Upper Bound on $\lambda_1(\Lambda^{\perp}(\mathbf{A}))$

Huiwen Jia[1], Yupu Hu[2], Chunming Tang[1], Yanhua Zhang[3]
hwjia@gzhu.edu.cn, yphu@mail.xidian.edu.cn, ctang@gzhu.edu.cn, and yhzhang@zzuli.edu.cn

[1] Key Laboratory of Information Security, School of Mathematics and Information Science, Guangzhou University, Guangzhou, China
[2] ISN Laboratory, Xidian University, Xi'an, China
[3] Zhengzhou University of Light Industry, Zhengzhou, China

Abstract. It has been shown that, for appropriate parameters, solving the SIS problem in the average case is at least as hard as approximating certain lattice problems in the worst case to within polynomial factor $\beta \cdot \widetilde{O}(\sqrt{n})$, where typically $\beta = O(\sqrt{n \log n})$ such that random SIS instances admit a solution. In this work, we show that $\beta = O(1)$, i.e., β is essentially upper-bounded by a constant. This directly gives us a poly-time exhaustive search algorithm for solving the SIS problem (resulting in approximating certain worst-case lattice problems to within $\widetilde{O}(\sqrt{n})$ factor). Although the exhaustive search algorithm is rather inefficient for typical setting of parameters, our result indicates that lattice-based cryptography is not secure, at least in an asymptotical sense. Our work is based on an observation of the lower/upper bounds on the smoothing parameter for lattices.

Keywords: security of lattice-based cryptography, shortest vector problem, worst-cast/average-case reduction, the smoothing parameter

1 Introduction

Lattice-based cryptography has been one of the most attractive area in mathematical cryptography over the decades. In the seminal work [1], Ajtai introduced the small integer solution (SIS) problem which asks, given a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for appropriate parameters, to find a nontrivial combination of columns of $\mathbf{A}$ that sums to zero. More importantly, he showed that the SIS problem in the average case is at least as hard as approximating certain lattice problems in the worst case to within polynomial factor. This remarkable worst-case/average-case connection is of particular importance, because cryptography inherently requires hard problems for which random instances are hard to solve. If one can design provably secure cryptographic constructions based on the SIS problem, then they are infeasible to break, unless all instances of certain lattice problems are easy to solve. Ajtai's breakthrough not only initiates the study of modern lattice-based cryptography, but also motivates the following two lines of research.

Tighter worst-case/average-case reduction. Since up to now there is no polytime algorithms known for solving any of the worst-case problems used in Ajtai's reduction (even by quantum computers), it is reasonable to conjecture that they are hard for any polynomial approximation factor. However, as these lattice problems get easier and easier as the approximation factor increases, it is theoretically interesting and practically important to reduce the factor as small as possible, thus improving the security guarantee of the the underlying cryptographic constructions as much as possible. Therefore, subsequent to [1], several efforts has been made [5, 4, 13, 16, 8, 15] to reduce the reduction factor. Specifically, in [16], Micciancio and Regev reduced the factor down to $\beta \cdot \widetilde{O}(\sqrt{n})$ for the first time, where $\beta \ll q$ is the upper bound on the l_2 norm of the solutions to the SIS instances. As shown in [16], by the pigeon-hole principle, we can set $\beta = \sqrt{m}q^{n/m}$ such that SIS instances are guaranteed to have solutions. For typical choice of parameters, say $\beta = O(\sqrt{n \log n})$, the worst-case/average-case reduction factor can be as small as $\widetilde{O}(n)$, almost linear in n.

Showing hardness of lattice problems. As mentioned above, generally certain worst-case lattice problems are conjectured hard to solve within polynomial approximation factor. In fact, for much smaller factors, a long sequence of works [3, 2, 6, 12, 7, 10, 9] have been made to show their NP hardness. The stat of the art is that, for some $c > 0$, no efficient algorithm can approximate lattice problems to within $n^{c/\log\log n}$, unless P = NP or another unlikely event occurs.

