, has had a tremendous impact on many other subjects (e.g., integer programming [19], or the disproof of the Mertens conjecture [22]). However, underlying so much fascination and activity is the belief, yet not a proof, that many of the well-known algorithmic problems for lattices are computationally hard for P.

Regarding NP-hardness, Lagarias [16] showed that the shortest vector problem is NP-hard for the l_∞-norm, but it is not known whether it is NP-hard under any other l_p norm. Van Emde Boas [24] showed that finding the nearest vector is NP-hard under all l_p norms, $p \geq 1$. From [3] it is known that finding an approximate solution to within any constant factor for the nearest vector problem for any l_p norm, and, for the shortest vector problem in the l_∞-norm, are both NP-hard. There are no known polynomial-time algorithms to find approximate solutions to these problems within any polynomial factor, even probabilistically. The celebrated Lovász basis reduction algorithm finds a short vector within a factor of $2^{n/2}$ in P-time. Schnorr's algorithm gets a bound of $(1+\epsilon)^n$, but the running time badly depends on ϵ in the exponent [23]. Babai gave an algorithm that approximates the nearest vector by a factor of $(3/\sqrt{2})^n$ [4].

The recent breakthrough by Ajtai [1] has its motivations from cryptography, and the connection between average-case and worst-case complexity in general. It has been realized for some time that the security of a cryptographic protocol depends on the intractability of a certain computational problem *on the average*. As noted by Ajtai [1], the most desirable guarantee for the required security would be a mathematical proof of hardness, either in an asymptotic sense or for specific values of parameters. Unfortunately as yet we have no such proofs for any problem in NP. The next best thing to an absolute lower bound would be a proof that breaking the protocol is NP-hard. However, if one can have neither, as it is currently the case, then as the next alternative, one would like to have a cryptographic protocol based on a sufficiently "famous" problem, such as factoring, for which the most able minds have labored long and hard, and have found no polynomial time algorithms. It is suggested in [1] that some of the well-known lattice problems also fit this description.

Note that, however, even a proof of hardness for a certain problem usually only refers to its worst-case complexity, and says nothing about its average-case complexity. Thus, e.g., even a proof that factoring is not solvable in P or in BPP, or is NP-hard, would not imply that it is hard on the average. (In fact, in some reasonable sense half the integers are divisible by 2. Moreover, it is not known to follow from any hardness assumption for factoring in the worst-case that factoring numbers of the form pq, for p and q primes, is hard on the average.)

In a beautiful paper, Ajtai [1] found the first explicit connection between, in a certain technical sense, the worst-case and the average-case complexity of the shortest lattice vector problem. Ajtai [1] established, among other things, a reduction from the problem of approximating a short lattice basis in the worst-case, to the problem of finding a short lattice vector v for a uniformly chosen lattice in a certain random class of lattices. The reduction is probabilistic. Moreover, the connection involves a rather large polynomial factor blow-up n^c. This factor represents the increase from the length of v in the average-case, to the approximation factor of the length of the longest basis vector computed, with respect to the best basis possible, in the worst case. (Technical definitions will be given in Section 2.) More precisely, for a certain naturally defined class of lattices Λ Ajtai showed that:

> If there is a prob. poly-time algorithm $\mathcal{A}$ which finds a short vector v of length at most n, for a uniformly chosen lattice in the class Λ indexed by n, with non-trivial probability, then, there is a prob. poly-time algorithm $\mathcal{B}$ which, given any lattice of dimension N by its basis vectors, will find with high probability, a good basis $b_1, b_2, \ldots, b_N$, such that the maximum length $\max_{1 \le i \le N} ||b_i||$ is within a fixed polynomial factor N^c from the best possible over all bases.

This factor N^c is the crucial performance guarantee in any intended application to cryptography. It is also of intrinsic interest as it is the provable connection between the worst-case and the average-case complexity of lattice problems over these naturally defined lattices. It is this exponent c that we improve significantly in this paper. While no explicit value of c was given in [1], Ajtai's proof shows that the exponent c can be taken to be 8. This might look rather large, nonetheless it is the first time such a reduction is proved for a problem in NP. [1] Our main result is to improve this exponent c from 8 to around 3.5 .

Theorem 1 *For any constant $\epsilon > 0$, if there exists a probabilistic polynomial time algorithm $\mathcal{A}$ that when given a random lattice $\Lambda(X)$, indexed by n, m, q, where $q = \Theta(n^3)$ and $m = \Theta(n)$, with probability $\frac{1}{n^{O(1)}}$ returns a vector of the lattice $\Lambda(X)$ of length $\le n$, then, there exists a probabilistic polynomial time algorithm $\mathcal{B}$ which when given a basis $a_1, \ldots, a_n \in \mathbf{Z}^n$ for a lattice $L = L(a_1, \ldots, a_n)$, outputs another basis for L, $b_1, \ldots, b_n$, so that,*

$$\max_{i=1}^{n} \|b_i\| \le \Theta(n^{3.5+\epsilon}) \min_{\text{all bases } b'_1,\ldots,b'_n \text{ for } L} \max_{i=1}^{n} \|b'_i\|.$$

Our algorithm $\mathcal{B}$ has a similar structural design as that of Ajtai's, but many of the steps and their proofs are different. The heart of the reduction is an iterative process on a set of independent lattice vectors S. Assuming that the current set S is not already sufficiently short compared to the shortest basis possible, this process successively replaces S with another independent set where the longest member is reduced by a constant factor in length. One

[1]Related to this are the known random self-reducibilities for problems such as discrete logarithm modulo a prime p, or quadratic residue modulo $m = pq$, or the Permanent function. In the first two cases, the random self-reducibility only applies for a fixed p, resp. m; for varying choices of p, resp. m, there is no known random reduction. For the Permanent, the problem is not known, nor believed to be, in NP.

starts with the given set of basis vectors which defines the lattice L. When a sufficiently short independent set of lattice vectors is at hand, we convert it to a basis with a loss of a factor at most $\sqrt{n}$. It is not explicitly tested whether one has reached a sufficiently short independent set of lattice vectors, one simply tries the main iterative process till, probabilistically, no progress is being made, and then with high probability the current set S is already sufficiently short.

