The Remote Set Problem on Lattices

Ishay Haviv*

June 11, 2012

Abstract

We initiate studying the *Remote Set Problem* (RSP) on lattices, which given a lattice asks to find a set of points containing a point which is far from the lattice. We show a polynomial-time deterministic algorithm that on rank n lattice $\mathcal{L}$ outputs a set of points at least one of which is $\sqrt{\log n/n} \cdot \rho(\mathcal{L})$-far from $\mathcal{L}$, where $\rho(\mathcal{L})$ stands for the covering radius of $\mathcal{L}$ (i.e., the maximum possible distance of a point in space from $\mathcal{L}$). As an application, we show that the Covering Radius Problem with approximation factor $\sqrt{n/\log n}$ lies in the complexity class NP, improving a result of Guruswami, Micciancio and Regev by a factor of $\sqrt{\log n}$ (Computational Complexity, 2005).

Our results apply to any ℓ_p norm for $2 \leq p \leq \infty$ with the same approximation factors (except a loss of $\sqrt{\log \log n}$ for $p = \infty$). In addition, we show that the output of our algorithm for RSP contains a point whose ℓ_2 distance from $\mathcal{L}$ is at least $(\log n/n)^{1/p} \cdot \rho^{(p)}(\mathcal{L})$, where $\rho^{(p)}(\mathcal{L})$ is the covering radius of $\mathcal{L}$ measured with respect to the ℓ_p norm. The proof technique involves a theorem on balancing vectors due to Banaszczyk (Random Struct. Alg., 1998) and the 'six standard deviations' theorem of Spencer (Trans. AMS, 1985).

1 Introduction

An m-dimensional lattice of rank n is the set of all integer combinations of n linearly independent vectors in $\mathbb{R}^m$ called a basis. Lattices were investigated since the late 18th century by mathematicians, and during the last decades they have also attracted lots of attention from a computational point of view. On one hand, a long line of research shows that many fundamental lattice problems are hard and indicates that it is impossible to solve them in polynomial running time. On the other hand, lattices were shown to be useful as an algorithmic tool as well as applicable in cryptography (see, e.g., [28]). Interestingly, the use of lattices in constructions of cryptographic primitives enjoys strong security relied on the worst-case hardness of certain lattice problems, as was first shown by Ajtai [2]. Therefore, research on algorithms for lattice problems and on their hardness is highly motivated.

There are many important computational problems associated with lattices. The two most fundamental ones are the Shortest Vector Problem (SVP) and the Closest Vector Problem (CVP). In the former, for a lattice given by an *arbitrary* basis we are supposed to find (the length of) a shortest nonzero vector in the lattice. The problem CVP is an inhomogeneous variant of SVP, in which given a lattice and some target point one has to find (the distance from) the closest lattice

*School of Computer Science, The Academic College of Tel Aviv-Yaffo, Israel. Email: ishayhav@mta.ac.il.

point. Another lattice problem of interest is the Covering Radius Problem (CRP) in which given a lattice the goal is to find (a point in space which attains) the maximum possible distance from the lattice. This distance is referred to as the *covering radius* of the lattice. In all problems, the distance is measured relative to some fixed norm on $\mathbb{R}^m$. Usually it is the Euclidean norm ℓ_2 (to which we refer unless otherwise specified) but other ℓ_p norms for $1 \leq p \leq \infty$ are of interest as well (see, e.g., [32]). We note that all the mentioned problems have analogous intensively studied problems in the context of linear codes.

The first polynomial-time approximation algorithm for SVP was presented by Lenstra, Lenstra and Lovász (LLL) in 1982 and achieved an approximation factor of $2^{O(n)}$, where n is the rank of the lattice [23]. Using their algorithm, Babai came up with the nearest plane algorithm achieving the same approximation factor for CVP [6]. A few years later, Schnorr obtained a slightly sub-exponential approximation factor for SVP, namely $2^{O(n(\log\log n)^2/\log n)}$ [33], and this has since been improved by a randomized algorithm of [3]. Kannan presented deterministic algorithms solving SVP and CVP *exactly* requiring running time $n^{O(n)}$ [20], and this was improved to $2^{O(n)}$ more than two decades later by Micciancio and Voulgaris [29]. The algorithm of [29] was recently extended to any ℓ_p norm (and other norms) by Dadush, Peikert and Vempala [12].

On the hardness side, it is known that CVP is NP-hard to approximate to within $n^{c/\log\log n}$ [14] for some constant $c > 0$ and that (under randomized reductions) it is NP-hard to approximate SVP to within any constant [21]. Hardness of approximating SVP to within some $n^{c/\log\log n}$ factor is known to date only assuming some stronger (yet plausible) complexity assumptions [19, 26] (see [13] for stronger results for the ℓ_∞ norm). In contrast to the hardness results, there is a line of research showing limits on the hardness of lattice problems. For example, suitably defined gap versions of both SVP and CVP are known to lie in coNP for approximation factor of $\sqrt{n}$ [1] and in coAM for approximation factor of $\sqrt{n/\log n}$ [15]. Therefore, they are unlikely to be NP-hard to approximate to within $\sqrt{n/\log n}$, as this would imply the collapse of the polynomial-time hierarchy [11]. The results of [1] were extended by Peikert to SVP and CVP in the ℓ_p norm for $2 \leq p \leq \infty$ with essentially the same approximation factors [31].

The study of the Covering Radius Problem on lattices (CRP) from a computational point of view was initiated by Guruswami, Micciancio and Regev in [16]. Previously this problem was used by Micciancio to get tighter connections between the average-case and worst-case complexity of lattice problems [24]. It was shown in [16] that approximating CRP to within $\gamma(n)$ can be done in exponential time $2^{O(n)}$ for any constant $\gamma(n) > 1$ and in polynomial time for some $\gamma(n) = 2^{O(n\log\log n/\log n)}$.[1] In addition, they showed that CRP is in AM for $\gamma(n) = 2$, in coAM for $\gamma(n) = \sqrt{n/\log n}$, and in NP $\cap$ coNP for $\gamma(n) = \sqrt{n}$. Peikert showed in [31] that CRP in the ℓ_p norm for $2 \leq p \leq \infty$ lies in coNP for the same $\sqrt{n}$ approximation factor (except a loss of $\sqrt{\log n}$ for $p = \infty$). However, such an extension to ℓ_p norms is not known for NP and this was left as an open question in [16]. On the hardness side, very little is known. The decisional gap version of CRP (of deciding whether the covering radius is at most some given r) naturally lies in the complexity class Π_2 and is conjectured to be Π_2-hard [24]. However, Π_2-hardness is only known for CRP in the ℓ_p norm for any sufficiently large value of p [18].

Among the results mentioned above regarding CRP, the one saying that CRP is in AM for

[1] To be precise, the algorithms of [16] were randomized since they used randomized algorithms of [4]. However, the deterministic algorithm of [29] implies that the approximation obtained in [16] can be achieved deterministically.

$\gamma(n) = 2$ is unique for this lattice problem. The proof of this fact is relatively simple, and follows from the following AM protocol. Given a lattice $\mathcal{L}$ and a number r, the verifier sends to the prover a uniformly chosen random point in space and the prover has to provide a lattice point whose distance from the random point is at most r. Clearly, if the covering radius is at most r then the prover can act in a way that the verifier accepts with probability 1. On the other hand, the soundness is crucially based on an observation of [16] that random points in space are far from the lattice with high probability. More precisely, a uniformly chosen random point is with constant probability at least $\frac{1}{2} \cdot \rho(\mathcal{L})$-far from a lattice $\mathcal{L}$, where $\rho(\mathcal{L})$ stands for the covering radius of $\mathcal{L}$.

A natural question to ask is whether CRP with $\gamma(n) = 2$ (or with some other factor smaller than $\sqrt{n}$) can be shown to be in NP. Observe that if the verifier could *deterministically* pick a point in space which is quite far from the input lattice, then the protocol above could yield an NP verifier for CRP. Moreover, it can be seen that a deterministic algorithm which outputs polynomially many points at least one of which is quite far from the lattice could suffice for this purpose as well. This challenge is the driving force of the current work, in which we study deterministic polynomial-time algorithms which given a lattice find a set of points containing a point which is far from the lattice.

1.1 Our Contribution

In this paper we initiate studying the Remote Set Problem (RSP) on lattices. This problem can be viewed as a *generalized search* variant of the Covering Radius Problem studied in [24, 16, 18, 17]. In RSP the input is a rank n lattice given by a basis generating it. The goal is to find a set S of points in the span of B containing a point which is far from the lattice. This problem is analogous to a problem suggested for study by Alon, Panigrahy and Yekhanin in the context of linear codes [5] (see Section 1.3 for details).

