Lecture 7
Limits on inapproximability

Lecturer: Oded Regev
Scribe: Michael Khanevsky

Let us recall the promise problem GapCVP_γ.

DEFINITION 1 GapCVP_γ
YES *instances:* *triples* (B, v, d) *such that* $\operatorname{dist}(v, \mathcal{L}(B)) \le d$
NO *instances:* *triples* (B, v, d) *such that* $\operatorname{dist}(v, \mathcal{L}(B)) > \gamma d$
where B *is a basis for a lattice in* $\mathbb{Q}^n$*,* $v \in \mathbb{Q}^n$ *is a vector, and* $d \in \mathbb{Q}$ *is some number.*

It is easy to see that $\mathsf{GapCVP}_\gamma \in$ **NP** for any $\gamma \ge 1$. Indeed, a witness is a vector $u \in \mathcal{L}(B)$ such that $\|v - u\| \le d$. This witness is of polynomial size (since the length of the vector is at most $\|v\| + d$) and it can be verified efficiently (simply check that $\|v - u\| \le d$).

What about the complement of GapCVP_γ? In other words, for what values of γ is $\mathsf{GapCVP}_\gamma \in$ **coNP**? In order to show such a containment, we need to give a witness that $\operatorname{dist}(v, \mathcal{L}(B)) > \gamma d$. Some thought reveals that this is no longer a trivial task. After all, there is a huge number of lattice vectors that can potentially be very close to v. Some of the early results in this direction [5, 2] showed that $\mathsf{GapCVP}_n \in$ **coNP**. These results are based on the use of dual lattices. Later, Goldreich and Goldwasser [4] showed that $\mathsf{GapCVP}_{\sqrt{n/\log n}} \in$ **coAM** (we define the class **AM** later). More recently, Aharonov and Regev [1] showed that $\mathsf{GapCVP}_{\sqrt{n}} \in$ **coNP**. All of these results also hold for GapSVP since GapSVP is not harder than GapCVP (this was shown rigorously in the previous class).

To summarize, we have that $\mathsf{GapCVP}_{\sqrt{n/\log n}} \in$ **NP** $\cap$ **coAM** and $\mathsf{GapCVP}_{\sqrt{n}} \in$ **NP** $\cap$ **coNP**. One of the interesting implications of these results is the following. It is known that GapCVP_γ is **NP**-hard for $\gamma \le n^{O(1/\log\log n)}$ [3]. Can we hope to improve this **NP**-hardness result to, say, $\gamma = \sqrt{n}$? The above results imply that the answer is probably *no*: if GapCVP_γ is **NP**-hard for $\gamma \ge \sqrt{n/\log n}$ (even under Cook reductions) then the polynomial hierarchy collapses. The proof requires some care (especially for Cook reductions) and is discussed in Section 2.

Finally, another interest in the above results arises from lattice-based cryptographic constructions. All known constructions are based on the worst-case hardness of lattice problems such as GapSVP_{n^c} for some constant $c \ge 1$. Hence, the above results imply that these constructions are based on a problem in **NP**$\cap$**coNP** (like factoring).

1 The Goldreich-Goldwasser protocol

In this section we focus on the Goldreich-Goldwasser protocol.

THEOREM 1 $\mathsf{GapCVP}_{\sqrt{n/\log n}} \in$ **coAM**

REMARK 1 In fact, the proof implies a stronger result, namely, that $\mathsf{GapCVP}_{\sqrt{n/\log n}}$ is contained in a complexity class known as Statistical Zero Knowledge, or **SZK**.

For simplicity, we will show that $\mathsf{GapCVP}_{\sqrt{n}} \in$ **coAM**. A slightly more careful analysis of the same protocol yields a gap of $c\sqrt{n/\log n}$ for any constant $c > 0$. First, let us define the class **AM**.

DEFINITION 2 *A promise problem is in* **AM** *if there exists a protocol with a constant number of rounds between a* **BPP** *machine Arthur and a computationally unbounded Merlin, and two constants* $0 \le a < b \le 1$ *such that*

- *for any* YES *input, there exists a strategy for Merlin such that Arthur accepts with probability at least* b*, and*
- *for any* NO *input, and any strategy for Merlin, Arthur accepts with probability at most* a*.*

In order to prove Theorem 1, we present a protocol that allows Arthur to verify that a point is far from the lattice. Specifically, given (B, v, d), Arthur accepts with probability 1 if $\operatorname{dist}(v, \mathcal{L}(B)) \geq \sqrt{n}d$ and rejects with some positive probability if $\operatorname{dist}(v, \mathcal{L}(B)) < d$.

