



SVP_p is Deterministically NP-Hard for all $p > 2$, Even to Approximate Within a Factor of $2^{\log^{1-\varepsilon} n}$

Isaac M Hair*
UCSB, UCLA

Amit Sahai†
UCLA

Abstract

We prove that SVP_p is NP-hard to approximate within a factor of $2^{\log^{1-\varepsilon} n}$, for all constants $\varepsilon > 0$ and $p > 2$, under standard deterministic Karp reductions. This result is also the first proof that *exact* SVP_p is NP-hard in a finite ℓ_p norm. Hardness for SVP_p with p finite was previously only known if $NP \not\subseteq RP$, and under that assumption, hardness of approximation was only known for all constant factors. As a corollary to our main theorem, we show that under the Sliding Scale Conjecture, SVP_p is NP-hard to approximate within a small polynomial factor, for all constants $p > 2$.

Our proof techniques are surprisingly elementary; we reduce from a *regularized* PCP instance directly to the shortest vector problem by using simple gadgets related to Vandermonde matrices and Hadamard matrices.

*isaacmhair@gmail.com

†sahai@cs.ucla.edu

1 Introduction

A lattice \mathcal{L} is the set of all integral linear combinations of a set of basis vectors $\{\mathbf{b}_1, \dots, \mathbf{b}_n\}$. Lattices are a classical mathematical object, with research on the subject dating back centuries. In the last several decades, lattices have come to the forefront of study in theoretical computer science, due to their applications in algorithmic number theory [LLL82], integer programming [LJ83, Kan87, FT87, Sch98], coding theory [For88, dB02, Zam14], and perhaps most strikingly, post-quantum cryptography [Ajt98, NS01, MR07, Reg09, P⁺16]. Most cryptographic primitives based on lattice problems enjoy a special feature: while the problems used in actual cryptosystems are sampled from random distributions, the computational hardness of breaking these cryptosystems is implied by the difficulty of solving *worst-case* problems. Perhaps the most foundational of these worst-case problems is the gap shortest vector (GapSVP) problem. As the name suggests, this problem amounts to approximating the length of the shortest nonzero vector in a given lattice \mathcal{L} , where we measure the length of a vector $\mathbf{w} \in \mathcal{L}$ using the ℓ_p norm:¹

Problem 1 (γ -GapSVP _{p}). *Given a matrix $\mathbf{M} \in \mathbb{Z}^{n \times m}$ and a threshold k , distinguish between the following two cases, when one of them is promised to hold:*

1. *There exists a nonzero vector $\mathbf{v} \in \mathbb{Z}^n$ such that $\|\mathbf{v}\mathbf{M}\|_p \leq k$.*
2. *For all nonzero vectors $\mathbf{v} \in \mathbb{Z}^n$, we have $\|\mathbf{v}\mathbf{M}\|_p > \gamma k$.*

(Here and throughout the paper, all vectors are row vectors.)

There is a rich body of work exploring upper bounds [Kan87, AKS02, BN09, EV22] and lower bounds [Mic01, HR07, ASD18, BGLR24] for all choices of p , as well as the relationship between GapSVP instances in different ℓ_p norms [RR06, ACK⁺21]. Besides serving as a natural counterpart to the results for $p = 2$, GapSVP results for $p \neq 2$ have proved important historically because they often lead to results for the $p = 2$ case. This is perhaps most obvious when examining the chronology of GapSVP lower bounds.

The shortest vector problem (without an approximation gap) was first proven to be NP hard in the ℓ_∞ norm by van Emde Boas [vEB81], who conjectured that the same hardness result should hold for ℓ_2 . After nearly two decades, Ajtai [Ajt98] gave a breakthrough reduction showing that the exact shortest vector problem is hard in the ℓ_2 norm unless $\text{NP} \subseteq \text{RP}$, i.e. with the caveat that his reduction is *randomized*. Over the next few years, researchers improved upon Ajtai’s results [CN98, Mic01, Kho03, Kho05, HR07]. Many of the improvements were first discovered for ℓ_p with p being a sufficiently large constant, and then were generalized to all $p \geq 1$. The current state of affairs with regard to the assumption that $\text{NP} \not\subseteq \text{RP}$ is summarized in the following theorem.