Recall that the maximum possible distance of a point in space from a lattice $\mathcal{L}$ is called the covering radius of $\mathcal{L}$ and is denoted by $\rho(\mathcal{L})$. The quality of an algorithm for RSP depends on two parameters (to be minimized):

1. the *size* d of the set S constructed by the algorithm, and
2. the *remoteness parameter* which is defined as the minimum $\gamma \geq 1$ for which S contains a point whose distance from $\mathcal{L}$ is at least $\frac{1}{\gamma} \cdot \rho(\mathcal{L})$ for every input lattice $\mathcal{L}$.

As was mentioned before, for every lattice $\mathcal{L}$ a uniformly chosen random point in space has distance at least $\frac{1}{2} \cdot \rho(\mathcal{L})$ from $\mathcal{L}$ with a constant probability [16]. This implies that the efficient algorithm which uniformly and randomly picks a point in space (without even looking at the specific input) solves RSP with $d = 1$ and $\gamma = 2$ with a constant probability of success. Moreover, an algorithm that independently and randomly picks d points and outputs all of them solves RSP with parameters d and $\gamma = 2$ with failure probability which tends to 0 exponentially in d. However, the problem seems much more challenging if we require the algorithm to be *deterministic* (this is also the case for linear codes; see Section 1.3 and [5] for details).

To obtain a deterministic algorithm for RSP one can use an observation made in [7, 16] saying that for every lattice $\mathcal{L}$ there exists a point in $\frac{1}{2} \cdot \mathcal{L}$ whose distance from $\mathcal{L}$ is at least $\frac{1}{2} \cdot \rho(\mathcal{L})$. This implies that the algorithm, which outputs all the linear combinations of the basis vectors with all

coefficients in $\{0, \frac{1}{2}\}$, deterministically solves RSP with $\gamma = 2$. However, the number of points that this algorithm outputs is $d = 2^n$, where n is the rank of the input lattice, and, in particular, its running time is exponential in n.

In this paper we consider the task of finding an algorithm for RSP which is simultaneously deterministic and of polynomial running time. First, we observe that the LLL algorithm [23] can be used to deterministically and efficiently calculate a point whose distance from the lattice approximates the covering radius with an exponential factor.

Theorem 1.1. *There exists a deterministic polynomial-time algorithm for* RSP *with* $d = 1$ *and* $\gamma(n) = 2^{O(n)}$.

Our main result significantly improves the remoteness parameter γ achieved in Theorem 1.1 at the price of having d polynomial in the input size, as stated below.

Theorem 1.2. *There exists a deterministic polynomial-time algorithm for* RSP *with* $\gamma(n) = \sqrt{n/\log n}$.

Notice that the number d of points that the algorithm of Theorem 1.2 outputs is polynomial in the input size, as d clearly cannot be higher than the running time.

As alluded to before, besides being a natural lattice problem, studying RSP is motivated by research on the Covering Radius Problem (CRP). In the promise version of CRP with parameter $\gamma \geq 1$ the input consists of a lattice $\mathcal{L}$ and a number r, and the goal is to decide whether the covering radius $\rho(\mathcal{L})$ of $\mathcal{L}$ is at most r or larger than $\gamma \cdot r$. This problem lies in the complexity class Π_2 (for any γ), since $\rho(\mathcal{L}(B)) \leq r$ if and only if *for all* x in the span of $\mathcal{L}$ there *exists* $y \in \mathcal{L}$ such that the distance between x and y is at most r. For small values of γ the problem is conjectured to be Π_2-hard [24], however it is known that for $\gamma(n) = \sqrt{n}$ it lies in NP [16] (see also [27, Section 7]). In order to prove that CRP with certain $\gamma = \gamma(n)$ is in NP one should come up with an efficiently verifiable witness for instances with $\rho(\mathcal{L}(B)) \leq r$ which does not exist if $\rho(\mathcal{L}(B)) > \gamma \cdot r$. We claim that a deterministic and efficient algorithm for RSP can be useful for this purpose. Indeed, such an algorithm outputs a set S of points at least one of which is quite far from the lattice, hence in order to verify that the covering radius is small it suffices to verify that the points in S are close to the lattice. This can be easily done taking the witness which consists of the lattice points closest to the points in S. We combine this idea with Theorem 1.2 and obtain the following theorem which improves upon the $\sqrt{n}$ factor obtained in [16].

Theorem 1.3. CRP *with approximation factor* $\sqrt{n/\log n}$ *is in* NP.

Another motivation to study RSP comes from the connections between CRP and the Closest Vector Problem (CVP). Known connections between these problems were found useful in several results of [16] regarding the complexity of CRP, namely, the exponential-time approximation algorithm for any $\gamma > 1$ and the proof systems implying that CRP with approximation factors $\sqrt{n}$ and $\sqrt{n/\log n}$ are in coNP and coAM respectively. It turns out that algorithms for RSP imply reductions from CRP to CVP. Specifically, we show that our algorithm for RSP from Theorem 1.2 implies a deterministic rank-preserving polynomial-time Cook reduction from CRP to CVP with $\sqrt{n/\log n}$ loss in the approximation factor (see Corollary 4.5). The only similar result we are aware of is implied by a paper of Micciancio [25] and gives a $\sqrt{n}$ loss in the approximation factor.[2]

[2]Strictly speaking, Micciancio shows in [25] a gap-preserving Cook reduction from the Shortest Independent Vectors Problem (SIVP) to CVP which, combined with known relations between SIVP and CRP, implies a Cook reduction from CRP to CVP with $\sqrt{n}$ loss in the approximation factor.

We also show that Karp reductions from CRP to CVP can be derived from algorithms for RSP. For details see Section 4.2.

In the above discussion RSP and CRP were considered with respect to the Euclidean norm, but it is natural to consider them with respect to any other ℓ_p norm for $1 \leq p \leq \infty$. It is easy to prove that our results can be adapted to arbitrary ℓ_p norm, since in $\mathbb{R}^m$ all ℓ_p norms are within $\sqrt{m}$ from the ℓ_2 one. However, this introduces a $\sqrt{m}$ loss in the approximation factors (where m is the dimension of the lattice). We actually show that this loss is not necessary answering a question asked in [16]. We prove that Theorem 1.2 holds for any ℓ_p norm for $2 \leq p < \infty$. Namely, for every $2 \leq p < \infty$, there exists a deterministic polynomial-time algorithm that given a lattice whose covering radius with respect to the ℓ_p norm is r outputs a set of points guaranteed to contain a point whose ℓ_p distance from the lattice is at least $\sqrt{\log n/n} \cdot r$. Interestingly, we show that our algorithm can also be generalized to the ℓ_p norm in the following manner: given a lattice whose covering radius with respect to the ℓ_p norm is r it outputs a set of points guaranteed to contain a point whose ℓ_2 distance from the lattice is at least $(\log n/n)^{1/p} \cdot r$. Our results are similarly extended to the ℓ_∞ norm and imply a generalization of Theorem 1.3 to every ℓ_p norm for $2 \leq p \leq \infty$.

1.2 Techniques

The algorithm for RSP which yields Theorem 1.2 is quite simple. At the heart of its analysis lies a fact (which was mentioned before) saying that for every lattice $\mathcal{L}$ there exists a point in $\frac{1}{2} \cdot \mathcal{L}$ whose distance from $\mathcal{L}$ is close to its covering radius (see [7, 16] and Lemma 2.1). This suggests a deterministic construction of all the 2^n linear combinations of the basis vectors with all coefficients in $\{0, \frac{1}{2}\}$, since at least one of them is quite far from the lattice. In what follows we explain how we significantly decrease the number of points that the algorithm outputs.

Let us start with the simple algorithm which given a lattice basis $B = (b_1, \ldots, b_n)$ outputs the set $S = \{\frac{1}{2} \cdot b_1, \ldots, \frac{1}{2} \cdot b_n\}$. We claim that at least one of the points in S is quite far from the lattice $\mathcal{L}(B)$ generated by B. Indeed, it can be shown that if all the points in S are of distance at most r from $\mathcal{L}(B)$ then *every* sum of a subset of S is of distance at most $\sqrt{n} \cdot r$ from it. However, there exists a subset of S whose sum has distance from the lattice which is close to its covering radius. This implies that the distance from the lattice of at least one of the points in S approximates the covering radius to within a factor of $\sqrt{n}$.