Informally, the protocol is as follows. Arthur first flips a fair coin. If it comes up heads, he randomly chooses a 'uniform' point in the lattice $\mathcal{L}(B)$; if it comes up tails, he randomly chooses a 'uniform' point in the shifted lattice $v + \mathcal{L}(B)$. Let w denote the resulting point. Arthur randomly chooses a uniform point x from the ball of radius $\frac{1}{2}\sqrt{n}d$ around w and then sends x to Merlin. Merlin is supposed to tell Arthur if the coin came up heads or not.

The correctness of this protocol follows from the following two observations (see Figure 1). If $\operatorname{dist}(v, \mathcal{L}(B)) \geq \sqrt{n}d$ then the two distributions are disjoint and the Merlin can answer correctly with probability 1. On the other hand, if $\operatorname{dist}(v, \mathcal{L}(B)) < d$, then the overlay between the two distributions is too large and the prover must make a mistake with some positive probability.

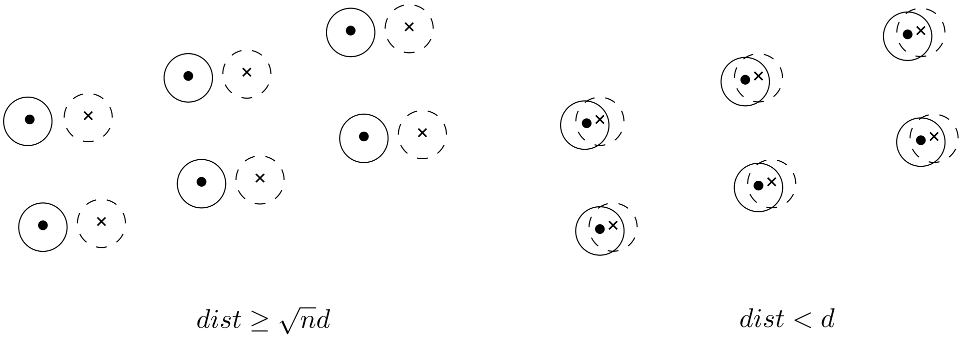

Figure 1: Two distributions

This informal description hides two technical problems. First, we cannot really work with the point x since it is chosen from a continuous distribution (and hence cannot be represented precisely in any finite number of bits). This is easy to take care of by working with an approximation of x with some polynomial number of bits. Another technical issue is the choice of a 'random' point from $\mathcal{L}(B)$. This is an infinite set and there is no uniform distribution on it. On possible solution is to take the uniform distribution on points in the intersection of $\mathcal{L}(B)$ with, say, some very large hypercube. This indeed solves the problem, but introduces some unnecessary complications to the proof since one needs to argue that the probability to fall close to the boundary of the hypercube is low. The solution we choose here is different and avoids this problem altogether by working with distribution on the basic parallelepiped of the lattice. We describe this solution in Subsection 1.3.

In the next few subsections, we present the necessary preliminaries for the proof.

1.1 Statistical Distance

DEFINITION 3 *The statistical distance between two distributions X, Y on some set Ω is defined as* $\Delta(X, Y) = \max_{A \subseteq \Omega} |\Pr(X \in A) - \Pr(Y \in A)|$.

One useful special case of this definition is the case where X and Y are discrete distributions over some countable set Ω. In this case, we have

$$\Delta(X, Y) = \frac{1}{2} \sum_{\omega \in \Omega} |\Pr(X = \omega) - \Pr(Y = \omega)|.$$

Another useful special case is when X and Y are distributions on $\mathbb{R}^n$ with density functions f, g. In this case, we have

$$\Delta(X,Y) = \frac{1}{2}\int_{\mathbb{R}^n} |f(x) - g(x)|\,\mathrm{d}x.$$

For any distributions X, Y, $\Delta(X,Y)$ obtains values between 0 and 1. It is 1 if and only if the supports of X and Y are disjoint[1]; it is 0 if and only if X and Y are equivalent distributions. It is helpful to consider the following interpretation of statistical distance. Assume we are given a sample that is taken from X with probability $\frac{1}{2}$ or from Y with probability $\frac{1}{2}$. Our goal is to decide which distribution the sample comes from. Then, it can be seen that our best strategy succeeds with probability $\frac{1}{2} + \frac{1}{2}\Delta(X,Y)$.