Now we outline the main iterative process. First we construct a suitable parallelepiped called a pseudo-cube with lattice points as vertices. Here we use a different rounding procedure with Gram-Schmidt orthogonalization instead of the tiling by fundamental domains used by Ajtai. This results in better geometric properties, which translate to a significant reduction in the exponent c. We next partition this pseudo-cube into a large number of sub-pseudo-cubes which forms a tiling of the whole pseudo-cube. Now we must handle the main difficulties in the proof of Theorem 1, with a series of technical lemmas. The basic idea is to prove that each sub-pseudo-cube has roughly the same number of lattice points. The proof relies heavily on the geometric properties of the set-up, in terms of *eigenvalues* and *singular values*. A recent theorem of Keith Ball [5], which gives a precise upper bound on the volume of the intersection of any hyperplane with the unit cube, also plays a role in the proof.

Once we achieved a reasonable level of uniformity in the number of lattice points in each sub-pseudo-cube, we devise a sampling procedure that samples with exact uniform distribution all the lattice points in the pseudo-cube, and thus inducing a distribution close to being uniform on the set of sub-pseudo-cubes. We then further *uniformize* this distribution by amplification techniques, so that the resulting distribution is almost uniform, and can be used as an input source X for the presumed algorithm $\mathcal{A}$. For each sub-pseudo-cube we choose its *center* as an *address*, and decompose each lattice vector in a sub-pseudo-cube into the sum of the address vector and a *remainder* vector. This choice of the center enables us to prove that the expectation of the *remainder* vector is zero. Finally, by a change of order of summation, a short lattice vector of $\Lambda(X)$ produced by the algorithm $\mathcal{A}$ will, with high probability, produce a short lattice vector of L as a linear combination of *remainder* vectors. Moreover, Ball's theorem [5] implies that one can get independent short lattice vectors by repeating this process.

Based on the reduction in [1], Ajtai and Dwork [2] have proposed a public-key cryptosystem with provable security guarantees based on worst-case hardness assumption. Another public-key system based on lattice problems was proposed in [10], although no proof was given for that system assuming worst-case hardness. As the security proof of Ajtai-Dwork system is based on the proof in [1], our result will automatically improve the security bound.

The paper is organized as follows. In Section 2, we will give the definitions, state some preliminary lemmas and describe a sampling procedure for sampling a lattice point uniformly. In Section 3, we will describe our algorithm. In Section 4 we will present some geometric and probabilistic theorems concerning volume estimates and number of lattice points, and discuss pseudorandom amplification of randomness. We also mention some additional results. Most proofs are omitted from this extended abstract due to space limitation.

2 Definitions and Preliminaries

We denote by $\mathbf{R}$ the field of real numbers, by $\mathbf{Z}$ the ring of integers and by $\mathbf{Z}_q$ the ring of integers mod q.

The Euclidean norm is denoted by $\|\cdot\|$. The Frobenius norm $\|A\|_F$ of a matrix A is $\|A\|_F = \sqrt{\sum_{i,j} |a_{ij}|^2}$. The length of a set of vectors is defined as the length of the longest vector in the set.

$P(v_1, \ldots, v_n) = \{\sum_{i=1}^n \beta_i v_i \mid \forall i \; 0 \le \beta_i \le 1\}$ denotes the parallelepiped defined by $v_1, \ldots, v_n$. $P^-(v_1, \ldots, v_n)$ is the half-open parallelepiped defined by $v_1, \ldots, v_n$, *i.e.*, $\{\sum_{i=1}^n \beta_i v_i \mid \forall i \; 0 \le \beta_i < 1\}$. The *volume* vol $(P(v_1, \ldots, v_n))$ of a parallelepiped $P(v_1, \ldots, v_n)$ is $|\det\,(v_1, \ldots, v_n)|$. The *minimal height* H of a parallelepiped $P = P(v_1, \ldots, v_n)$ is the minimum value of the ratio $\frac{V}{F_i}$, where $V =$ vol (P) and F_i is the volume of the face of P defined by $v_1, \ldots, v_{i-1}, v_{i+1}, \ldots v_n$, *i.e.*, the $(n-1)$-dimensional volume of the parallelepiped $P(v_1, \ldots, v_{i-1}, v_{i+1}, \ldots v_n)$.

If $a_1, \ldots, a_n$ are linearly independent vectors in $\mathbf{R}^n$ then the set of all integral linear combinations of the a_i forms a (n-dimensional) lattice, denoted by $L(a_1, \ldots, a_n)$, and the a_i are called a basis of that lattice. A lattice can also be abstractly defined as a discrete additive subgroup of $\mathbf{R}^n$. For a lattice L, det L denotes the *determinant* of the lattice. If $b_1, \ldots, b_n$ is a basis for L, then, det $L =$ vol $(P(b_1, \ldots, b_n)) = |\det(b_1, \ldots, b_n)|$. It is invariant under a change of basis. The length of the shortest non-zero vector in L is denoted by sh(L). Define the *length* of a basis $b_1, \ldots, b_n$ as $\max_{i=1}^n \|b_i\|$. Then bl(L) denotes the minimum of the lengths of all bases of L. Let $a_1^*, \ldots, a_n^*$ be the Gram-Schmidt orthogonalization of $a_1, \ldots, a_n$, then $B = P^-(a_1^*, \ldots, a_n^*) - \sum_{i=1}^n \frac{1}{2} a_i^* = \{\sum_{i=1}^n \gamma_i a_i^* \mid -\frac{1}{2} \le \gamma_i < \frac{1}{2}\}$ is called a *fundamental brick* of the lattice. Note that vol $B =$ vol $(P(a_1^*, \ldots, a_n^*)) =$ det L.