Intuitively, the above algorithm achieves an approximation factor of $\sqrt{n}$ since there exists a point far from the lattice which can be written as a sum of at most n points in S. A possible approach to improve the $\sqrt{n}$ factor is to output a set S for which a point far from the lattice can be written as a sum of fewer points in S. To do so, our algorithm arbitrarily partitions the n basis vectors into $n/\log n$ sets of size $\log n$ each denoted $B_1, \ldots, B_{n/\log n}$. The algorithm outputs the set S of all points which form a linear combination of vectors that belong to only one of the B_i's with all coefficients in $\{0, \frac{1}{2}\}$. Since the number of vectors in every B_i is logarithmic in n, it is possible to construct these points in polynomial running time. In addition, there is a point whose distance from the lattice is close to the covering radius, which can be written as a sum of at most $n/\log n$ points in S, hence it can be shown that the distance of one of them from the lattice approximates the covering radius to within a factor of $\sqrt{n/\log n}$.

The reasoning above can be extended to any ℓ_p norm for $2 \leq p \leq \infty$. Here, the bound on

the distance from the lattice of a sum of points in S is based on what is known as the "type 2" property of ℓ_p norms for $2 \leq p < \infty$ (see, e.g., [30]) and on the celebrated 'six standard deviations' theorem of Spencer for $p = \infty$ [34]. To obtain additional extensions of our results we use a theorem on balancing vectors involving the ℓ_2 and other ℓ_p norms due to Banaszczyk [8]. The interest in this theorem of Banaszczyk comes from a famous conjecture of Komlós whose implication to our setting is presented in Section 3.2.

Some of the techniques used in the analysis of our algorithm were applied by Banaszczyk in [7] to relate the covering radius of a lattice to another lattice parameter (namely, the nth successive minimum). For the interested reader we overview such relations in Appendix A. In addition, we note that the technique of applying an exponential running-time routine on sublattices of logarithmic rank in order to approximate a parameter of the whole lattice was previously used in lattice algorithms. For example, recall that the LLL algorithm for SVP of [23] is based on a basis reduction which given a lattice basis transforms it into some other basis of the same lattice. Roughly speaking, this reduction causes every two consecutive vectors in the basis to have a certain property, which was shown in [23] to be beneficial for the purpose of approximating SVP. In an improvement of Schnorr [33], the requirement on every two consecutive basis vectors was replaced by a requirement on every (roughly) $\log n$ consecutive basis vectors (namely, being a Korkine-Zolotarev basis of a certain lattice [22]). This enabled him to apply an exponential-time algorithm to every $\log n$ consecutive basis vectors to get a polynomial-time algorithm for SVP with an improved approximation factor.

1.3 The Remote Set Problem on Linear Codes

The current paper is concerned with the lattice analogous of the Remote Set Problem (RSP) on linear codes which was introduced by Alon, Panigrahy and Yekhanin [5]. Here we give a short overview on this problem. In RSP on linear codes, given a linear space $L \subseteq \mathbb{F}_2^n$ of dimension k the goal is to find a set of $O(n)$ points at least one of which is far from L with respect to the Hamming distance. Besides the connection of this problem to the Nearest Codeword Problem (which is the analogous of CVP for linear codes), this problem is motivated by the matrix rigidity approach to circuit lower bounds in computational complexity theory, as explained below.

In 1977, Valiant [35] considered the problem of finding an *explicit* set S of $O(n)$ points in $\mathbb{F}_2^n$ such that for every linear space $L \subseteq \mathbb{F}_2^n$ of dimension $\frac{n}{2}$ there exists a point in S whose Hamming distance from L is at least n^ε for some fixed $\varepsilon > 0$. Whereas such a set is known to exist (even for distance $\Omega(n)$), no polynomial-time deterministic algorithm for constructing such a set is known. Valiant showed that such construction implies a circuit lower bound for an explicit function. Nevertheless, three decades later this problem is still open. This led the authors of [5] to suggest the study of the relaxed question RSP in which the goal is to deterministically find a set containing a point which is far from *one* linear space given as input (instead of satisfying this property for *every* linear space of dimension $\frac{n}{2}$). They presented a deterministic polynomial-time algorithm for RSP achieving remoteness $\Omega(\log n)$. Note that their algorithm is for a special case of RSP, called the Remote Point Problem, in which the algorithm is required to output only one point.

1.4 Open Questions

Our work raises several open questions. It will be interesting to understand for which parameters the Remote Set Problem can be deterministically solved in polynomial time. We have shown that there exists a deterministic polynomial-time algorithm that given a lattice outputs a set of points one of which has distance at least $\sqrt{\log n/n}$ times the covering radius. Can the guarantee on the distance be improved? Can it be improved to the factor $1/2$ for which this can be achieved by a randomized algorithm [16]? Can one achieve the $\sqrt{\log n/n}$ factor (or even a $1/\sqrt{n}$ factor) by an algorithm which outputs only one point instead of polynomially many points? Can it be achieved for the ℓ_p norm for $1 \leq p < 2$?

1.5 Outline

The paper is organized as follows. In Section 2 we gather all the definitions on lattices and computational lattice problems that we need in the paper. In Section 3 we present our algorithms for RSP proving Theorems 1.1 and 1.2 and some extensions of them. In Section 4 we present applications of RSP to CRP including the proof of Theorem 1.3 and reductions from CRP to CVP. In Appendix A we overview some inequalities stemming from [7, 8] involving the covering radius of lattices in ℓ_p norms.

2 Preliminaries

Notations. For $1 \leq p < \infty$ the ℓ_p norm of a vector $x \in \mathbb{R}^m$ is defined as $\|x\|_p = (\sum_{i=1}^m |x_i|^p)^{1/p}$ and for $p = \infty$ it is defined as $\|x\|_\infty = \max_{1\leq i\leq m} |x_i|$. The ℓ_p distance between two vectors $x, y \in \mathbb{R}^m$ is defined as $\mathrm{dist}_p(x,y) = \|x-y\|_p$. The ℓ_p distance of a vector $x \in \mathbb{R}^m$ from a set $S \subseteq \mathbb{R}^m$ is defined as $\mathrm{dist}_p(x,S) = \min_{y\in S} \mathrm{dist}_p(x,y)$. We say that x is *r-far* from S if $\mathrm{dist}_p(x,S) \geq r$. When we omit the subscript p (or a superscript (p)) we refer to the the Euclidean norm ℓ_2.

Lattices. A *lattice* is a discrete additive subgroup of $\mathbb{R}^m$. Equivalently, it is the set of all integer combinations

$$\mathcal{L}(b_1,\ldots,b_n) = \left\{ \sum_{i=1}^n x_i b_i : x_i \in \mathbb{Z} \text{ for all } 1 \leq i \leq n \right\}$$

of n linearly independent vectors $b_1,\ldots,b_n$ in $\mathbb{R}^m$ $(n \leq m)$. If the lattice *rank* n equals its *dimension* m we say that the lattice is *full-rank*. The set $(b_1,\ldots,b_n)$ is called a *basis* of the lattice. Note that a lattice has many possible bases. We often represent a basis by an m by n matrix B having the basis vectors as columns, and we say that the basis B *generates* the lattice $\mathcal{L}$. In such case we write $\mathcal{L} = \mathcal{L}(B)$. The linear space spanned by B is denoted $\mathrm{span}(B) = \{\sum_{i=1}^n x_i b_i : x_i \in \mathbb{R} \text{ for all } 1 \leq i \leq n\}$. A *sublattice* of $\mathcal{L}$ is a lattice $\mathcal{L}(S) \subseteq \mathcal{L}$ generated by some linearly independent lattice vectors $S \subseteq \mathcal{L}$.

One basic parameter of a lattice $\mathcal{L}$, denoted by $\lambda_1^{(p)}(\mathcal{L})$, is the minimum ℓ_p norm of a nonzero vector in it. Equivalently, $\lambda_1^{(p)}(\mathcal{L})$ is the minimum ℓ_p distance between two distinct points in the lattice $\mathcal{L}$. This definition can be generalized to define the ith *successive minimum* as the smallest r

such that $\mathcal{B}^{(p)}(r)$ contains i linearly independent lattice points, where $\mathcal{B}^{(p)}(r)$ denotes the ℓ_p ball of radius r centered at the origin. More formally, $\lambda_i^{(p)}(\mathcal{L}) = \min\{r : \dim(\mathrm{span}(\mathcal{L} \cap \mathcal{B}^{(p)}(r))) \geq i\}$.