Although there still remains a gap between the factors that certain problems are known to be NP-hard and the factors that the worst-case/average-case reduction can be established, everything seems to be going well. Since if this gap can be filled, then the worst-case/average-case reduction can help us achieve the ambitious goal that constructs cryptosystems whose security is based solely on the P $\neq$ NP conjecture. *But on the contrary, if we are able to solve the* SIS *problem, then the reduction also gives us a possibility to solve the worst-case lattice problems within polynomial factor. Further, it may help us understand the relation between* P *and* NP

Results and main ideas In this work, we show that for commonly used parameters in lattice-based cryptography, the length of the shortest nonzero vector in lattices of the form $\Lambda^{\perp}(\mathbf{A})$ for uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is upper-bounded by a small constant, i.e., $\lambda_1\big(\Lambda^{\perp}(\mathbf{A})\big) = O(1)$, far from the upper bound $\lambda_1\big(\Lambda^{\perp}(\mathbf{A})\big) \leq \beta = \sqrt{n \log n}$ given by the pigeon hole argument used in previous reductions. At first glance, a tighter worst-case/average-case reduction with approximate factor $\widetilde{O}(\sqrt{n})$ immediately follows and lattice-based cryptography seems to enjoy stronger security guarantee. However, our result also directly gives us a simple exhaustive search algorithm with running time polynomial in m (thus in n) that finds the shortest nonzero vector on random lattice $\Lambda^{\perp}(\mathbf{A})$, resulting in approximating certain worst-case problems on any n-dimensional lattice to within $\widetilde{O}(\sqrt{n})$ factor. Although for typical setting of parameters the exhaustive search algorithm is rather inefficient and seems has no effect on the

security of current underlying cryptographic applications, our result still indicates that lattice-based cryptography is not secure, at least in an asymptotical sense.

Our main result is closely related to an observation of the lower/upper bound on the smoothing parameter for lattices. This notion was first introduced by Micciancio and Regev in [16] and plays a crucial role in their reduction. Roughly speaking, the smoothing parameter is the smallest amount of Gaussian noise that can "smooth out" the discrete structure of the lattice. Micciancio and Regev also showed that this parameter could be roughly bounded by λ_n, the nth successive minima of the lattice. In [17], Peikert presented a new bound: the smoothing parameter can be roughly upper-bounded by $1/\lambda_1^{\infty}(\mathcal{L}^*)$, where the denominator is l_∞ norm of the shortest nonzero vector in the dual lattice $\mathcal{L}^*$. Then in [8], Gentry et al. related the smoothing parameter to a certain lattice quantity, called the Gram-Schmidt minimum and defined as $\widetilde{bl}(\mathcal{L}) = \min_{\mathbf{B}} \|\widetilde{\mathbf{B}}\|$, where the minimum is taken over all bases $\mathbf{B}$ of $\mathcal{L}$.

We now start with a simple observation. Let $\mathcal{L}(\mathbf{B})$ be a lattice with basis $\mathbf{B}$ such that $\|\widetilde{\mathbf{B}}\| = \widetilde{bl}(\mathcal{L})$. As in the proof of Lemma 3.1 in [8], a rigid rotations and reflections are applied to the lattice $\mathcal{L}$ (and its dual $\mathcal{L}^*$) such that the orthogonal Gram-Schmidt vectors $\widetilde{\mathbf{b}}_i$ are parallel to the standard basis vectors $\mathbf{e}_i \in \mathbb{R}^n$. Here we denote by $\bar{\mathcal{L}}(\bar{\mathbf{B}})$ the transformed version of $\mathcal{L}(\mathbf{B})$, i.e., $\widetilde{\bar{\mathbf{b}}}_i = \|\widetilde{\mathbf{b}}_i\| \cdot \mathbf{e}_i$. Notice that such transformations do not affect the values of the quantities defined with respect to the l_2 norm, e.g., the smoothing parameter, $\lambda_1(\mathcal{L})$, and $\widetilde{bl}(\mathcal{L})$ et al., because the structures of $\mathcal{L}$ (and $\mathcal{L}^*$) are preserved. However, the quantities relative to the l_∞ norm, say $\lambda_1^{\infty}(\mathcal{L}^*)$, are sensitive to rotations. This simple observation serves as a good inspiration for us: to get a new bound as tight as possible, by the bound in [17], one may rotate the lattices $\mathcal{L}$ (and $\mathcal{L}^*$) such that $\lambda_1^{\infty}(\mathcal{L}^*)$ achieves its maximum value, making $1/\lambda_1^{\infty}(\mathcal{L}^*)$ as small as possible. Actually, when $\mathcal{L}$ is of the form $\Lambda^{\perp}(\mathbf{A})$, our main line of thought is to show that,

$$\lambda_1(\mathcal{L}) \leq \widetilde{bl}(\mathcal{L}) \leq \frac{1}{\lambda_1^{\infty}(\mathcal{L}^*)} \leq 4.$$