Let $\mathbf{Z}_q^{n \times m}$ denote the set of $n \times m$ matrices over $\mathbf{Z}_q$. For every n, m, q, $\Omega_{n,m,q}$ denotes the uniform distribution on $\mathbf{Z}_q^{n \times m}$. For every $X \in \mathbf{Z}_q^{n \times m}$, the set $\Lambda(X) = \{y \in \mathbf{Z}^m \mid Xy \equiv 0 \bmod q\}$ defines a lattice of dimension m. $\Lambda = \Lambda_{n,m,q}$ denotes the probability space of lattices consisting of $\Lambda(X)$ by choosing X according to $\Omega_{n,m,q}$. By Minkowski's Theorem it can be proved that, $\forall c \; \exists c'$ s.t. $\forall \Lambda(X) \in \Lambda_{n,c'n,n^c} \; \exists v \; (v \in \Lambda(X)$ and $0 < \|v\| \le n)$. (In fact, the bound n can be reduced to $n^{1/2+\epsilon}$. The important point is, the bound $\|v\| \le n$ implies that the assumption on the hypothetical algorithm $\mathcal{A}$ is not vacuous.) We now present some preliminary lemmas.

Lemma 1 *Let $u_1, \ldots, u_n$ be linearly independent vectors in a lattice L satisfying $\|u_i\| \le M$. Then any vector $w \in \mathbf{R}^n$ can be expressed as a sum of two vectors, v and δ, where $v \in L$ and $\|\delta\| \le \frac{\sqrt{n}M}{2}$. Moreover if all the vectors are integral, then v and δ can be computed in polynomial time.*

Lemma 2 *Let $L = L(a_1, \ldots, a_n)$ be a lattice. Let B be the fundamental brick of L. Then the whole space $\mathbf{R}^n$ can be tiled up as a disjoint union of copies of B:*

$$\mathbf{R}^n = \bigcup_{l \in L} (B + l).$$

The next lemma proves that if $n \ge 4$ then from a short set of linearly independent vectors in a lattice one can construct a basis with only a $\frac{\sqrt{n}}{2}$ blow-up in size.

Lemma 3 *Let $L = L(a_1, \ldots, a_n)$ be a lattice in $\mathbf{Z}^n$. Let $r_1, \ldots, r_n$ be linearly independent vectors in L with $\|r_i\| \le M$. Then a basis $b_1, \ldots, b_n$ of L can be constructed in P-time so that for every i, $r_i = \sum_{j=1}^i \alpha_{ij} b_j$ where the α_{ij} are integers, $\alpha_{ii} > 0$, and $\|b_i\| \le \max\{1, \frac{\sqrt{n}}{2}\} M$.*

Next we present an algorithm that samples lattice points uniformly from the half open parallelepiped $P = P^-(v_1, \ldots, v_n)$ where v_i are any linearly independent lattice vectors. Let $\chi = \sum_{i=1}^n x_i v_i$, where $0 \le x_i < 1$, be a lattice point in P. By Lemma 3,

$$(x_1\ x_2\ \ldots\ x_n)\begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix} = (x_1\ x_2\ \ldots\ x_n)\begin{pmatrix} \alpha_1 & & & \\ \alpha_{21} & \alpha_2 & & \\ \vdots & \vdots & \ddots & \\ \alpha_{n1} & \alpha_{n2} & \cdots & \alpha_n \end{pmatrix}\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix}.$$

Write χ in terms of b_i, the coefficient of b_n is $x_n\alpha_n$ which has to be an integer. Choose x_n uniformly from the set $\{\frac{0}{\alpha_n}, \frac{1}{\alpha_n}, \ldots, \frac{\alpha_n - 1}{\alpha_n}\}$. Then the coefficient of b_{n-1} is $x_{n-1}\alpha_{n-1} + x_n\alpha_{n,n-1}$. Let x_{n-1}^0 be the root of the equation $x\alpha_{n-1} + x_n\alpha_{n,n-1} = 0$. x_{n-1} is set to the value $(x_{n-1}^0 + y) \bmod 1$ where y is uniformly chosen in $\{\frac{0}{\alpha_{n-1}}, \frac{1}{\alpha_{n-1}}, \ldots, \frac{\alpha_{n-1}-1}{\alpha_{n-1}}\}$, etc. It can be shown that this sampling procedure samples all the lattice points in P uniformly.

3 The Algorithm

As described in Section 1 the heart of the algorithm $\mathcal{B}$ of Theorem 1 is an iterative process $\mathcal{B}'$. At all times we maintain a set S of n linearly independent lattice vectors. At each iteration if the length of S is greater than $n^{3+\epsilon}\ \mathrm{bl}(L)$ then with non-trivial probability we update S, reducing its length by $\frac{1}{2}$. If we are unable to do this at some step then we use Lemma 3 to produce a basis that is at most a $\sqrt{n}$ factor longer. We start this process with $S = \{a_1, \ldots, a_n\}$. The following lemma summarizes this iterative process $\mathcal{B}'$.