Another parameter associated with lattices is the covering radius. For a lattice basis $B = (b_1, \ldots, b_n)$ the *covering radius* of $\mathcal{L}(B)$ with respect to the ℓ_p norm is defined as

$$\rho^{(p)}(\mathcal{L}(B)) = \max_{x \in \mathrm{span}(B)} \mathrm{dist}_p(x, \mathcal{L}(B)).$$

Hence, $\rho^{(p)}(\mathcal{L}(B)) \leq r$ means that for any $x \in \mathrm{span}(B)$ there exists a lattice point $y \in \mathcal{L}(B)$ such that $\mathrm{dist}_p(x, y) \leq r$. Conversely, $\rho^{(p)}(\mathcal{L}(B)) > r$ means that there exists some $x \in \mathrm{span}(B)$ such that any lattice point $y \in \mathcal{L}(B)$ satisfies $\mathrm{dist}_p(x, y) > r$. A *deep hole* of $\mathcal{L}(B)$ is a point $x \in \mathrm{span}(B)$ at distance $\mathrm{dist}_p(x, \mathcal{L}(B)) = \rho^{(p)}(\mathcal{L}(B))$ from the lattice.

The following lemma shows that in order to find a point quite far from a lattice $\mathcal{L}(B)$ it suffices to consider linear combinations of vectors in B with coefficients in $\{0, \frac{1}{2}\}$. This lemma (in more general forms) was proved in [7, 16], and we repeat its proof here for completeness.

Lemma 2.1. *For every $1 \leq p \leq \infty$ and any lattice basis $B = (b_1, \ldots, b_n)$ there exists a vector*

$$v = a_1 \cdot b_1 + \ldots + a_n \cdot b_n$$

with $a_j \in \{0, \frac{1}{2}\}$ for all $1 \leq j \leq n$ such that $\mathrm{dist}_p(v, \mathcal{L}(B)) \geq \frac{1}{2} \cdot \rho^{(p)}(\mathcal{L}(B))$.

Proof: Let w be a deep hole of the lattice $\mathcal{L}(B)$ with respect to the ℓ_p norm. Consider the point $2w$ and observe that, like any point in $\mathrm{span}(B)$, its ℓ_p distance from $\mathcal{L}(B)$ is at most $\rho^{(p)}(\mathcal{L}(B))$. This means that there exists a lattice point $u \in \mathcal{L}(B)$ such that $\mathrm{dist}_p(u, 2w) \leq \rho^{(p)}(\mathcal{L}(B))$ and hence $\mathrm{dist}_p(\frac{1}{2} \cdot u, w) \leq \frac{1}{2} \cdot \rho^{(p)}(\mathcal{L}(B))$. Now, by triangle inequality,

$$\mathrm{dist}_p(\frac{1}{2} \cdot u, \mathcal{L}(B)) \geq \mathrm{dist}_p(w, \mathcal{L}(B)) - \mathrm{dist}_p(\frac{1}{2} \cdot u, w) \geq \frac{1}{2} \cdot \rho^{(p)}(\mathcal{L}(B)).$$

Finally, observe that $\frac{1}{2} \cdot u \in \frac{1}{2} \cdot \mathcal{L}(B)$, so by reducing modulo 1 its coefficients as a linear combination of B, we obtain a vector of the required form with the same ℓ_p distance from $\mathcal{L}(B)$. ∎

For a sequence of vectors $(b_1, \ldots, b_n)$ define the corresponding *Gram-Schmidt orthogonalized vectors* $\tilde{b}_1, \ldots, \tilde{b}_n$ by

$$\tilde{b}_i = b_i - \sum_{j=1}^{i-1} \mu_{i,j} \tilde{b}_j, \quad \mu_{i,j} = \frac{\langle b_i, \tilde{b}_j \rangle}{\langle \tilde{b}_j, \tilde{b}_j \rangle}.$$

In words, $\tilde{b}_i$ is the component of b_i orthogonal to $b_1, \ldots, b_{i-1}$.

Computational Lattice Problems. In what follows we define the lattice computational problems considered in this paper. For an in-depth introduction to the computational aspects of lattices we refer the reader to the book of Micciancio and Goldwasser [27]. For any $1 \leq p \leq \infty$ and any approximation factor $\gamma \geq 1$ (which is usually considered as a function of the lattice rank n) we define the following computational problems.

Definition 2.2 (Covering Radius Problem)**.** *An instance of* $\mathsf{GapCRP}_{\gamma}^{(p)}$ *is a pair* (B, r) *where* $B \in \mathbb{Q}^{m \times n}$ *is a rank* n *lattice basis and* $r \in \mathbb{Q}$ *is a rational number. In* YES *instances* $\rho^{(p)}(\mathcal{L}(B)) \leq r$ *and in* NO *instances* $\rho^{(p)}(\mathcal{L}(B)) > \gamma \cdot r$.

Definition 2.3 (Closest Vector Problem)**.** *An instance of* $\mathsf{GapCVP}_{\gamma}^{(p)}$ *is a triple* (B, t, r) *where* $B \in \mathbb{Q}^{m \times n}$ *is a rank* n *lattice basis,* $t \in \mathbb{Q}^m$ *is a target point, and* $r \in \mathbb{Q}$ *is a rational number. In* YES *instances* $\mathrm{dist}_p(t, \mathcal{L}(B)) \leq r$ *and in* NO *instances* $\mathrm{dist}_p(t, \mathcal{L}(B)) > \gamma \cdot r$.

The main problem under study in the present paper is the Remote Set Problem (RSP), which can be viewed as a generalized search variant of CRP and is defined for any $1 \leq p \leq \infty$ and $d, \gamma \geq 1$ as follows.

Definition 2.4 (Remote Set Problem)**.** *An instance of* $\mathsf{RSP}_{d,\gamma}^{(p)}$ *is a rank* n *lattice basis* $B \in \mathbb{Q}^{m \times n}$*. The goal is to find a set* $S \subseteq \mathrm{span}(B)$ *of size* $|S| \leq d$ *containing a point* v *such that*

$$\mathrm{dist}_p(v, \mathcal{L}(B)) \geq \frac{1}{\gamma} \cdot \rho^{(p)}(\mathcal{L}(B)).$$

Balancing Vectors. The analysis of the main algorithm presented in this work relies on upper bounds on the length of linear combinations with ± 1 coefficients of a given set of vectors. In the following we provide the needed background.

In Banach spaces theory, a normed space X is said to have *type* 2 if there exists a constant $T < \infty$ such that for every n and $x_1, \ldots, x_n \in X$,

$$\Big(\mathbb{E} \, \| \sum_{i=1}^{n} \varepsilon_i \cdot x_i \|_X^2 \Big)^{1/2} \leq T \cdot \big(\sum_{i=1}^{n} \|x_i\|_X^2 \big)^{1/2}, \tag{1}$$

where the expectation is over a uniform choice of signs $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, +1\}$. For example, it is easy to see that the Euclidean space ℓ_2 has type 2, since for ℓ_2 equality holds in (1) with $T = 1$ as follows from the parallelogram law. It is well-known that for every $2 \leq p < \infty$ the ℓ_p normed space has type 2 with $T = c \cdot \sqrt{p}$ for some absolute constant $c > 0$ (see, e.g., [30]). In particular, for every n vectors $x_1, \ldots, x_n$ there exists some choice of signs for which the corresponding linear combination has ℓ_p norm at most $O(\sqrt{n})$ times the maximum ℓ_p norm of the x_i's. This is stated in the following lemma.

Lemma 2.5. *For every* $2 \leq p < \infty$ *there exists a constant* $c_p > 0$ *for which the following holds. For every* n *vectors* $x_1, \ldots, x_n \in \mathbb{R}^m$ *there exist* $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, +1\}$ *such that*

$$\| \sum_{i=1}^{n} \varepsilon_i \cdot x_i \|_p \leq c_p \cdot \sqrt{n} \cdot \max_{1 \leq i \leq n} \|x_i\|_p.$$

A similar statement, motivated by questions on set systems in combinatorial discrepancy, is known for the ℓ_∞ norm. By a simple probabilistic argument it can be seen that every set of n vectors in $\mathbb{R}^m$ has a linear combination with ± 1 coefficients whose ℓ_∞ norm is at most $O(\sqrt{n \log m})$ times the maximum ℓ_∞ norm of the vectors. Interestingly, Spencer showed in 1985 that this can be improved to $O(\sqrt{n \log(2m/n)})$ [34]. For the special case of $m = n$ he showed a bound of $6\sqrt{n}$, commonly referred to as the 'six standard deviations' theorem. In a recent breakthrough, Bansal [9] gave algorithmic results related to Spencer's bound.