Notice that the first and the last inequalities have been shown in [8], so all we need to show is the second one. At first, we tried to prove this inequality directly and made some progress. In more detail, for any lattice $\mathcal{L}$, we can show that $\widetilde{bl}(\bar{\mathcal{L}}) \leq 1/\lambda_1^{\infty}(\bar{\mathcal{L}}^*)$. Combining with the inequality $1/\lambda_1^{\infty}(\bar{\mathcal{L}}^*) \leq \widetilde{bl}(\bar{\mathcal{L}})$ shown in [8], we have $\widetilde{bl}(\bar{\mathcal{L}}) = 1/\lambda_1^{\infty}(\bar{\mathcal{L}}^*)$. However, we cannot go further to show that $\lambda_1^{\infty}(\bar{\mathcal{L}}^*) \geq \lambda_1^{\infty}(\mathcal{L}^*)$. Hence, to get our main result, we have to take a slightly "tortuous" approach.

The main idea underlying our result is simple and clear. Specifically, we show that

$$\widetilde{bl}(\bar{\mathcal{L}}) \cdot \sqrt{\ln(2/\epsilon)/\pi} < \eta_\epsilon(\bar{\mathcal{L}}) = \eta_\epsilon(\mathcal{L}) \leq \frac{\sqrt{\ln(2n(1+1/\epsilon))/\pi}}{\lambda_1^{\infty}(\mathcal{L}^*)}$$

where $\eta_\epsilon(\cdot)$ is the smoothing parameter with respect to any $\epsilon \in (0,1)$ (see Definition 4). Note that the fist inequality indicates that the upper bound given in

[8] is essentially tight. Then we have

$$\lambda_1(\mathcal{L}) \leq \widetilde{bl}(\mathcal{L}) = \widetilde{bl}(\bar{\mathcal{L}}) < \frac{1}{\lambda_1^{\infty}(\mathcal{L}^*)} \cdot \frac{\sqrt{\ln(2n(1+1/\epsilon))/\pi}}{\sqrt{\ln(2/\epsilon)/\pi}}.$$

Notice that the multiplication term attached to $1/\lambda_1^{\infty}(\mathcal{L}^*)$ trends to 1 as $\epsilon \to 0$. As a result, the inequality $\lambda_1(\Lambda^{\perp}(\mathbf{A})) \leq 4$ follows.

Organization In what follows, we start by briefly recalling some basic definitions and results. Then we describe our main work in Sec. 3.

2 Preliminaries

Notations We denote by $\mathbb{Z}$ the integers. Vectors are assumed to be in column form and are written using bold lower-case letters, e.g. $\mathbf{x}$. The ith component of $\mathbf{x}$ will be denoted by x_i. The l_2 and l_∞ norms of a vector $\mathbf{x}$ are $\|\mathbf{x}\| = \sqrt{\sum_i x_i^2}$ and $\|\mathbf{x}\|_\infty = \max_i |x_i|$, respectively. Matrices are written as bold capital letters, e.g. $\mathbf{X}$, and the ith column vector of a matrix $\mathbf{X}$ is denoted by $\mathbf{x}_i$. The length of a matrix $\|\mathbf{X}\| = \max_i \|\mathbf{x}_i\|$ is the norm of its longest column. For notational convenience, we use $\mathbf{X}^{-t} = (\mathbf{X}^{-1})^t$ to represent the inverse and transpose of $\mathbf{X}$.