Lemma 4 *Let $\epsilon > 0$ be any constant. Assume there exists an algorithm $\mathcal{A}$ that when given a random value X of $\Omega_{n,m,q}$, where $q = \Theta(n^3)$ and $m = \Theta(n)$, with probability greater than $\frac{1}{n^{O(1)}}$ returns a vector of the lattice $\Lambda(X)$ of length $\le n$, then there exists an algorithm $\mathcal{B}'$ which when given two sets of linearly independent vectors $a_1, \ldots, a_n$ and $u_1, \ldots, u_n$, $u_i, a_i \in \mathbf{Z}^n$, $u_i \in L = L(a_1, \ldots, a_n)$, $\|u_i\| \le M$, with high probability returns n linearly independent vectors $b_1, \ldots, b_n$ such that either* $\max_i \|b_i\| \le n^{3+\epsilon}\ \mathrm{bl}(L)$ *or* $\max_i \|b_i\| \le \frac{M}{2}$.

Now we describe the algorithm $\mathcal{B}'$ in detail.

Step 1: Constructing the pseudo-cube

Let e_i be the unit vector that has its i^{th} coordinate 1 and all other coordinates zero. Let $w_i = (n^{1.5}M)e_i$. Thus, w_i are mutually orthogonal vectors and they define a perfect cube of side $n^{1.5}M$. The w_i are not necessarily lattice vectors and we would like to find lattice vectors v_i that are not too far away from the w_i so that the parallelepiped they define is close to a perfect cube. Now applying Lemma 1 each w_i can be written as the sum of two vectors v_i and δ_i such that $v_i \in L$ and $\|\delta_i\| \le \frac{\sqrt{n}M}{2}$. This implies $\|v_i\| \le (n^{1.5} + \frac{\sqrt{n}}{2})M$. As noted in Lemma 1, such v_i and δ_i can be computed efficiently. $P(v_1, \ldots, v_n)$ is the pseudo-cube constructed.

Step 2: Sampling lattice points

We work with an expanded and shifted version of the pseudo-cube constructed in Step 1. Consider the parallelepiped $P = P^-(2v_1, \ldots, 2v_n) - \sum_{i=1}^n v_i = \{\sum_{i=1}^n z_i v_i \mid -1 \le z_i < 1\}$.

We partition P into q^n sub-pseudo-cubes. Assume q is odd. (The case for even q is similar but slightly more involved, and is omitted in this extended abstract.) Consider the sub-pseudo-cube $Q = P^-(\frac{2v_1}{q}, \ldots, \frac{2v_n}{q}) - \sum_{i=1}^n \frac{v_i}{q} = \{\sum_{i=1}^n z_i v_i \mid -\frac{1}{q} \le z_i < \frac{1}{q}\}$. Tile up P with copies of this sub-pseudo-cube, *i.e.* with sub-pseudo-cubes of the form $Q + \sum_{i=1}^n \frac{2t_i}{q} v_i = \{\sum_{i=1}^n \frac{\gamma_i}{q} v_i \mid 2t_i - 1 \le \gamma_i < 2t_i + 1\}$, for integers t_i, $-\frac{q-1}{2} \le t_i \le \frac{q-1}{2}$. Each lattice point in P has an *address* depending on where in P it lies. Here is how we define the address of a lattice point. The sub-pseudo-cube $Q + \sum_{i=1}^n \frac{2t_i}{q} v_i$ has the address $(2t_1 \bmod q, \ldots, 2t_n \bmod q)$. Note that as integers t_i run through $-\frac{q-1}{2}, \ldots, \frac{q-1}{2}$, the reduced moduli $2t_i \bmod q$ run through each value of $\mathbf{Z}_q$ exactly once, since q is odd. Now for a lattice point x, if it lies in the interior of a sub-pseudo-cube, then the address of x is the address of the sub-pseudo-cube. If it lies on the surface of two or more sub-pseudo-cubes, then one sub-pseudo-cube among these is chosen with equal probability and the address is then calculated as above. *However*, in considering which sub-pseudo-cubes share boundary points, any boundary point of P, $x = \sum_{i=1}^n z_i v_i$, with $z_i = -1$ is identified with the point "at the other end" $x' = \sum_{i=1}^n z_i' v_i$ with $z_i' = 1$ and $z_j' = z_j$ for $j \ne i$. Abstractly, we are making an identification on the boundary of P, which can be viewed as taking the quotient space $\mathbf{R}^n$ modulo the lattice $L = L(2v_1, 2v_2, \ldots, 2v_n)$. This identification creates an n-dimensional *torus* out of $P \cong \mathbf{R}^n / L(2v_1, 2v_2, \ldots, 2v_n)$.

For example, if the point sampled is $-v_1 - v_2$ then the following 4 lattice points are chosen with equal probability: $-v_1 - v_2$, $-v_1 + v_2$, $v_1 - v_2$ and $v_1 + v_2$. The address is then $(1, 1, 0, \ldots, 0)$, $(1, q-1, 0, \ldots, 0)$, $(q-1, 1, 0, \ldots, 0)$ and $(q-1, q-1, 0, \ldots, 0)$ respectively, as t_1, and t_2 take on values $-\frac{q-1}{2}$ and $\frac{q-1}{2}$. The address space is $\mathbf{Z}_q^n$. We represent a lattice point l in P by the tuple (σ, δ), where σ is its address and $\delta \in Q$ and $\exists$ unique *even* integers $c_1, \ldots, c_n$, $-(q-1) \le c_i \le q-1$, s.t. $(c_1, \ldots, c_n) \equiv \sigma \bmod q$ and $l = \sum_{i=1}^n \frac{c_i}{q} v_i + \delta$. This δ is called the *remainder* vector of l.