One important fact concerning the statistical distance is that cannot increase by the application of a possibly randomized function. In symbols, $\Delta(f(X), f(Y)) \le \Delta(X,Y)$ for any (possibly randomized) function f. This fact is easy to deduce from the above interpretation of Δ.

1.2 Balls in n-dimensional Space

FACT 1 *The volume of the unit ball in n dimensions is*

$$V_n := \frac{\pi^{n/2}}{(n/2)!}$$

where we define $n! = n(n-1)!$ for $n \ge 1$ and $\frac{1}{2}! = \frac{1}{2}\sqrt{\pi}$.

It can be shown that

$$\frac{(n+\frac{1}{2})!}{n!} \approx \frac{n!}{(n-\frac{1}{2})!} \approx \sqrt{n}.$$

LEMMA 2 *For any $\varepsilon > 0$ and any vector v of length $\|v\| \le \varepsilon$, the relative volume of the intersection of two unit balls whose centers are separated by v satisfies*

$$\frac{\mathrm{vol}(\mathbf{B}(0,1) \cap \mathbf{B}(v,1))}{\mathrm{vol}(\mathbf{B}(0,1))} \ge \varepsilon \frac{(1-\varepsilon^2)^{\frac{n-1}{2}}}{3}\sqrt{n}$$

PROOF: As shown in Figure 2, the above intersection contains a cylinder of height ε and radius $\sqrt{1-\varepsilon^2}$ centered around $v/2$. Hence, the volume of the intersection satisfies:

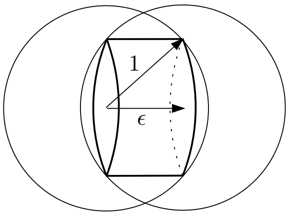

Figure 2: A cylinder in the intersection of two balls

$$\frac{\mathrm{vol}(\mathbf{B}(0,1) \cap \mathbf{B}(v,1))}{\mathrm{vol}(\mathbf{B}(0,1))} > \frac{\varepsilon V_{n-1}(\sqrt{1-\varepsilon^2})^{n-1}}{V_n} = \varepsilon(1-\varepsilon^2)^{\frac{n-1}{2}} \frac{\pi^{\frac{n-1}{2}}/(\frac{n-1}{2})!}{\pi^{\frac{n}{2}}/(\frac{n}{2})!} \approx \varepsilon(1-\varepsilon^2)^{\frac{n-1}{2}} \frac{\sqrt{n/2}}{\sqrt{\pi}}$$

[1] More precisely, in the continuous case, the intersection needs to be a set of measure zero.

□

Notice that for $\varepsilon \leq \frac{2}{\sqrt{n}}$, the right hand side of the expression in Lemma 2 is at least some positive constant independent of n. This yields the following corollary.

COROLLARY 3 *There exists a constant $\delta > 0$ such that for any $d > 0$ and any $y \in \mathbb{R}^n$ such that $\|y\| \leq d$,*

$$\Delta\big(U(\mathbf{B}(0, \tfrac{1}{2}\sqrt{n}d)),\ U(\mathbf{B}(y, \tfrac{1}{2}\sqrt{n}d))\big) < 1 - \delta,$$

where $U(\cdot)$ denotes the uniform distribution on a set.

PROOF: This statistical distance is exactly the volume of the symmetric difference of two balls divided by the sum of their volumes. According to the above lemma, this is bounded away from 1. □

REMARK 2 For any constant c and any $\varepsilon \leq c\sqrt{\log n}$, the right hand side of the expression in Lemma 2 is still greater than some $1/\text{poly}(n)$. Using this, one can obtain the improved result $\mathsf{GapCVP}_{c\sqrt{n/\log n}} \in$ **coAM**.

1.3 Working with periodic distributions

In the informal description above, we talked about the 'uniform distribution' on the lattice. This is clearly not defined. One possible solution is to restrict our attention to some large enough cube $[-K, K]^n$. While possible, this solution introduces some technical annoyances as one has to argue that the probability to fall too close to the boundary of the cube (where the protocol might behave badly) is small.