Theorem 2.6 ([34]). *There exists a constant $c_\infty > 0$ such that for every n vectors $x_1, \ldots, x_n \in \mathbb{R}^m$ ($m \geq n$) there exist $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, +1\}$ such that*

$$\| \sum_{i=1}^{n} \varepsilon_i \cdot x_i \|_\infty \leq c_\infty \cdot \sqrt{n \cdot \log(2m/n)} \cdot \max_{1 \leq i \leq n} \|x_i\|_\infty.$$

3 Algorithms for the Remote Set Problem

In this section we present our deterministic polynomial-time algorithms for RSP. We first prove Theorem 1.2 (and some extensions) and then turn to consider a special case of RSP, called the Remote Point Problem, proving Theorem 1.1.

3.1 Proof of Theorem 1.2

We start with the following statement from which we derive Theorem 1.2.

Theorem 3.1. *For every $2 \leq p < \infty$ and every $k = k(n) \geq 1$ there exists a deterministic $2^k \cdot s^{O(1)}$ time algorithm for $\mathsf{RSP}^{(p)}_{d,\gamma}$ with $d(n) = O(\frac{n}{k} \cdot 2^k)$ and $\gamma(n) = O(\sqrt{\frac{n}{k}})$, where n denotes the lattice rank and s denotes the input size. The same holds for $p = \infty$ with $\gamma(n, m) = O(\sqrt{\frac{n}{k} \cdot \log(2mk/n)})$, where m denotes the lattice dimension.*

Proof: Assume for simplicity that $k = k(n)$ divides n. We consider the algorithm that given a lattice basis $B = (b_1, \ldots, b_n)$ first partitions its vectors into $\frac{n}{k}$ sets of size k each. Then the algorithm outputs all vectors in space which form a linear combination with all coefficients in $\{0, \frac{1}{2}\}$ of vectors in one of these sets. More precisely, for every $1 \leq i \leq \frac{n}{k}$ let S_i be the set of all vectors of the form

$$a_1 \cdot b_{(i-1)k+1} + \ldots + a_k \cdot b_{ik}$$

where $a_j \in \{0, \frac{1}{2}\}$ for all j. Our algorithm outputs the union $S = \cup_{i=1}^{n/k} S_i$ (see Figure 1). Observe that $|S| \leq \frac{n}{k} \cdot 2^k$ and that S can be constructed in time $2^k \cdot s^{O(1)}$ where s is the input size.

Remote Set Problem(B)
Input: A lattice basis $B = (b_1, \ldots, b_n) \in \mathbb{Q}^{m \times n}$.
Output: A set S of $\frac{n}{k} \cdot 2^k$ vectors in $\mathrm{span}(B)$ at least one of which is far from $\mathcal{L}(B)$.

- For every $1 \leq i \leq \frac{n}{k}$,
  1. Define $B_i = (b_{(i-1)k+1}, \ldots, b_{ik})$.
  2. Construct the set S_i of all vectors that form a linear combination with all coefficients in $\{0, \frac{1}{2}\}$ of the vectors in B_i.
- Output $S = \cup_{i=1}^{n/k} S_i$.

Figure 1: **An Algorithm for the Remote Set Problem**

Fix some $2 \leq p < \infty$. We claim that there exists a vector in S whose ℓ_p distance from $\mathcal{L}(B)$ is at least $\frac{1}{2 \cdot c_p} \cdot \sqrt{\frac{k}{n}} \cdot \rho^{(p)}(\mathcal{L}(B))$, where $c_p > 0$ is the constant from Lemma 2.5 which depends solely on p. Assume for contradiction that this is not the case. By Lemma 2.1, there exists a vector

$$v = a_1 \cdot b_1 + \ldots + a_n \cdot b_n$$

with $a_j \in \{0, \frac{1}{2}\}$ for all $1 \leq j \leq n$ such that $\mathrm{dist}_p(v, \mathcal{L}(B)) \geq \frac{1}{2} \cdot \rho^{(p)}(\mathcal{L}(B))$. Write

$$v = \frac{1}{2}(v_1 + \ldots + v_{n/k})$$

where for every $1 \leq i \leq \frac{n}{k}$, $v_i = 2 \cdot (a_{(i-1)k+1} \cdot b_{(i-1)k+1} + \ldots + a_{ik} \cdot b_{ik})$. Since $\frac{1}{2} \cdot v_i \in S$ our assumption implies that there exists a lattice vector $u_i \in \mathcal{L}(B)$ such that

$$\|\frac{1}{2} \cdot v_i - u_i\|_p < \frac{1}{2 \cdot c_p} \cdot \sqrt{\frac{k}{n}} \cdot \rho^{(p)}(\mathcal{L}(B)). \tag{2}$$

For every $1 \leq i \leq \frac{n}{k}$, denote $s_i = \frac{1}{2} \cdot v_i - u_i$, and apply Lemma 2.5 to obtain $\varepsilon_1, \ldots, \varepsilon_{n/k} \in \{-1, +1\}$ such that

$$\|\sum_{i=1}^{n/k} \varepsilon_i \cdot s_i\|_p \leq c_p \cdot \sqrt{\frac{n}{k}} \cdot \max_{1 \leq i \leq n/k} \|s_i\|_p < c_p \cdot \sqrt{\frac{n}{k}} \cdot \frac{1}{2 \cdot c_p} \cdot \sqrt{\frac{k}{n}} \cdot \rho^{(p)}(\mathcal{L}(B)) = \frac{1}{2} \cdot \rho^{(p)}(\mathcal{L}(B)),$$

as follows from (2). Finally, observe that the difference between v and $\sum_{i=1}^{n/k} \varepsilon_i \cdot s_i$ is a lattice vector, hence

$$\mathrm{dist}_p(v, \mathcal{L}(B)) = \mathrm{dist}_p(\sum_{i=1}^{n/k} \varepsilon_i \cdot s_i, \mathcal{L}(B)) \leq \|\sum_{i=1}^{n/k} \varepsilon_i \cdot s_i\|_p < \frac{1}{2} \cdot \rho^{(p)}(\mathcal{L}(B)),$$

in contradiction to our choice of v.

The analysis for $p = \infty$ is almost identical to the analysis described above. The only difference is in applying Spencer's theorem (Theorem 2.6) instead of Lemma 2.5 to find a short ± 1 combination of the s_i's. ∎

Notice that in the ℓ_∞ case the remoteness parameter γ obtained in Theorem 3.1 does not depend only on the rank n but also on the dimension m. Hence, let us state it again for the special case of full-rank lattices (i.e., $m = n$) which is usually considered.

Theorem 3.2. *For every $k = k(n) \geq 1$ there exists a deterministic $2^k \cdot s^{O(1)}$ time algorithm for $\mathsf{RSP}^{(\infty)}_{d,\gamma}$ on full-rank lattices with $d(n) = O(\frac{n}{k} \cdot 2^k)$ and $\gamma(n) = O(\sqrt{\frac{n \cdot \log(2k)}{k}})$, where s denotes the input size.*

Now Theorem 1.2 is easily derived from Theorem 3.1 by choosing $k = c \log n$ where n is the lattice rank and c is a constant, as stated in the following corollaries. We note that one can obtain a slightly stronger version of these corollaries by choosing $k = O(\log s)$ where s is the input size.

Corollary 3.3. *For every $2 \leq p < \infty$ and every constant $c \geq 1$, there exists a deterministic polynomial-time algorithm for $\mathsf{RSP}^{(p)}_{d,\gamma}$ with $d(n) = n^{O(c)}$ and $\gamma(n) = O(\sqrt{\frac{n}{c \log n}})$.*

Corollary 3.4. *For every constant $c \geq 1$, there exists a deterministic polynomial-time algorithm for $\mathsf{RSP}^{(\infty)}_{d,\gamma}$ on full-rank lattices with $d(n) = n^{O(c)}$ and $\gamma(n) = O(\sqrt{\frac{n \cdot \log\log n}{c \log n}})$.*

3.2 Extensions of Theorem 1.2

In the analysis of our algorithm for RSP we applied Lemma 2.5 and Theorem 2.6 which roughly speaking say that every set of vectors has a linear combination with ± 1 coefficients of small ℓ_p norm compared to the maximum ℓ_p norm of the vectors in the set. It turns out that similar questions were studied where the goal is to minimize the ℓ_p norm of the linear combination compared to the maximum ℓ_2 norm of the vectors in the set. This is stated in the following theorem which stems from a paper of Banaszczyk [8] (see also [10, Propositions 24, 25]).