To sample a lattice point in P, first sample a lattice point l uniformly in the parallelepiped $P^-(2v_1, \ldots, 2v_n)$ as described in our sampling algorithm given in Section 2. The uniformly sampled lattice point in P is then $l - \sum_{i=1}^n v_i$. We want to pick independently $m = \Theta(n)$ lattice points in this parallelepiped in such a way that the distribution induced on their addresses is close to uniform. As will be clear in the next section it is not enough for our purposes to just sample m points directly from P. If so, the distribution induced on their addresses will not be as close to the uniform distribution as we want. So we employ a pseudorandom amplification technique as follows. Using our sampling algorithm we first sample $k = \lceil \frac{2}{\epsilon} \rceil$ independent samples from P, $\chi^{(1)}, \chi^{(2)}, \ldots, \chi^{(k)}$ where $\chi^{(j)} = (\sigma^{(j)}, \eta^{(j)})$, and $\sigma^{(1)}, \sigma^{(2)}, \ldots, \sigma^{(k)}$ are the corresponding addresses and $\eta^{(1)}, \eta^{(2)}, \ldots, \eta^{(k)}$ are the remainder vectors. Let $\alpha = \left(\sum_{j=1}^k \sigma^{(j)}\right) \bmod q$, ($\alpha \in \mathbf{Z}_q^n$, with each coordinate reduced modulo q,) and $\eta = \sum_{j=1}^k \eta^{(j)}$. There is a unique sub-pseudo-cube whose center coordinates are congruent to α coordinate-wise modulo q. Let $\chi_c = \sum_{i=1}^n \frac{c_i}{q} v_i$ be the center of this sub-pseudo-cube, where $(c_1, \ldots, c_n) \equiv \alpha \bmod q$. Then $\chi = \chi_c + \eta$ is our constructed lattice point.

It is important to note that χ, while not necessarily equal to $\sum_{j=1}^k \chi^{(j)}$, is always a lattice point. In fact, $\chi^{(j)} = \sum_{i=1}^n \frac{c_i^{(j)}}{q} v_i + \eta^{(j)}$, where $(c_1^{(j)}, c_2^{(j)}, \ldots, c_n^{(j)}) \equiv \sigma^{(j)} \bmod q$. Thus, $\alpha \equiv \sum_{j=1}^k \sigma^{(j)} \equiv (\sum_{j=1}^k c_1^{(j)}, \ldots, \sum_{j=1}^k c_n^{(j)}) \bmod q$. And therefore, $\sum_{j=1}^k c_i^{(j)} \equiv c_i \bmod q$. Finally, notice that v_i are lattice vectors, by exchanging the order of summation,

$$\begin{aligned}\sum_{j=1}^{k} \chi^{(j)} &= \sum_{i=1}^{n} \left(\sum_{j=1}^{k} c_i^{(j)} \right) \frac{v_i}{q} + \sum_{j=1}^{k} \eta^{(j)} \\ &= \sum_{i=1}^{n} \left(\frac{c_i}{q} \right) v_i + \text{ a lattice vector } + \eta.\end{aligned}$$

Since each $\chi^{(j)} \in L$, it follows that $\chi = \sum_{i=1}^{n} \frac{c_i}{q} v_i + \eta$ is also a lattice vector. Note that, χ may lie outside the sub-pseudo-cube $Q + \sum_{i=1}^{n} \frac{c_i}{q} v_i$, or even outside P. But that doesn't matter. We still call α the address of χ and η its remainder vector. We do this m times to get the lattice points χ_i with address α_i and remainder vector η_i, $1 \leq i \leq m$.

Step 3: Calling $\mathcal{A}$

If $M > n^{3+\epsilon}\mathrm{bl}(L)$, the α_i are distributed almost uniformly on the address space (for a proof sketch see next section), and so when the matrix $X = (\alpha_1, \alpha_2, \ldots, \alpha_m)$ is given to algorithm $\mathcal{A}$, with non-trivial probability, it returns a vector $(h_1, h_2, \ldots, h_m)$ in $\Lambda(X)$ of norm $\leq n$.

The output is $g = \sum_{j=1}^{m} h_j \eta_j$. Crucially, g is always a lattice vector of L ([1]). This can be seen using a similar exchange argument as above. We have $\sum_{j=1}^{m} h_j \alpha_j \equiv 0 \bmod q$. Thus, $\sum_{j=1}^{m} h_j c_{ij} \equiv 0 \bmod q$, for each i, where, $(c_{1j}, c_{2j}, \ldots, c_{nj}) \equiv \alpha_j \bmod q$ are the coordinates of the jth sub-pseudo-cube. Hence,

$$\sum_{j=1}^{m} h_j \chi_j = \sum_{i=1}^{n} \left(\sum_{j=1}^{m} \frac{h_j c_{ij}}{q} \right) v_i + \sum_{j=1}^{m} h_j \eta_j = \text{ a lattice vector } + g.$$

This shows that g, being the difference of two lattice vectors, is itself a lattice vector.

We need to repeat the above $\Theta(n)$ times to produce n linearly independent lattice vectors. We next prove that with high probability $\|g\| \leq \frac{M}{2}$ when $q = \Theta(n^3)$ is appropriately chosen.

The expected length of the output

As noted above, when given X, with probability $\frac{1}{n^{O(1)}}$, $\mathcal{A}$ returns a vector $h = (h_1, h_2, \ldots, h_m) \in \Lambda(X)$ of length $\leq n$. In case $\mathcal{A}$ fails to produce such a vector after $n^{O(1)}$ tries we set h to the all-zero vector. So in all cases we can assume that $\|h\| \leq n$.

Let $g = \sum_{i=1}^{m} h_i \eta_i$. We intend to show that with high probability this vector has length $\Theta(\frac{n^3 M}{q})$. Therefore a choice of $q = \Theta(n^3)$ ensures that with high probability g has a length not more than $\frac{M}{2}$. The key is to evaluate the expectation $E\ [\|g\|^2]$.