Throughout the paper all quantities are implicitly functions of the security parameter n. We use standard big-O notation to classify the growth of functions, and say that $f(n) = \widetilde{O}(g(n))$ if $f(n) = O(g(n) \cdot \log^c n)$ for some fixed constant c. A negligible function is an $f(n)$ such that $f(n) = o(n^{-c})$ for every fixed constant c.

Lattices A (full-rank) lattice is defined as the set of all integer combinations of n linearly independent vectors $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\} \subset \mathbb{R}^n$:

$$\mathcal{L}(\mathbf{B}) = \Big\{ \sum_i z_i \cdot \mathbf{b}_i : \ z_i \in \mathbb{Z} \Big\}.$$

The matrix $\mathbf{B}$ is known as a basis of the lattice and is not unique. The determinant $\det(\mathcal{L})$ of the lattice $\mathcal{L}$, given by $|\det(\mathbf{B})|$, is independent of the choice of the basis.

The Gram-Schmidt orthogonalization $\widetilde{\mathbf{B}}$ of $\mathbf{B}$ is defined iteratively in the following way: $\widetilde{\mathbf{b}}_1 = \mathbf{b}_1$, and for each $i = 2, \ldots, n$, $\widetilde{\mathbf{b}}_i$ is the component of $\mathbf{b}_i$ orthogonal to $\mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_{i-1})$. Clearly, $\|\widetilde{\mathbf{b}}_i\| \leq \|\mathbf{b}_i\|$.

The dual lattice $\mathcal{L}^*$ of a lattice $\mathcal{L}$ is the set

$$\mathcal{L}^* = \{\mathbf{x} \in \mathrm{span}(\mathcal{L}) : \ \forall\ \mathbf{y} \in \mathcal{L},\ \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$$

of all vectors that have integer scalar product with all lattice vectors. The dual of a lattice is also a lattice, and if a lattice $\mathcal{L}(\mathbf{B})$ is generated by basis $\mathbf{B}$, then $\mathbf{B}^* = \mathbf{B}(\mathbf{B}^t\mathbf{B})^{-1}$ is a basis for the dual lattice. We have the following fact that relates $\widetilde{\mathbf{B}}$ and $\widetilde{\mathbf{B}}^*$.

Lemma 1 ([8]). *Let $\{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$ be an (ordered) basis. Let $\{\mathbf{d}_n, \ldots, \mathbf{d}_1\}$ be its dual basis in reverse order and $\widetilde{\mathbf{d}}_n, \ldots, \widetilde{\mathbf{d}}_1$ be its Gram-Schmidt orthogonalization (using this order). Then for all $i \in [n]$, $\widetilde{\mathbf{d}}_i = \widetilde{\mathbf{b}}_i / \|\widetilde{\mathbf{b}}_i\|^2$.*

Here we recall the shortest vector problem. For more details and other problems, please refer to [11, 16].

Definition 1 (Shortest Vector Problem, SVP). *Given an arbitrary basis $\mathbf{B}$ of some lattice $\mathcal{L}(\mathbf{B})$, find a shortest nonzero lattice vector, i.e., a vector $\mathbf{v}$ such that $\|\mathbf{v}\| = \lambda_1(\mathcal{L})$, where $\lambda_1(\mathcal{L})$ denotes the length of the shortest vector in $\mathcal{L}$.*

Definition 2 (Decisional Approximate SVP, GapSVP_{γ}). *Given a pair of $(\mathbf{B}, d)$ where $\mathbf{B}$ is a basis of an n-dimensional lattice $\mathcal{L}$ and d is a rational number. In* YES *input $\lambda_1(\mathcal{L}) \leq d$, and in* NO *input $\lambda_1(\mathcal{L}) > \gamma \cdot d$.*

Then we recall the average case problem SIS [1].

Definition 3 (Smallt Integer Solution ($\mathsf{SIS}_{n,m,q,\beta}$) Problem). *Given uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, find a nonzero vector $\mathbf{z} \in \mathbb{Z}^m$ with $\|\mathbf{z}\| \leq \beta$ such that*

$$\mathbf{Az} = \sum_i \mathbf{a}_i \cdot z_i = \mathbf{0} \in \mathbb{Z}_q^n.$$

Equivalently, the SIS problem can be seen as an average-case SVP on the family of "q-ary" m-dimensional lattice

$$\Lambda^{\perp}(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{Az} = \mathbf{0} \in \mathbb{Z}_q^n\} \supseteq q\mathbb{Z}^m.$$

Typically, we often set $\beta = O(\sqrt{n \log q})$ by the pigeon-hole principle such that the SIS problem is hard enough and nontrivial. The following theorem is one of the security foundations of modern lattice-based cryptography.