Theorem 3.5 ([8]). *For every $2 \leq p \leq \infty$ there exists a constant $c_p > 0$ for which the following holds. For every n vectors $x_1, \ldots, x_n \in \mathbb{R}^m$ there exist $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, +1\}$ such that for $2 \leq p < \infty$*

$$\| \sum_{i=1}^{n} \varepsilon_i \cdot x_i \|_p \leq c_p \cdot n^{1/p} \cdot \max_{1 \leq i \leq n} \|x_i\|_2,$$

and for $p = \infty$,

$$\| \sum_{i=1}^{n} \varepsilon_i \cdot x_i \|_\infty \leq c_\infty \cdot \sqrt{1 + \log n} \cdot \max_{1 \leq i \leq n} \|x_i\|_2.$$

For $p = \infty$, a famous conjecture of Komlós asserts the following.

Conjecture 3.6. *[Komlós Conjecture] There exists a constant $c > 0$ such that for every n vectors $x_1, \ldots, x_n \in \mathbb{R}^m$ there exist $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, +1\}$ such that*

$$\| \sum_{i=1}^{n} \varepsilon_i \cdot x_i \|_\infty \leq c \cdot \max_{1 \leq i \leq n} \|x_i\|_2.$$

Now we observe that Theorem 3.5 can be used to prove an additional property of the output of our algorithm for RSP. The use of Lemma 2.5 and Theorem 2.6 in the proof implied that at least one of points in the output has large ℓ_p distance from the lattice compared to the covering radius in the ℓ_p norm. However, applying Theorem 3.5 in the proof yields that at least one of the vectors has large ℓ_2 distance from the lattice, still compared to the covering radius in the ℓ_p norm. Theorems 3.7 and 3.8 below follow from the algorithm presented in the proof of Theorem 3.1 (see Figure 1) for $k = \Theta(\log n)$ and $k = 1$ respectively. We omit the proof details.

Theorem 3.7. *For every $2 \leq p < \infty$ there exists a constant $c_p > 0$ for which the following holds. For every $c \geq 1$ there exists a deterministic polynomial-time algorithm that given a rank n lattice $\mathcal{L}$ outputs a set of $n^{O(c)}$ points at least one of which has ℓ_2 distance at least $c_p \cdot (\frac{c \log n}{n})^{1/p} \cdot \rho^{(p)}(\mathcal{L})$ from $\mathcal{L}$.*

Theorem 3.8. *There exists a constant $c > 0$ and a deterministic polynomial-time algorithm that given a rank n lattice $\mathcal{L}$ outputs a set of n points at least one of which has ℓ_2 distance at least $\frac{c}{\sqrt{1+\log n}} \cdot \rho^{(\infty)}(\mathcal{L})$ from $\mathcal{L}$. Assuming Conjecture 3.6, one of the points has ℓ_2 distance at least $c \cdot \rho^{(\infty)}(\mathcal{L})$ from $\mathcal{L}$.*

3.3 Proof of Theorem 1.1

The Remote Point Problem is a special case of the Remote Set Problem, in which the goal is to output *one* point which is far from the lattice (i.e., $\mathsf{RSP}_{d,\gamma}$ with $d = 1$). In the following we observe that the LLL algorithm [23] can be used to deterministically and efficiently find a point in space

whose distance from the lattice approximates the covering radius to within a factor exponential in the rank n. We present the result for the ℓ_2 norm, but notice that similar results can be derived for an arbitrary ℓ_p norm by standard relations between ℓ_p norms.

Theorem 3.9. *There exists a deterministic polynomial-time algorithm for* $\mathsf{RSP}_{d,\gamma}$ *with* $d(n) = 1$ *and* $\gamma(n) = 2^{n/2}$.

Proof: It is well-known that given a lattice the LLL algorithm [23] constructs a reduced basis $B = (b_1, \ldots, b_n)$ that generates it and satisfies for every i, $\|\tilde{b}_{i+1}\|^2 \geq \frac{1}{2} \cdot \|\tilde{b}_i\|^2$, where $\tilde{b}_1, \ldots, \tilde{b}_n$ are the Gram-Schmidt orthogonalized vectors. In addition, we recall that Babai's nearest plane algorithm for CVP [6], on input lattice basis B and a target point t, produces a lattice vector whose squared distance from t is at most $\frac{1}{4} \cdot \sum_{i=1}^{n} \|\tilde{b}_i\|^2$. This implies that

$$\rho(\mathcal{L}(B)) \leq \sqrt{\frac{1}{4} \cdot \sum_{i=1}^{n} \|\tilde{b}_i\|^2} \leq \sqrt{\frac{1}{4} \cdot \sum_{i=1}^{n} 2^{n-i} \|\tilde{b}_n\|^2} \leq 2^{n/2-1} \cdot \|\tilde{b}_n\|.$$

For the Remote Point Problem consider the algorithm which given a lattice calculates an LLL-reduced basis $B = (b_1, \ldots, b_n)$ generating it and outputs $\frac{1}{2} \cdot \tilde{b}_n$. Since the projection of every vector in $\mathcal{L}(B)$ to $\mathrm{span}(\tilde{b}_n)$ is $c \cdot \tilde{b}_n$ for some $c \in \mathbb{Z}$ we obtain

$$\mathrm{dist}(\frac{1}{2} \cdot \tilde{b}_n, \mathcal{L}(B)) \geq \frac{1}{2} \cdot \|\tilde{b}_n\| \geq \frac{1}{2^{n/2}} \cdot \rho(\mathcal{L}(B)).$$

■

4 On the Complexity of the Covering Radius Problem

In this section we describe some applications of deterministic and efficient algorithms for RSP to the complexity of the Covering Radius Problem (CRP).

4.1 Proof of Theorem 1.3

The following lemma relates RSP to proving that CRP with certain approximation factors is in NP.

Lemma 4.1. *For every* $1 \leq p \leq \infty$, $d = d(n)$ *and* $\gamma = \gamma(n)$, *if there exists a deterministic polynomial-time algorithm for* $\mathsf{RSP}^{(p)}_{d,\gamma}$ *then* $\mathsf{GapCRP}^{(p)}_{\gamma}$ *is in* NP.

Proof: Let A be a deterministic polynomial-time algorithm for $\mathsf{RSP}^{(p)}_{d,\gamma}$. On input $\mathsf{GapCRP}^{(p)}_{\gamma}$ instance (B, r) we consider the following NP verifier: first run algorithm A on the lattice basis B to obtain d points $t_1, \ldots, t_d \in \mathrm{span}(B)$ such that at least one of them is $\frac{\rho^{(p)}(\mathcal{L}(B))}{\gamma}$-far from $\mathcal{L}(B)$ with respect to ℓ_p distance. Then guess (non-deterministically) d vectors $u_1, \ldots, u_d$, and accept if and only if u_i belongs to $\mathcal{L}(B)$ and satisfies $\mathrm{dist}_p(u_i, t_i) \leq r$ for every $1 \leq i \leq d$. Notice that this can be done in polynomial time, as d is bounded from above by the running time of A.

We turn to prove correctness. If (B, r) is a YES instance of $\mathsf{GapCRP}^{(p)}_{\gamma}$ then $\rho^{(p)}(\mathcal{L}(B)) \leq r$ and therefore there exist $u_1, \ldots, u_d$ for which the verifier accepts. On the other hand, if (B, r) is a NO instance of $\mathsf{GapCRP}^{(p)}_{\gamma}$ then $\rho^{(p)}(\mathcal{L}(B)) > \gamma \cdot r$. Since there exists some $1 \leq j \leq d$ for which $\mathrm{dist}_p(t_j, \mathcal{L}(B)) \geq \frac{\rho^{(p)}(\mathcal{L}(B))}{\gamma} > r$, there is no guess for which the verifier accepts. ■

The following theorems are immediate consequences of Lemma 4.1 and Corollaries 3.3 and 3.4 confirming Theorem 1.3.

Theorem 4.2. *For every $2 \leq p < \infty$ and every constant $c \geq 1$, $\mathsf{GapCRP}_{\gamma}^{(p)}$ is in NP for $\gamma(n) = \sqrt{\frac{n}{c \log n}}$.*

Theorem 4.3. *For every constant $c \geq 1$, $\mathsf{GapCRP}_{\gamma}^{(\infty)}$ on full-rank lattices is in NP for $\gamma(n) = \sqrt{\frac{n \log \log n}{c \log n}}$.*

4.2 Reductions from CRP to CVP

We now show how algorithms for RSP can be used for reducing CRP instances to CVP instances. Lemmas 4.4 and 4.6 deal with Cook and Karp reductions respectively.