A different, yet distributionally equivalent, way to uniformly sample lattice points in P is to first choose an address, that is, choose a sub-pseudo-cube, with a probability that is proportional to the number of lattice points in the sub-pseudo-cube and then to uniformly sample a lattice point in that sub-pseudo-cube. This process however cannot be carried out efficiently. But the distribution this induces on the addresses is identical to the one induced by our sampling algorithm. Note that the output of $\mathcal{A}$ depends only on the addresses of the lattice points chosen and not on the remainder vectors. So with this equivalent way of looking at things we can evaluate the expectation $E\ [\|g\|^2]$, by first randomizing α_i, and then for any fixed output $(h_1, h_2, \ldots, h_m)$ by $\mathcal{A}$.

The pseudo-cube P is symmetric about the origin. If x is a lattice point so is $-x$. Since lattice points are chosen uniformly the probability that x is chosen is the same as the probability that $-x$ is chosen. This is also true for lattice points on the boundary of P, where x and $-x$ are chosen with equal probability. Moreover, the set of center points of all sub-pseudo-cubes is also invariant under the map $x \mapsto -x$. This means η and $-\eta$ are equally likely to occur as the remainder vector. Thus $E\ [\eta] = 0$. Now,

$$E\ [\|\sum_{i=1}^{m} h_i \eta_i\|^2] = E\ [\sum_{i,l=1}^{m} h_i h_l \langle \eta_i, \eta_l \rangle] = \sum_{i,l=1}^{m} h_i h_l E\ [\langle \eta_i, \eta_l \rangle].$$

and

$$E\ [\langle \eta_i, \eta_l \rangle] = E\ [\sum_{p,q=1}^{k} \langle \eta_i^{(p)}, \eta_l^{(q)} \rangle] = \sum_{p,q=1}^{k} E\ [\langle \eta_i^{(p)}, \eta_l^{(q)} \rangle].$$

If $p \neq q$ or $i \neq l$, $\eta_i^{(p)}$ and $\eta_l^{(q)}$ are independent, we have $E\ [\langle \eta_i^{(p)}, \eta_l^{(q)} \rangle] = \langle E\ [\eta_i^{(p)}], E\ [\eta_l^{(q)}] \rangle = 0$.

Therefore, $E\ [\| \sum_{i=1}^{m} h_i \eta_i \|^2] = \sum_{i=1}^{m} h_i^2 E\ [\langle \eta_i, \eta_i \rangle] = \sum_{i=1}^{m} h_i^2 \sum_{p=1}^{k} E\ [\|\eta_i^{(p)}\|^2]$. It can be shown that (proof omitted) the diagonal of Q has length at most $O(\sqrt{n}(\frac{n^{1.5}M}{q}))$. Therefore, $\|\eta_i^{(p)}\|$ is at most half that. Using this as an upper bound for $E\ [\|\eta_i^{(p)}\|^2]$, and by Markov's inequality, we get with high probability $\|g\| = O(\frac{n^3 M}{q})$, and thus $\|g\| \leq \frac{M}{2}$ when q is chosen to be $\Theta(n^3)$.

4 Some geometric and probabilistic lemmas

In this section we give some sketch of the proof that our sampling procedure which samples lattice points uniformly from a pseudo-cube induces a distribution close to uniform on the addresses. The key to this proof is some volume estimate using eigenvalue and singular value techniques. The volume bounds will then be used to estimate the number of lattice points in a pseudo-cube which in turn will be used to show that if M is larger than $n^{3+\epsilon}$ bl(L), the distribution induced on the address space by our sampling algorithm is close to the uniform distribution. But, it is not close enough! We will then use amplification to reduce the distance between the two distributions. We will also prove an upper bound on the number of lattice points lying on a hyperplane intersecting a pseudo-cube. This will be used to show that a small number of independent tries are sufficient to produce n linearly independent lattice vectors. Most of the Lemmas and proofs are omitted due to space limitation. We first state and prove a lemma about the volume of a pseudo-cube that is close to a unit cube. Any pseudo-cube can be suitably scaled down and this lemma applies.

Lemma 5 *Let $e_1, \ldots, e_n$ be the standard unit orthogonal vectors. Let $u_1, \ldots, u_n$ be linearly independent vectors such that $\|u_i - e_i\| \leq \epsilon$. Then*

$$1 - n\epsilon \ \leq\ \mathrm{vol}(P(u_1, \ldots, u_n)) \ \leq\ (1+\epsilon)^n.$$

Proof The upper bound is an easy consequence of Hadamard's inequality. Since $\|u_i\| \leq 1+\epsilon$, $\mathrm{vol}(P(u_1, \ldots, u_n)) = |\det(u_1, \ldots, u_n)| \leq \prod_{i=1}^{n} \|u_i\| \leq (1+\epsilon)^n$.

To prove the lower bound, we note that the matrix $(u_1, \ldots, u_n)$ can be written as a sum of the unit matrix I and a perturbation matrix $A = (a_1, \ldots, a_n)$, i.e., $(u_1, \ldots, u_n) = I + A$. Since the determinant is the product of the eigenvalues, and the i^{th} eigenvalue $\lambda_i(I + A) = 1 + \lambda_i(A)$ for a scalar matrix I, we have vol $(P(u_1, \ldots, u_n)) = |\det(u_1, \ldots, u_n)| = |\prod\ (\lambda_i(I + A))| = |\prod\ (1 + \lambda_i(A))| = \prod\ |1 + \lambda_i(A)| \geq \prod\ (1 - |\lambda_i(A)|)$.

By Schur's decomposition, there exists a unitary matrix U s.t. UAU^* is an upper triangular matrix. Since U is unitary, $\|Ua_i\| = \|a_i\|$. Thus, $\|UA\|_F^2 = \sum_{i=1}^n \|Ua_i\|^2 = \sum_{i=1}^n \|a_i\|^2 = \|A\|_F^2$. Similarly $\|UAU^*\|_F^2 = \|UA\|_F^2 = \|A\|_F^2$, since U^* is also unitary. Furthermore, a unitary transformation $A \mapsto UAU^*$ preserves eigenvalues, hence $\lambda_i(A)$ appear on the diagonal of UAU^*. Thus, $\sum |\lambda_i|^2 \leq \|UAU^*\|_F^2 = \|A\|_F^2 \leq n\epsilon^2$.