Theorem 1.1 ([Kho05]). *For all $p \geq 1$, there is no randomized polynomial time algorithm for γ -GapSVP _{p} , for $\gamma = O(1)$, unless $\text{NP} \subseteq \text{RP}$.*

Hardness of approximation for larger (but still subpolynomial) factors are known only under significantly complexity stronger assumptions like $\text{NP} \not\subseteq \text{RTIME}(2^{(\log n)^{O(1)}})$ or $\text{NP} \not\subseteq \bigcap_{\delta > 0} \text{RTIME}(2^{n^\delta})$ [Kho05, RR06, HR07]. All known reductions are randomized, and the best approximation factor we can achieve under $\text{NP} \subseteq \text{RP}$ is still just $\gamma = O(1)$. There is a single exception: Dinur showed that γ -GapSVP _{∞} is NP hard under a deterministic reduction, where $\gamma = n^{c/\log \log n}$ for a sufficiently small constant $c > 0$ [Din02].

Thus, a central 40-year-old open question is whether we can get a deterministic NP-hardness reduction when p is finite, which was explicitly asked by many authors including van Emde Boas [vEB81], Micciancio [Mic01, Mic12], Haviv and Regev [HR07], and Bennett and Peikert [BP22, Ben23].

¹Recall that for finite p , we define the ℓ_p norm of a vector $\mathbf{w} \in \mathbb{Z}^m$ as $\|\mathbf{w}\|_p := (\sum_{i=1}^m |\mathbf{w}_i|^p)^{1/p}$, and for ℓ_∞ we define $\|\mathbf{w}\|_\infty := \max_i |\mathbf{w}_i|$. We use the ℓ_0 pseudo-norm to denote the number of nonzero entries in a given vector.

Question 1. ([BP22]) “Indeed, it is a notorious, long-standing open problem to prove that GapSVP_p is NP-hard, **even in its exact form**,² under a deterministic reduction for some finite p .”

Another natural question is whether we can improve the NP hardness of approximation factor, even under randomized reductions.

Question 2. (*Implicit in the complexity community*) Can we find a polynomial time (potentially randomized) reduction from an NP complete problem to $\omega(1)$ - GapSVP_p , for some finite p ?

One explanation for the lack of improvement since Haviv and Regev’s work is that essentially all known hardness of approximation results for GapSVP_p (with the exception of $p = \infty$) boil down to a similar technique: They first achieve a constant gap via a (randomized) polynomial time reduction from the closest vector problem, which is already known to be NP-hard to approximate [ABSS97], by using a special mathematical object called a *locally dense lattice*. Then they apply tensoring to boost the approximation factor past a constant. Due to the critical usage of tensoring, this approach appears to be incapable of yielding hardness results for $\gamma = \omega(1)$ under the sole assumption that $\text{NP} \subseteq \text{RP}$. (Each tensoring operation increases the lattice dimension by a multiplicative factor of $\text{poly}(n)$, and we need a superconstant number of such operations to reach $\gamma = \omega(1)$.) Derandomization also appears to be quite difficult; despite decades of research, we do not have a deterministic construction of locally dense lattices.

A New More Direct Approach. In this paper, we resolve both questions above in the affirmative. We present a *deterministic, nearly-polynomial factor NP hardness of approximation result for γ - GapSVP_p for all constants $p > 2$* , based on a direct reduction from a carefully pre-processed probabilistically checkable proof (PCP) for NP. Our reduction bypasses the closest vector problem entirely; its closest relative appears to be Dinur’s 2002 NP hardness of approximation result for γ - GapSVP_∞ . Before presenting our results more formally, we give some background on PCPs.