Theorem 1 (Worst-case/average-case Reduction [16, 8, 15]). *For any m, $\beta = \mathrm{poly}(n)$ and for any $q \geq \beta \cdot \omega(\sqrt{n \log n})$, solving a random instance of $\mathsf{SIS}_{n,m,q,\beta}$ problem with non-negligible probability is at least at hard as approximating certain problems on any n-dimensional lattice to within $\gamma = \beta \cdot \widetilde{O}(\sqrt{n})$ factors.*

Gaussian measures For any vectors $\mathbf{c}$, $\mathbf{x}$ and any $s > 0$, let

$$\rho_{s,\mathbf{c}}(\mathbf{x}) = \exp(-\pi \|(\mathbf{x} - \mathbf{c})/s\|^2)$$

be a Gaussian function centred in $\mathbf{c}$ scaled by a factor of s. we can define the discrete Gaussian distribution over the n-dimensional lattice Λ as

$$\forall\ \mathbf{x} \in \Lambda,\ D_{\Lambda,s,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{s,\mathbf{c}(\mathbf{x})}}{\rho_{s,\mathbf{c}(\Lambda)}} \propto \rho_{s,\mathbf{c}(\mathbf{x})}.$$

The smoothing parameter Micciancio and Regev in [16] proposed a fundamental quantity for lattices called the smoothing parameter.

Definition 4 (The Smoothing Parameter [16]). *For any n-dimensional lattice Λ and positive real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest real $s > 0$ such that $\rho_{1/s}(\Lambda^*) \leq 1 + \epsilon$.*

Note that since $\rho_{1/s}(\Lambda^*)$ is a continuous and strictly decreasing function of s, the '$\leq$' in above definition can be replaced by '='. Micciancio and Regev [16] also showed that this parameter can be roughly bounded by λ_n, the nth successive minima. In [17], Peikert first relates the smoothing parameter of a lattice to the minimum distance of its dual lattice in the l_∞ norm.

Lemma 2 ([17]). *For any n-dimensional lattice Λ and real $\epsilon > 0$, we have*

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

Then for any $\omega(\sqrt{\ln n})$ function, there is a negligible $\epsilon(n)$ for which $\eta_\epsilon(\Lambda) \leq \omega(\sqrt{\ln n})/\lambda_1^\infty(\Lambda^*)$.

In [8], Gentry et al. presented a new smoothing parameter bound, which related the smoothing parameter to the longest Gram-Schmidt vector in any basis of the lattice.

Lemma 3 ([8]). *For any n-dimensional lattice Λ and real $\epsilon > 0$, we have*

$$\eta_\epsilon(\Lambda) \leq \widetilde{bl}(\Lambda) \cdot \sqrt{\ln(2n(1+1/\epsilon))/\pi},$$

where $\widetilde{bl} = \min_{\mathbf{B}} \|\widetilde{\mathbf{B}}\|$ is called the Gram-Schmidt minimum. Then for any $\omega(\sqrt{\ln n})$ function, there is a negligible $\epsilon(n)$ for which $\eta_\epsilon(\Lambda) \leq \widetilde{bl}(\Lambda) \cdot \omega(\sqrt{\ln n})$.

In [8], Gentry et al. also showed that for lattice of the form $\Lambda^\perp(\mathbf{A})$, it has a small smoothing parameter.

Lemma 4 ([8]). *Let n and q be positive integers with q prime, and let $m \leq 2n \log q$. Then for all but an at most q^{-n} fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we have $\lambda_1^\infty\big((\Lambda^\perp(\mathbf{A}))^*\big) \geq 1/4$.*

In particular, for such $\mathbf{A}$ and for any $\omega(\sqrt{\log m})$, there is a negligible function $\epsilon(m)$ such that $\eta_\epsilon\big(\Lambda^\perp(\mathbf{A})\big) \leq \omega(\sqrt{\log m})$.