By Cauchy-Schwarz, $\sum |\lambda_i| \leq n\epsilon$. We are left with the problem of minimizing the product $\prod_{i=1}^n (1 - x_i)$ subject to the conditions $x_i \geq 0$, $\sum x_i \leq n\epsilon$. An easy induction shows that the minimum occurs at $x_i = n\epsilon$, $x_j = 0$ for $j \neq i$. Thus vol$(P) \geq 1 - n\epsilon$. □

Denote $\mu = \mathrm{bl}(L)$. The next lemma proves that the number of lattice points in a parallelepiped of volume V is closely approximated by the ratio $\frac{V}{\det L}$, when the minimum height H is large enough compared to μ.

Lemma 6 *Let $L = L(a_1, \ldots, a_n)$ be a lattice in $\mathbf{R}^n$, $\|a_i\| \leq \mu$, $g_1, \ldots, g_n$ linearly independent vectors in $\mathbf{R}^n$, $b \in \mathbf{R}^n$. Let $P_b = b + P(g_1, \ldots, g_n)$. Let k_0 (resp. k_1) be the number of lattice points in P_b (resp. in its interior). Let H be the minimal height and V be the volume of P_b. Then for $j = 0, 1$:*

1. $\left(1 - \frac{2\mu\sqrt{n}}{H}\right)^n \frac{V}{\det L} \leq k_j \leq \left(1 + \frac{2\mu\sqrt{n}}{H}\right)^n \frac{V}{\det L}$.
2. *If in addition, there exist mutually orthogonal vectors $w_1, \ldots, w_n, \|w_i\| = Y, \|w_i - g_i\|/Y = O\left(\frac{1}{n}\right)$. Let F be any hyperplane, then the number of lattice points in $F \cap P_b$ is at most $c(H^{n-1}) 2\mu\sqrt{n} \left(1 + \frac{2\mu\sqrt{n}}{H}\right)^{n-1} (\det L)^{-1}$ for some constant c.*

Denote $W = P_b$. Let W_e be the parallelepiped obtained from W by expanding it by a factor $\left(1 + \frac{2\mu\sqrt{n}}{H}\right)$, and let W_c be W contracted by $\left(1 - \frac{2\mu\sqrt{n}}{H}\right)$. Let $\mathcal{B} = P(a_1^*, \ldots, a_n^*) - \sum_{i=1}^n \frac{1}{2} a_i^*$ be a fundamental brick of this lattice. Tile up the whole space with copies of $\mathcal{B}$. Any two points of $\mathcal{B}$ are at most a distance $\mu\sqrt{n}$ apart. Therefore any brick that intersects W has to lie completely inside W_e, and any brick that intersects the parallelepiped W_c lies completely inside W. Clearly the number of bricks that intersect W is an upper bound, and the number of bricks that lie completely inside W is a lower bound, for the number of lattice points in W.

The proof of part *2.* of Lemma 6 uses the following lemma. Its proof, in turn, uses singular values, Courant-Fischer inequality, and a recent theorem of Ball [5] which states that the volume of the intersection of any hyperplane with the unit cube (in whatever dimension) has the precise upper bound of $\sqrt{2}$. The proof of Lemma 7 is omitted here for space limitation.

Lemma 7 *Let $e_1, \ldots, e_n$ be the standard unit orthogonal vectors. Let $u_1, \ldots, u_n$ be linearly independent vectors such that $\|u_i - e_i\| \leq \epsilon$. Let H be a hyperplane. Then the area of the surface $P(u_1, \ldots, u_n) \cap H$ is at most $\sqrt{2}e(1 + \epsilon)^{n-1}$.*

In the previous section we proved that, with high probability, each output vector g of algorithm $\mathcal{B}'$ is short, ($\|g\| \leq \frac{M}{2}$). In order to eventually output n linearly independent lattice vectors we also need to show that with non-trivial probability the $(j+1)^{st}$ vector output does not lie in the linear span of the previous j vectors. The above lemmas and some additional lemmas can be used to prove that. (Details are omitted here.) We also need the following lemma for the uniformity of distribution on the address space.

Lemma 8 *If $M > n^{3+\epsilon}\mu$ and with an appropriate choice of $q = \Theta(n^3)$, there exists a uniformly distributed random variable ζ which takes values from the address space and with probability greater than $1 - \frac{1}{n^\epsilon}$ agrees with the actual address of a lattice point chosen randomly according to our sampling algorithm.*

We need to sample m lattice points and ensure that the matrix formed by the m addresses as column vectors is close to the uniform distribution so that $\mathcal{A}$ behaves nicely on it. The above lemma says that each column vector directly sampled is close to being uniform. But since $m = \Theta(n)$ the $1 - \frac{1}{n^\epsilon}$ bound above for each of the m addresses is not good enough. To decrease the distance between the two distributions, we construct a lattice point by first sampling $k = \lceil \frac{2}{\epsilon} \rceil$ points and then combining them as described earlier. The next lemma states that the distribution induced on the address space by this amplification is now much better.

Lemma 9 *There exists a uniformly distributed random variable ρ which takes values from the address space and with probability greater than $1 - \frac{1}{n^2}$ agrees with the actual address of a lattice point chosen randomly by thus combining $\lceil \frac{2}{\epsilon} \rceil$ lattice points.*

Apart from constructing a relatively short basis with high probability for any lattice in $\mathbf{Z}^n$, we are also able to use the hypothesis of Theorem 1 to approximate (with high probability) the length of a shortest non-zero vector in any lattice in $\mathbf{Z}^n$ within a better polynomial factor. This is achieved by using an improved connection between a lattice and its dual. The dual L^* of a lattice L in $\mathbf{R}^n$ is defined as $L^* = \{y \mid \forall x \in L \ \langle x, y \rangle \in \mathbf{Z}\}$, where $\langle \cdot, \cdot \rangle$ denotes the standard inner product. We state the connection first.