Instead, our solution will be to work with only one period of the distribution. To demonstrate this approach, let us first consider the one-dimensional case. Assume we want to represent the distribution intuitively described as follows: choose a random point from $2\pi\mathbb{Z}$ and add to it a number chosen uniformly from $[-0.1, 0.1]$. The first solution above would require us to take some large segment, say, $[-1000, 1000]$, and to restrict our distribution on it. Instead, we take one period of the distribution, say the segment $[0, 2\pi]$, and consider the distribution on it. Hence, we obtain the uniform distribution on $[0, 0.1] \cup [2\pi - 0.1, 2\pi]$. Notice that we could take another period, say the segment $[-2\pi, 0]$, and work with it instead. Crucially, the transformation from one representation to another can be performed efficiently (in this case, by subtracting or adding 2π).

A similar idea works for higher dimensions. If we want to represent a periodic distribution on a lattice, we consider it as a distribution on some period, say, $\mathcal{P}(B)$ for some basis B. As before, we have several possible representation, depending on the choice of basis B. The transformation from a representation using B_1 to one using B_2 can be done efficiently by reducing points modulo $\mathcal{P}(B_2)$. Mathematically speaking, the objects we work with are distributions on the quotient $\mathbb{R}^n/\mathcal{L}(B)$, and $\mathcal{P}(B)$ is its set of representatives.

We emphasize that it is much easier to imagine 'periodic distributions' on $\mathbb{R}^n$. However, technically, it is much easier to work with distributions on $\mathcal{P}(B)$.

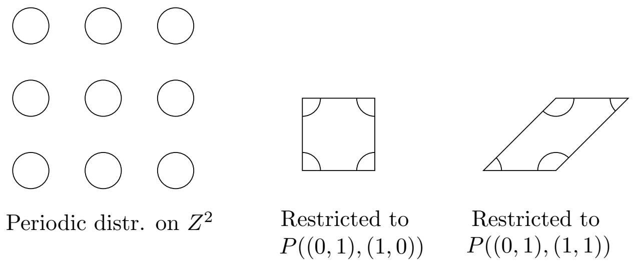
Periodic distr. on Z^2
Restricted to $P((0,1),(1,0))$
Restricted to $P((0,1),(1,1))$

1.4 The protocol

We now prove Theorem 1. First, recall the following definition.

DEFINITION 4 *For $x \in \mathbb{R}^n$, $x \bmod \mathcal{P}(B)$ is the unique $y \in \mathcal{P}(B)$ such that $x - y \in \mathcal{L}(B)$.*

The **AM** protocol:

1. Arthur selects $\sigma \in \{0,1\}$ uniformly and a random point t in the ball $B(0, \frac{1}{2}\sqrt{n}d)$. He then sends $x = (\sigma v + t) \bmod \mathcal{P}(B)$ to Merlin.
2. Merlin checks if $\operatorname{dist}(x, \mathcal{L}(B)) < \operatorname{dist}(x, v + \mathcal{L}(B))$. If so, he responds with $\tau = 0$; otherwise, he responds with $\tau = 1$.
3. Arthur accepts if and only if $\tau = \sigma$.

REMARK 3 For simplicity, we ignore issues of finite precision; these can be dealt with by standard techniques. One issue that we do want to address is how to choose a point from the ball $B(0,R)$ uniformly at random. One option is to use known algorithms for sampling (almost) uniformly from arbitrary convex bodies, and apply them to the case of a ball. A simpler solution is the following. Take n independent samples $v_1, \ldots, v_n \in \mathbb{R}$ from the standard normal distribution and let v be the vector $(v_1, \ldots, v_n) \in \mathbb{R}^n$. Then v is distributed according to the standard n-dimensional Gaussian distribution, which is rotationally invariant. Now, choose r from the distribution on $[0,R]$ whose probability density function is proportional to r^{n-1} (this corresponds to the $n-1$-dimensional surface area of a sphere of radius r). The vector $\frac{r}{\|v\|}v$ is distributed uniformly in $B(0,R)$.