3 Upper Bound on $\lambda_1\big(\Lambda^\perp(\mathbf{A})\big)$

For any lattice $\mathcal{L}(\mathbf{B})$ with basis $\mathbf{B}$ such that $\|\widetilde{\mathbf{B}}\| = \widetilde{bl}(\mathcal{L})$, denote by $\bar{\mathcal{L}}(\bar{\mathbf{B}})$ its transformed version (by applying rigid rotations and reflections as in the proof of Lemma 3.1 in [7]), such that the orthogonal Gram-Schmidt vectors $\widetilde{\bar{\mathbf{b}}}_i$ are parallel to the standard basis vectors $\mathbf{e}_i$, i.e., $\widetilde{\bar{\mathbf{b}}}_i = \|\widetilde{\bar{\mathbf{b}}}_i\| \cdot \mathbf{e}_i$. Notice that such transformations do not affect the values of the quantities defined with respect to the l_2 norm, e.g., $\eta_\epsilon(\mathcal{L}) = \eta_\epsilon(\bar{\mathcal{L}})$ and $\widetilde{bl}(\mathcal{L}) = \widetilde{bl}(\bar{\mathcal{L}}) = \|\widetilde{\bar{\mathbf{B}}}\|$, because the structures of $\mathcal{L}$ (and $\mathcal{L}^*$) are preserved.

Theorem 2 (Main Result). *Let n, q be positive integers with q prime, and let $m \geq 2n\log q$. Then for uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, we have $\lambda_1(\Lambda^{\perp}(\mathbf{A})) = O(1)$ with overwhelming probability.*

Proof. As mentioned before, for any lattice $\mathcal{L}$, if we can show that

$$\widetilde{bl}(\bar{\mathcal{L}}) \cdot \sqrt{\ln(2/\epsilon)/\pi} < \eta_\epsilon(\bar{\mathcal{L}}) = \eta_\epsilon(\mathcal{L}) < \frac{\sqrt{\ln(2n(1+1/\epsilon))/\pi}}{\lambda_1^{\infty}(\mathcal{L}^*)},$$

i.e.,

$$\widetilde{bl}(\bar{\mathcal{L}}) < \frac{1}{\lambda_1^{\infty}(\mathcal{L}^*)} \cdot \frac{\sqrt{\ln(2n(1+1/\epsilon))/\pi}}{\sqrt{\ln(2/\epsilon)/\pi}},$$

then by Lemma 4 and the fact that $\sqrt{\ln(2n(1+1/\epsilon))/\pi}/\sqrt{\ln(2/\epsilon)/\pi} \to 1$ as $\epsilon \to 0$, our main result

$$\lambda_1(\mathcal{L}) \leq \widetilde{bl}(\mathcal{L}) = \widetilde{bl}(\bar{\mathcal{L}}) < \frac{1}{\lambda_1^{\infty}(\mathcal{L}^*)} \cdot \frac{\sqrt{\ln(2n(1+1/\epsilon))/\pi}}{\sqrt{\ln(2/\epsilon)/\pi}} = O(1)$$

follows when $\mathcal{L}$ is of the form $\Lambda^{\perp}(\mathbf{A})$.

We now show that $\widetilde{bl}(\bar{\mathcal{L}}) \cdot \sqrt{\ln(2/\epsilon)/\pi} < \eta_\epsilon(\bar{\mathcal{L}})$.

On one hand, $\widetilde{bl}(\bar{\mathcal{L}}) \cdot \sqrt{\ln(2/\epsilon)/\pi} < \eta_\epsilon\big(\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z}\big)$.

By the definition of $\eta_\epsilon(\cdot)$ and the fact that $\rho_{1/s}(\cdot)$ is a continuous and strictly decreasing function of s, we only need to show that $\rho_{1/s}\big((\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z})^*\big) > 1+\epsilon$ for $s = \widetilde{bl}(\bar{\mathcal{L}}) \cdot \sqrt{\ln(2/\epsilon)/\pi}$. Indeed,

$$\begin{aligned}
\rho_{1/s}\big((\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z})^*\big) &= \rho_{1/s}\big(1/\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z}\big) \\
&= \rho\big(\sqrt{\ln(2/\epsilon)/\pi} \cdot \mathbb{Z}\big) \\
&= \sum_{i\in\mathbb{Z}} (\epsilon/2)^{i^2} \\
&> 1+\epsilon.
\end{aligned}$$

On the other hand, $\eta_\epsilon(\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z}) < \eta_\epsilon(\bar{\mathcal{L}})$.