Lemma 10 *If L is a lattice in $\mathbf{R}^n$ and L^* is its dual lattice then,*

$1 \leq \mathrm{sh}(L)\mathrm{bl}(L^*) \leq cn^{1.5}$ *for some constant c.*

The upper bound above is better than the upper bound proved in [1] by a $\sqrt{n}$ factor. This, together with our improvement in Theorem 1, enables us to prove the following,

Theorem 2 *For any constant $\epsilon > 0$, if there exists a probabilistic polynomial time algorithm $\mathcal{A}$ that when given a random lattice $\Lambda(X)$, indexed by n, m, q, where $q = \Theta(n^3)$ and $m = \Theta(n)$, with probability $\frac{1}{n^{O(1)}}$ returns a vector of the lattice $\Lambda(X)$ of length $\leq n$, then, there exists a probabilistic polynomial time algorithm $\mathcal{C}$ which when given a basis $a_1, \ldots, a_n \in \mathbf{Z}^n$ for a lattice $L = L(a_1, \ldots, a_n)$, returns a number l such that, $l \leq \mathrm{sh}(L) \leq O(n^{5+\epsilon})\ l$.*

Algorithm $\mathcal{C}$ applies algorithm $\mathcal{B}$ from Theorem 1 to L^*, thus obtaining l^*, an approximation to $\mathrm{bl}(L^*)$. Then, $l = \frac{1}{l^*}$ approximates $\mathrm{sh}(L)$ to the claimed factor. A similar improvement is obtained for the n^c-unique shortest vector problem stated in [1].

Acknowledgement

We wish to thank M. Ajtai, T. Cusick, A. Frieze, R. Kannan, J. Komlós, L. Lovász, E. Szemeredi and A. Yao for interesting discussions on the subject. We also thank the participants of the theory seminar run by the first author at University of Buffalo, P. Aduri, A. Agarwal and P. Stanica.

References

[1] M. Ajtai. Generating hard instances of lattice problems. In *Proc. 28th Annual ACM Symposium on the Theory of Computing*, 1996. Full version available from ECCC, *Electronic Colloquium on Computational Complexity*, at `http://www.uni-trier.de/eccc/`.

[2] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. 1996. Available from ECCC, *Electronic Colloquium on Computational Complexity*, at `http://www.uni-trier.de/eccc/`.

[3] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. In *Proc. 34th IEEE Symposium on Foundations of Computer Science (FOCS)*, 1993, 724-733.

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

[5] K. Ball. Cube slicing in $\mathbf{R}^\mathbf{n}$. *Proceedings of the American Mathematical Society*, 97(3):465–473, 1986.

[6] P. G. L. Dirichlet. Über die Reduktion der positiven quadratischen Formen mit drei unbestimmten ganzen Zahlen. *Journal für die Reine und Angewandte Mathematik*, 40:209–227, 1850.

[7] M. Dyer, A. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *Journal of the ACM*, 38(1):1–17, 1991.

[8] C. F. Gauss. *Disquisitiones Arithmeticae.* Transl. by A. A. Clarke. Yale University Press, 1966.

[9] O. Goldreich, S. Goldwasser, and S. Halevi. Collision-free hashing from lattice problems. 1996. Available from ECCC, *Electronic Colloquium on Computational Complexity*, at `http://www.uni-trier.de/eccc/`.

[10] O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. 1996. Available from ECCC, *Electronic Colloquium on Computational Complexity*, at `http://www.uni-trier.de/eccc/`.

[11] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer Verlag, 1988.

[12] P. M. Gruber. *Handbook of Convex Geometry.* Elsevier Science Publishers B.V., 1993.

[13] P. M. Gruber and C. G. Lekkerkerker. *Geometry of Numbers*. North-Holland, 1987.

[14] C. Hermite. Extraits de lettres de M. Ch. Hermite à M. Jacobi sur différents objets de la théorie des nombres. *Journal für die Reine und Angewandte Mathematik*, 40:261–278, 279–290, 291–307, 308–315, 1850.

[15] R. Kannan, L. Lovász, and M. Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13, 1995.

[16] J. C. Lagarias. The computational complexity of simultaneous diophantine approximation problems. In *Proc. 23rd IEEE Symposium on Foundations of Computer Science (FOCS)*, 1982, 32-39.

[17] J. C. Lagarias and A. M. Odlyzko. Solving low-density subset sum problems. In *Proc. 24th IEEE Symposium on Foundations of Computer Science (FOCS)*, 1983, 1-10.

[18] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982.

[19] H. W. Lenstra, Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8:538–548, 1983.

[20] L. Lovász. *An Algorithmic Theory of Numbers, Graphs and Convexity*. SIAM, Philadelphia, 1986.

[21] L. Lovász and M. Simonovits. The mixing rate of Markov chains, an isoperimetric inequality, and computing the volume. In *Proc. 31st IEEE Symposium on Foundations of Computer Science (FOCS)*, 1990. 346–354.

[22] A. Odlyzko and H.J.J. te Riele. Disproof of the Mertens conjecture. *Journal für die Reine und Angewandte Mathematik*, 357:138–160, 1985.

[23] C. P. Schnorr. A hierarchy of polynomial time basis reduction algorithms. *Theory of Algorithms*, pages 375–386, 1985.

[24] P. van Emde Boas. Another NP-complete partition problem and the complexity of computing short vectors in lattices. Technical Report 81-04, Mathematics Department, University of Amsterdam, 1981.