CLAIM 4 (COMPLETENESS) *If $\operatorname{dist}(v, \mathcal{L}(B)) > \sqrt{n}d$ then Arthur accepts with probability 1.*

PROOF: Assume $\sigma = 0$. Then

$$\operatorname{dist}(x, \mathcal{L}(B)) = \operatorname{dist}(t, \mathcal{L}(B)) \le \|t\| \le \frac{1}{2}\sqrt{n}d.$$

On the other hand,

$$\operatorname{dist}(x, v + \mathcal{L}(B)) = \operatorname{dist}(t, v + \mathcal{L}(B)) = \operatorname{dist}(t - v, \mathcal{L}(B)) \ge \operatorname{dist}(v, \mathcal{L}(B)) - \|t\| > \frac{1}{2}\sqrt{n}d.$$

Hence, Merlin answers correctly and Arthur accepts. The case $\sigma = 1$ is similar. □

CLAIM 5 (SOUNDNESS) *If $\operatorname{dist}(v, \mathcal{L}(B)) \le d$ then Arthur rejects with some constant probability.*

PROOF: Let y be the difference between v and its closest lattice point. So y is such that $v - y \in \mathcal{L}(B)$ and $\|y\| \le d$. Let η_0 be the uniform distribution on $\mathbf{B}(0, \frac{1}{2}\sqrt{n}d)$ and let η_1 be the uniform distribution on $\mathbf{B}(y, \frac{1}{2}\sqrt{n}d)$. Notice that the point Arthur sends can be equivalently seen as a point chosen from η_σ reduced modulo $\mathcal{P}(B)$. Accordingly to Corollary 3, $\Delta(\eta_0, \eta_1)$ is smaller than $1 - \delta$. Hence,

$$\Delta(\eta_0 \bmod \mathcal{P}(B), \eta_1 \bmod \mathcal{P}(B)) \le \Delta(\eta_0, \eta_1) < 1 - \delta$$

and Arthur rejects with probability at least δ. □

2 NP-hardness

In this section we show that, based on the following theorem of [1], $\mathsf{GapCVP}_{\sqrt{n}}$ is unlikely to be **NP**-hard, even under Cook reductions.

THEOREM 6 $\mathsf{GapCVP}_{\sqrt{n}} \in$ **NP** $\cap$ **coNP**

A similar proof shows that, based on Theorem 1, $\mathsf{GapCVP}_{\sqrt{n/\log n}}$ is unlikely to be **NP**-hard. However, for simplicity, we show this only for a $\sqrt{n}$ gap.

First, let us consider the simpler case of Karp reductions. If a problem in **coNP** is **NP**-hard under a Karp reduction (i.e., there is a many-to-one reduction from SAT to our problem) then the following easy claim shows that **NP** $\subseteq$ **coNP** (and hence the polynomial hierarchy collapses).

CLAIM 7 *If a promise problem* $\Pi = (\Pi_{\mathrm{YES}}, \Pi_{\mathrm{NO}})$ *is in* **coNP** *and is* **NP***-hard under Karp reductions, then* **NP** $\subseteq$ **coNP***.*

PROOF: Take any language L in **NP**. By assumption, there exists an efficient procedure R that maps any $x \in L$ to $R(x) \in \Pi_{\mathrm{YES}}$ and any $x \notin L$ to $R(x) \in \Pi_{\mathrm{NO}}$. Since $\Pi \in$ **coNP**, we have an **NP** verifier V such that for any $y \in \Pi_{\mathrm{NO}}$ there exists a w such that $V(y, w)$ accepts, and for any $y \in \Pi_{\mathrm{YES}}$ and any w, $V(y, w)$ rejects. Consider the verifier $U(x, \omega)$ given by $V(R(x), \omega)$. Notice that for all $x \notin L$ there exists a w such that $U(x, w)$ accepts and moreover, for all $x \in L$ and all w $U(x, w)$ rejects. Hence, $L \in$ **coNP**. □

The case of Cook reductions requires some more care. For starters, there is nothing special about a problem in **coNP** that is NP-hard under Cook reductions (for example, coSAT is such a problem). Instead, we would like to show that if a problem in **NP** $\cap$ **coNP** is **NP**-hard under Cook reductions, the polynomial hierarchy collapses. This implication is not too difficult to show for *total* problems (i.e., languages). However, we are dealing with *promise* problems and for such problems this implication is not known to hold (although still quite believable). In a nutshell, the difficulty arises because a Cook reduction might perform queries that are neither a YES instance nor a NO instance and for such queries we have no witness.