Lemma 4.4. *For every $1 \leq p \leq \infty$, $d = d(n)$, $\gamma = \gamma(n)$ and $\gamma' = \gamma'(n)$, if there exists a deterministic polynomial-time algorithm for $\mathsf{RSP}_{d,\gamma}^{(p)}$ then there exists a deterministic rank-preserving polynomial-time Cook reduction from $\mathsf{GapCRP}_{\gamma'}^{(p)}$ to $\mathsf{GapCVP}_{\gamma'/\gamma}^{(p)}$.*

Proof: Let A be a deterministic polynomial-time algorithm for $\mathsf{RSP}_{d,\gamma}^{(p)}$ and fix some γ'. We use A to describe a deterministic rank-preserving polynomial-time Cook reduction from $\mathsf{GapCRP}_{\gamma'}^{(p)}$ to $\mathsf{GapCVP}_{\gamma'/\gamma}^{(p)}$. Let (B, r) be a $\mathsf{GapCRP}_{\gamma'}^{(p)}$ instance, where the lattice $\mathcal{L}(B)$ has rank n. The reduction runs A on B to obtain d points $t_1, \ldots, t_d \in \mathrm{span}(B)$ such that for some $1 \leq j \leq d$ the point t_j satisfies $\mathrm{dist}_p(t_j, \mathcal{L}(B)) \geq \frac{\rho^{(p)}(\mathcal{L}(B))}{\gamma}$. Then, the reduction calls the $\mathsf{GapCVP}_{\gamma'/\gamma}^{(p)}$ oracle on every input (B, t_i, r), $1 \leq i \leq d$, and accepts if and only if the oracle accepts all these inputs.

We now prove the correctness of the reduction. Assume that (B, r) is a YES instance of $\mathsf{GapCRP}_{\gamma'}^{(p)}$, that is $\rho^{(p)}(\mathcal{L}(B)) \leq r$. This implies that for every $1 \leq i \leq d$, $\mathrm{dist}_p(t_i, \mathcal{L}(B)) \leq r$, and hence all the oracle calls are accepted. On the other hand, assume that (B, r) is a NO instance of $\mathsf{GapCRP}_{\gamma'}^{(p)}$, that is $\rho^{(p)}(\mathcal{L}(B)) > \gamma' \cdot r$. This implies that there exists some $1 \leq j \leq d$ for which

$$\mathrm{dist}_p(t_j, \mathcal{L}(B)) \geq \frac{\rho^{(p)}(\mathcal{L}(B))}{\gamma} > \frac{\gamma' \cdot r}{\gamma},$$

and hence at least one of the oracle calls to $\mathsf{GapCVP}_{\gamma'/\gamma}^{(p)}$ is rejected. ∎

The following corollary, easily derived from Lemma 4.4 and Corollaries 3.3 and 3.4, says that there exists a Cook reduction from CRP to CVP with an $O(\sqrt{\frac{n}{\log n}})$ loss in the approximation factor. As was mentioned before, the only other similar result we are aware of is implied by a paper of Micciancio [25] and gives a $\sqrt{n}$ loss in the approximation factor for the Euclidean norm.

Corollary 4.5. *For every $2 \leq p < \infty$, every constant $c \geq 1$ and every $\gamma = \gamma(n)$ there exists a deterministic rank-preserving polynomial-time Cook reduction from $\mathsf{GapCRP}_{\gamma}^{(p)}$ to $\mathsf{GapCVP}_{\gamma \cdot \sqrt{c \log n / n}}^{(p)}$. For $p = \infty$, the reduction is from $\mathsf{GapCRP}_{\gamma}^{(\infty)}$ to $\mathsf{GapCVP}_{\gamma \cdot \sqrt{c \log n / (n \log \log n)}}^{(\infty)}$ restricted to full-rank lattices.*

We turn to prove that algorithms for RSP can be used to obtain Karp reductions from CRP to CVP. Here, in contrast to Lemma 4.4, the parameter d of RSP affects the loss in the approximation factor.

Lemma 4.6. *For every $1 \leq p < \infty$, $d = d(n)$, $\gamma = \gamma(n)$ and $\gamma' = \gamma'(n)$, if there exists a deterministic polynomial-time algorithm for $\mathsf{RSP}^{(p)}_{d,\gamma}$ then there exists a deterministic polynomial-time Karp reduction from $\mathsf{GapCRP}^{(p)}_{\gamma'}$ to $\mathsf{GapCVP}^{(p)}_{\widetilde{\gamma}}$ where $\widetilde{\gamma} = \frac{\gamma'}{\gamma \cdot d^{1/p}}$. For $p = \infty$ the same holds for $\widetilde{\gamma} = \gamma'/\gamma$.*

Proof: Fix some $1 \leq p < \infty$. Let A be a deterministic polynomial-time algorithm for $\mathsf{RSP}^{(p)}_{d,\gamma}$ and fix some γ'. We use A to describe a deterministic polynomial-time Karp reduction from $\mathsf{GapCRP}^{(p)}_{\gamma'}$ to $\mathsf{GapCVP}^{(p)}_{\widetilde{\gamma}}$. Let (B, r) be a $\mathsf{GapCRP}^{(p)}_{\gamma'}$ instance, where the lattice $\mathcal{L}(B)$ has rank n. The reduction runs A on B to obtain a set S of d points $t_1, \ldots, t_d \in \mathrm{span}(B)$ such that for some $1 \leq j \leq d$ the point t_j satisfies $\mathrm{dist}_p(t_j, \mathcal{L}(B)) \geq \frac{\rho^{(p)}(\mathcal{L}(B))}{\gamma}$. The reduction outputs the instance (B', t, r') where $B' = B \otimes I_d$ (i.e., a basis of the direct sum of d copies of $\mathcal{L}(B)$), $t = (t_1, \ldots, t_d)$ is the concatenation of the points of S, and $r' = d^{1/p} \cdot r$. Clearly, the reduction can be implemented in polynomial time.

We turn to prove the correctness of the reduction. Assume that (B, r) is a YES instance of $\mathsf{GapCRP}^{(p)}_{\gamma'}$, that is $\rho^{(p)}(\mathcal{L}(B)) \leq r$. This implies that for every $1 \leq i \leq d$ there exists a lattice vector $u_i \in \mathcal{L}(B)$ such that $\|t_i - u_i\|_p \leq r$. Define $u = (u_1, \ldots, u_d)$ and notice that u belongs to $\mathcal{L}(B')$ and that $\|t - u\|_p \leq d^{1/p} \cdot r = r'$. This implies that (B', t, r') is a YES instance of $\mathsf{GapCVP}^{(p)}_{\widetilde{\gamma}}$. On the other hand, assume that (B, r) is a NO instance of $\mathsf{GapCRP}^{(p)}_{\gamma'}$, that is $\rho^{(p)}(\mathcal{L}(B)) > \gamma' \cdot r$. This implies that there exists some $1 \leq j \leq d$ for which

$$\mathrm{dist}_p(t_j, \mathcal{L}(B)) \geq \frac{\rho^{(p)}(\mathcal{L}(B))}{\gamma} > \frac{\gamma' \cdot r}{\gamma} = \frac{\gamma'}{\gamma \cdot d^{1/p}} \cdot r' = \widetilde{\gamma} \cdot r',$$

and hence $\mathrm{dist}_p(t, \mathcal{L}(B')) > \widetilde{\gamma} \cdot r'$, so (B', t, r') is a NO instance of $\mathsf{GapCVP}^{(p)}_{\widetilde{\gamma}}$.

Finally, it is easy to see that for $p = \infty$ the same reduction reduces $\mathsf{GapCRP}^{(\infty)}_{\gamma'}$ to $\mathsf{GapCVP}^{(\infty)}_{\gamma'/\gamma}$.

■

Acknowledgement

We would like to deeply thank Oded Regev for valuable and fruitful discussions.

References

[1] D. Aharonov and O. Regev. Lattice problems in NP intersect coNP. *Journal of the ACM*, 52(5):749–765, 2005. Preliminary version in FOCS'04.

[2] M. Ajtai. Generating hard instances of lattice problems. In *Complexity of computations and proofs*, volume 13 of *Quad. Mat.*, pages 1–32. Dept. Math., Seconda Univ. Napoli, Caserta, 2004.

[3] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proc. 33th ACM Symp. on Theory of Computing (STOC)*, pages 601–610, 2001.

[4] M. Ajtai, R. Kumar, and D. Sivakumar. Sampling short lattice vectors and the closest lattice vector problem. In *Proc. of 17th IEEE Annual Conference on Computational Complexity (CCC)*, pages 53–57, 2002.

[5] N. Alon, R. Panigrahy, and S. Yekhanin. Deterministic approximation algorithms for the nearest codeword problem. In *APPROX-RANDOM*, volume 5687 of *Lecture Notes in Computer Science*, pages 339–351. Springer, 2009.

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

[7] W. Banaszczyk. Balancing vectors and convex bodies. *Studia Math.*, 106(1):93–100, 1993.