For the same reason, it is enough to show that $\rho_{1/s}\big((\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z})^*\big) < \rho_{1/s}(\bar{\mathcal{L}}^*)$. Let $\bar{\mathbf{B}} = \mathbf{QDU}$, then by the definition of Gram-Schmidt orthogonolization, $\mathbf{Q} = \mathbf{I}$ is the identity matrix, $\mathbf{D}$ is a diagonal matrix with diagonal entries $\{\|\widetilde{\mathbf{b}}_i\|\}_{i=1}^n$ and $\mathbf{U}$ is a upper triangular matrix with 1's on the diagonal.

When $\|\widetilde{\mathbf{b}}_n\| = \widetilde{bl}(\bar{\mathcal{L}})$, we have

$$\begin{aligned}
\rho_{1/s}(\bar{\mathcal{L}}^*) &= \rho_{1/s}(\bar{\mathbf{B}}^{-t} \cdot \mathbb{Z}^n) \\
&= \rho_{1/s}(\mathbf{D}^{-t} \cdot \mathbf{U}^{-t} \cdot \mathbb{Z}^n)
\end{aligned}$$

$$
\begin{aligned}
&> \rho_{1/s}(1/\|\widetilde{\bar{\mathbf{b}}}_n\| \cdot \mathbb{Z}) \\
&= \rho_{1/s}\big((\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z})^*\big),
\end{aligned}
$$

since $\mathbf{D}^{-t}$ is a diagonal matrix with diagonal entries $\{1/\|\widetilde{\bar{\mathbf{b}}}_i\|\}_{i=1}^n$ and $\mathbf{U}^{-t}$ is a lower triangular matrix with 1's on the diagonal.

When $\|\widetilde{\bar{\mathbf{b}}}_j\| = \widetilde{bl}(\bar{\mathcal{L}})$ for some $j < n$, consider the sub-lattice $\bar{\mathcal{L}}_j(\bar{\mathbf{B}}_j)$ of $\bar{\mathcal{L}}(\bar{\mathbf{B}})$ generated by $\bar{\mathbf{B}}_j = (\bar{\mathbf{b}}_1, \ldots, \bar{\mathbf{b}}_j)$. Then we have $\bar{\mathbf{B}}_j = \mathbf{Q}_j \mathbf{D}_j \mathbf{U}_j$, where $\mathbf{Q}_j$ consists of the first j columns of $\mathbf{Q} = \mathbf{I}$, $\mathbf{D}_j$ is a diagonal matrix with diagonal entries $\{\|\widetilde{\bar{\mathbf{b}}}_i\|\}_{i=1}^j$ and $\mathbf{U}_j$ consists of the first $j \times j$ entries of $\mathbf{U}$. Then

$$
\begin{aligned}
\rho_{1/s}(\bar{\mathcal{L}}^*) &= \rho_{1/s}(\mathbf{D}^{-t} \cdot \mathbf{U}^{-t} \cdot \mathbb{Z}^n) \\
&> \rho_{1/s}(\mathbf{Q}_j^t \cdot \mathbf{D}^{-t} \cdot \mathbf{Q}_j \cdot \mathbf{Q}_j^t \cdot \mathbf{U}^{-t} \cdot \mathbf{Q}_j \cdot \mathbb{Z}^j) \\
&= \rho_{1/s}(\mathbf{D}_j^{-t} \cdot \mathbf{U}_j^{-t} \cdot \mathbb{Z}^j) \\
&> \rho_{1/s}(1/\|\widetilde{\bar{\mathbf{b}}}_j\| \cdot \mathbb{Z}) \\
&= \rho_{1/s}\big((\widetilde{bl}(\bar{\mathcal{L}}) \cdot \mathbb{Z})^*\big),
\end{aligned}
$$

where $\mathbf{D}_j^{-t}$ is also a diagonal matrix with diagonal entries $\{1/\|\widetilde{\bar{\mathbf{b}}}_i\|\}_{i=1}^j$, and $\mathbf{U}_j^{-t} \in \mathbb{R}^{j \times j}$ just consists of the first $j \times j$ entries of $\mathbf{U}^{-t}$. This completes the proof of our main result.