This issue can be resolved by using the fact that not only $\mathsf{GapCVP}_{\sqrt{n}} \in$ **NP** but also $\mathsf{CVP} \in$ **NP**. In other words, no promise is needed in order to show that a point is close to the lattice. In the following, we show that any problem with the above properties is unlikely to be NP-hard.

LEMMA 2 *Let* $\Pi = (\Pi_{\mathrm{YES}}, \Pi_{\mathrm{NO}})$ *be a promise problem and let* Π_{MAYBE} *denote all instances outside* $\Pi_{\mathrm{YES}} \cup \Pi_{\mathrm{NO}}$*. Assume that* Π *is in* **coNP** *and that the (non-promise) problem* $\Pi' = (\Pi_{\mathrm{YES}} \cup \Pi_{\mathrm{MAYBE}}, \Pi_{\mathrm{NO}})$ *is in* **NP***. Then, if* Π *is NP-hard under Cook reductions then* **NP** $\subseteq$ **coNP** *and the polynomial hierarchy collapses.*

PROOF: Take any language L in **NP**. By assumption, there exists a Cook reduction from L to Π. That is, there exists a polynomial time procedure T that solves L given access to an oracle for Π. The oracle answers YES on queries in Π_{YES} and NO on queries in Π_{NO}. Notice, however, that its answers on queries from Π_{MAYBE} are arbitrary and should not affect the output of T.

Since $\Pi \in$ **coNP**, there exists a verifier V_1 and a witness $w_1(x)$ for every $x \in \Pi_{\mathrm{NO}}$ such that V_1 accepts $(x, w_1(x))$. Moreover, V_1 rejects (x, w) for any $x \in \Pi_{\mathrm{YES}}$ and any w. Similarly, since $\Pi' \in$ **NP**, there exists a verifier V_2 and a witness $w_2(x)$ for every $x \in \Pi_{\mathrm{YES}} \cup \Pi_{\mathrm{MAYBE}}$ such that V_2 accepts $(x, w_2(x))$. Moreover, V_2 rejects (x, w) for any $x \in \Pi_{\mathrm{NO}}$ and any w.

We now show that L is in **coNP** by constructing an **NP** verifier. Let Φ be an input to L and let $x_1, \ldots, x_k$ be the set of oracle queries which T performs on input Φ. Our witness consists of k pairs, one for each x_i. For $x_i \in \Pi_{\mathrm{NO}}$ we include the pair $(\mathrm{NO}, w_1(x_i))$ and for $x_i \in \Pi_{\mathrm{YES}} \cup \Pi_{\mathrm{MAYBE}}$ we include the pair $(\mathrm{YES}, w_2(x_i))$. The verifier simulates T; for each query x_i that T performs, the verifier reads the pair corresponding to x_i in the witness. If the pair is of the form (YES, w) then the verifier checks that $V_2(x_i, w)$ accepts and then returns YES to T. Similarly, if the pair is of the form (NO, w) then the verifier checks that $V_1(x_i, w)$ accepts and then returns NO to T. If any of the calls to V_1 or V_2 rejects, then the verifier rejects. Finally, if T decides that $\Phi \in L$, the verifier rejects and otherwise it accepts.

The completeness follows easily. More specifically, if $\Phi \notin L$ then the witness described above will cause the verifier to accept. In order to prove soundness, assume that $\Phi \in L$ and let us show that the verifier rejects. Notice that for each query $x_i \in \Pi_{\text{NO}}$ the witness must include a pair of the form (NO, w) because otherwise V_2 would reject. Similarly, for each query $x_i \in \Pi_{\text{YES}}$ the witness must include a pair of the form (YES, w) because otherwise V_1 would reject. This implies that T receives the correct answers for all of its queries inside $\Pi_{\text{NO}} \cup \Pi_{\text{YES}}$ and must therefore output the correct answer, i.e., that $\Phi \in L$ and then the verifier rejects. □

References

[1] D. Aharonov and O. Regev. Lattice problems in NP intersect coNP. In *Proc. 45th Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 362–371, 2004.

[2] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.

[3] 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.

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

[5] J. C. Lagarias, H. W. Lenstra, Jr., and C.-P. Schnorr. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990.