[8] W. Banaszczyk. Balancing vectors and gaussian measures of n-dimensional convex bodies. *Random Struct. Algorithms*, 12(4):351–360, 1998.

[9] N. Bansal. Constructive algorithms for discrepancy minimization. In *FOCS*, pages 3–10, 2010.

[10] F. Barthe, O. Guédon, S. Mendelson, and A. Naor. A probabilistic approach to the geometry of the ℓ_p^n-ball. *The Annals of Probability*, 33(2):pp. 480–513, 2005.

[11] R. Boppana, J. Håstad, and S. Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25:127–132, May 1987.

[12] D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *FOCS*, pages 580–589, 2011.

[13] I. Dinur. Approximating SVP_∞ to within almost-polynomial factors is NP-hard. *Theoretical Computer Science*, 285(1):55–71, 2002.

[14] 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. Preliminary version in FOCS 1998.

[15] O. Goldreich and S. Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. System Sci.*, 60(3):540–563, 2000.

[16] V. Guruswami, D. Micciancio, and O. Regev. The complexity of the covering radius problem on lattices and codes. *Computational Complexity*, 14(2):90–121, 2005. Preliminary version in CCC'04.

[17] I. Haviv, V. Lyubashevsky, and O. Regev. A note on the distribution of the distance from a lattice. *Discrete and Computational Geometry*, 41(1):162–176, 2009.

[18] I. Haviv and O. Regev. Hardness of the covering radius problem on lattices. In *Proc. of 21th IEEE Annual Conference on Computational Complexity (CCC)*, pages 145–158, 2006.

[19] I. Haviv and O. Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 469–477, 2007.

[20] R. Kannan. Minkowski's convex body theorem and integer programming. *Math. Oper. Res.*, 12:415–440, 1987.

[21] S. Khot. Hardness of Approximating the Shortest Vector Problem in Lattices. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 126–135, 2004.

[22] A. Korkine and G. Zolotareff. Sur les formes quadratiques. *Mathematische Annalen*, 6:366–389, 1873.

[23] A. Lenstra, H. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261:515–534, 1982.

[24] D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM Journal on Computing*, 34(1):118–169, 2004. Preliminary version in STOC 2002.

[25] D. Micciancio. Efficient reductions among lattice problems. In *SODA*, pages 84–93, 2008.

[26] D. Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Electronic Colloquium on Computational Complexity (ECCC)*, 19, 2012.

[27] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: A Cryptographic Perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, MA, 2002.

[28] D. Micciancio and O. Regev. Lattice-based cryptography. In D. J. Bernstein and J. Buchmann, editors, *Post-quantum Cryprography*. Springer, 2008.

[29] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 351–358, 2010.

[30] V. D. Milman and G. Schechtman. *Asymptotic theory of finite dimensional normed spaces*. Springer-Verlag New York, Inc., New York, NY, USA, 1986.

[31] C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, 2008. Preliminary version in CCC'07.

[32] O. Regev and R. Rosen. Lattice problems and norm embeddings. In *Proc. 38th ACM Symp. on Theory of Computing (STOC)*, pages 447–456, 2006.

[33] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 53(2-3):201–224, 1987.

[34] J. Spencer. Six standard deviations suffice. *Trans. Amer. Math. Soc.*, 289(2):679–706, 1985.

[35] L. G. Valiant. Graph-theoretic arguments in low-level complexity. In *Mathematical foundations of computer science (Proc. Sixth Sympos., Tatranská Lomnica, 1977)*, pages 162–176. Lecture Notes in Comput. Sci., Vol. 53. Springer, Berlin, 1977.

A The Covering Radius and the nth Successive Minimum in ℓ_p Norms

In this section we prove several inequalities relating the covering radius and the nth successive minimum in ℓ_p norms of rank n lattices. The presented inequalities follow from the papers [7, 8]

of Banaszczyk. Since the analysis of our algorithm for RSP involves similar techniques, we have decided to include them in this appendix.

We start with the following theorem which says that for any ℓ_p norm for $2 \leq p < \infty$ the covering radius of a rank n lattice cannot, up to some multiplicative constant, be higher than $\sqrt{n}$ times the nth successive minimum of the lattice. We note that for the special case of $p = 2$ this follows from Babai's nearest plane algorithm [6] (see, e.g., [27, Theorem 7.9, page 138]).

Theorem A.1. *For every $2 \leq p < \infty$ there exists a constant $c_p > 0$ such that for every rank n lattice $\mathcal{L}$,*

$$\rho^{(p)}(\mathcal{L}) \leq c_p \cdot \sqrt{n} \cdot \lambda_n^{(p)}(\mathcal{L}).$$

Proof: By definition of $\lambda_n^{(p)}$, there exist n linearly independent lattice vectors $b_1, \ldots, b_n \in \mathcal{L}$ such that $\|b_i\|_p \leq \lambda_n^{(p)}(\mathcal{L})$ for all $1 \leq i \leq n$. Consider the sublattice $\mathcal{L}' = \mathcal{L}(b_1, \ldots, b_n) \subseteq \mathcal{L}$. By Lemma 2.1, there exists a vector $v = \sum_{i=1}^n a_i \cdot b_i$ with $a_j \in \{0, \frac{1}{2}\}$ for all $1 \leq j \leq n$ such that $\mathrm{dist}_p(v, \mathcal{L}') \geq \frac{1}{2} \cdot \rho^{(p)}(\mathcal{L}')$. Apply Lemma 2.5 to the vectors $a_1 \cdot b_1, \ldots, a_n \cdot b_n$ to obtain $\varepsilon_1, \ldots, \varepsilon_n \in \{-1, +1\}$ satisfying

$$\|\sum_{i=1}^n \varepsilon_i \cdot a_i \cdot b_i\|_p \leq c_p \cdot \sqrt{n} \cdot \max_{1 \leq i \leq n} \|a_i \cdot b_i\|_p \leq \frac{c_p \cdot \sqrt{n}}{2} \cdot \lambda_n^{(p)}(\mathcal{L}).$$

Since the lattice $\mathcal{L}$ and its sublattice $\mathcal{L}'$ have the same span, we have $\rho^{(p)}(\mathcal{L}) \leq \rho^{(p)}(\mathcal{L}')$. Observe that the difference between v and $\sum_{i=1}^n \varepsilon_i \cdot a_i \cdot b_i$ is a lattice vector of $\mathcal{L}'$, hence

$$\begin{aligned} \rho^{(p)}(\mathcal{L}) &\leq \rho^{(p)}(\mathcal{L}') \leq 2 \cdot \mathrm{dist}_p(v, \mathcal{L}') \\ &= 2 \cdot \mathrm{dist}_p(\sum_{i=1}^n \varepsilon_i \cdot a_i \cdot b_i, \mathcal{L}') \leq 2 \cdot \|\sum_{i=1}^n \varepsilon_i \cdot a_i \cdot b_i\|_p \leq c_p \cdot \sqrt{n} \cdot \lambda_n^{(p)}(\mathcal{L}). \end{aligned}$$

■

Theorem A.1 has an analogous theorem for the ℓ_∞ norm. As in the proof of Theorem 3.1, for ℓ_∞ we apply Spencer's theorem (Theorem 2.6) instead of Lemma 2.5. Since the proof is almost identical to the one of Theorem A.1 we state it below and omit the proof details.

Theorem A.2. *There exists a constant $c > 0$ such that for every m-dimensional lattice $\mathcal{L}$ of rank n,*

$$\rho^{(\infty)}(\mathcal{L}) \leq c \cdot \sqrt{n \cdot \log(2m/n)} \cdot \lambda_n^{(\infty)}(\mathcal{L}).$$

Finally, as in Section 3.2, Theorem 3.5 can be used to obtain the following statement, which relates the covering radius of a lattice in the ℓ_p norm to its nth successive minimum in the ℓ_2 norm.

Theorem A.3. *For every $2 \leq p < \infty$ there exists a constant $c_p > 0$ such that for every rank n lattice $\mathcal{L}$,*

$$\rho^{(p)}(\mathcal{L}) \leq c_p \cdot n^{1/p} \cdot \lambda_n^{(2)}(\mathcal{L}).$$

For $p = \infty$,

$$\rho^{(\infty)}(\mathcal{L}) \leq c_\infty \cdot \sqrt{1 + \log n} \cdot \lambda_n^{(2)}(\mathcal{L}),$$

for some constant $c_\infty > 0$. Assuming Conjecture 3.6,

$$\rho^{(\infty)}(\mathcal{L}) \leq c_\infty \cdot \lambda_n^{(2)}(\mathcal{L}).$$