Corollary 1. *There is a poly-time exhaustive search algorithm (in security parameter n) for solving the $\mathsf{SIS}_{n,m,q,4}$ problem, where n, m, q are as above. Specifically, by the worst-case/average-case reduction, the algorithm can be used to approximate* GapSVP *(among others) to within factor $\widetilde{O}(\sqrt{n})$ in the worst case.*

Remark 1. This simple algorithm runs in time roughly $\sum_{i=2}^{c^2} 2^{i-1} C_m^i = O(m^{c^2})$ for finding a vector $\mathbf{x} \in \mathbb{Z}^m$ with $\|\mathbf{x}\|_\infty = 1$ and $\|\mathbf{x}\| \leq c \leq 4$. In [14], the authors provide a strengthening version of Lemma 4, which can reduce the complexity of this algorithm and offer trade-offs between m and c.

Acknowledgments

References

1. M. Ajtai, Generating hard instances of lattice problems. Quaderni di Matematica, 13 (2004) 1-32. Preliminary version in STOC 1996.
2. M. Ajtai, The shortest vector problem in l2 is NP-hard for randomized reductions (extended abstract), Proccedings of STOC 1998, 10-19.
3. P. van Emde Boas, Another NP-complete problem and the complexity of computing short vectors in a lattice, Technical report, 1981.
4. J.Y. Cai, Applications of a new transference theorem to Ajtais connection factor. Proccedings of STOC 1999, 205-215.
5. J.Y. Cai, A. Nerurkar, An improved worst-case to average-case connection for lattice problems, Proccedings of FOCS 1997, 468-477.

6. J.Y. Cai, A. Nerurkar, Approximating the SVP to within a factor $(1 + 1/dim^{\epsilon})$ is NP-hard under randomized reductions, J. Comput. System Sci., 59(2):221-239 (1999).
7. I. Dinur, G. Kindler, R. Raz, and S. Safra, Approximating CVP to within almost-polynomial factors is NP-hard, Combinatorica, 23(2):205-243 (2003)
8. C. Gentry, C. Peikert, V. Vaikuntanathan, Trapdoors for hard lattices and new cryptographic constructions, Proccedings of STOC 2008, 197-206.
9. I. Haviv, O. Regev, Tensor-based hardness of the shortest vector problem to within almost polynomial factors, STOC 2007, 469-477
10. S. Khot, Hardness of approximating the shortest vector problem in lattices, Proccedings of FOCS 2004, 126-135.
11. D. Micciancio, S. Goldgasser, Complexity of Lattice Problems: a cryptograhpic perspective, volume 671 of The Kluwer International Series in Engineering and Computer Science. Kluwer Academic Publishers, Boston, Massachusetts, 2002.
12. D. Micciancio, The shortest vector problem is NP-hard to approximate to within some constant, SIAM Journal on Computing, 30(6):2008-2035 (2001). Preliminary version in FOCS 1998.
13. D. Micciancio, Almost perfect lattices, the covering radius problem, and applications to Ajtais connection factor. SIAM Journal on Computing 34 (1) (2004) 118-169. Preliminary version in STOC 2002.
14. D. Micciancio, C. Peikert, Trapdoors for lattices: simpler, tightr, faster, smaller, Proccedings of EUROCRYPT 2012, 700-718.
15. D. Micciancio, C. Peikert, Hardness of SIS and LWE with small parameters, Proccedings of CRYPTO 2013, 21-39.
16. D. Micciancio, O. Regev, Worst-case to average-case reductions based on Gaussian measure, SIAM Journal on Computing 37 (1) (2007) 267-302. Preliminary version in FOCS 2004.
17. C. Peikert, Limits on the hardness of lattice problems in l_p norms, Computational Complexity, 17(2): 300-351, 2007.