A PCP consists of a *verifier* which takes as input a sequence of symbols over alphabet Σ (the *proof*) and behaves as follows. The verifier uses r random bits to select q locations in the proof to read, and accepts iff the symbols at those locations satisfy some constraint (which may or may not be different for each choice of random bits). Finally, a (perfect correctness) PCP theorem provides a deterministic polynomial time reduction from instances of the NP complete problem circuit SAT to the *description* of a polynomial-time verifier, so that:

1. (Completeness) If the SAT instance was satisfiable, then there exists a proof which causes the verifier to accept with probability 1.
2. (Soundness) If the SAT instance was unsatisfiable, then every proof will cause the verifier to accept with probability at most s (the *soundness parameter*).

The strongest constant-query PCP we currently know how to construct is due to Dinur, Fischer, Kindler, Raz, and Safra:

Theorem 1.2 ([DFK⁺99]). *For every constant $\varepsilon > 0$, there exists a PCP for SAT instances of size n that uses $r = O(\log n)$ random bits to make $q = O(1)$ queries to a proof over alphabet Σ , where the alphabet size is $|\Sigma| = 2^{\Theta(\log^{1-\varepsilon} n)}$ and the soundness parameter is $s = 2^{-\Theta(\log^{1-\varepsilon} n)}$. The PCP can be constructed deterministically in $n^{O(1)}$ time.*

For our result, we will not leverage the PCP theorem exactly as stated above. Instead we first pre-process the PCP using a regularization technique due to Hirahara and Moshkovitz [HM23], which ensures that every position in the proof is read with equal probability. Details on this technique are given in Section 2.

²Emphasis added.

At a very high level, our reduction will make use of simple, deterministically constructable gadgets related to Vandermonde matrices and Hadamard matrices.³ Our analysis involves a novel approach for examining how these gadgets interact specifically within the context of a regularized PCP. The construction is not very technical; rather, it requires a paradigm shift in how we think about reductions to lattice problems.

1.1 Our Results

Below we give a basic version of our main theorem; for the full statement, see Theorem 4.3.

Theorem 1.3 (Informal). *Suppose there exists a PCP for SAT instances of size n that uses $r = O(\log n)$ random bits to make $q = O(1)$ queries to a proof over alphabet Σ and has soundness s , where $1/s \geq r^{\omega(1)}$. Suppose further that every proof location is read with the same probability, and that the PCP can be constructed deterministically in polynomial time. Then for all constants $p > 2$ (and for $p = \infty$), γ -GapSVP $_p$ on $n^{O(1)}$ dimensional lattices is NP hard, where $\gamma = (1/s)^{\Theta(1)}$.*

In Section 7, we show that the PCP from Theorem 1.2 can be regularized via the techniques of [HM23] to get a PCP that satisfies all preconditions of Theorem 1.3, where the alphabet size is $|\Sigma| = 2^{\Theta(\log^{1-\varepsilon} n)}$ and the soundness is $s = 2^{-\Theta(\log^{1-\varepsilon} n)}$. From this along with Theorem 1.3 we derive the following.

Corollary 1.4. *For all constants $\varepsilon > 0$ and $p > 2$, γ -GapSVP $_p$ on n dimensional lattices is NP hard, where $\gamma = 2^{\log^{1-\varepsilon} n}$.*

Remark 1.5. We suspect that, by adapting the PCP given by Dinur, Harsha, and Kindler [DHK15], the same NP hardness result should hold for slightly *subconstant* values of ε . We do not pursue such an improvement in this paper.

Corollary 1.4 succinctly resolves both Question 1 and Question 2 in the affirmative. Due to the elementary nature of our techniques (see Section 3 and Section 4), we believe that our approach should inspire hardness of approximation results for several other lattice problems, hopefully paving the way⁴ to a deterministic NP hardness of approximation result for γ -GapSVP $_2$.

Another consequence of Theorem 1.3 is that we can get polynomial factor NP hardness of approximation under the Sliding Scale Conjecture [BGLR93], which posits that there exists a PCP for SAT with the same parameters as in Theorem 1.2, except $|\Sigma| = n^{\Theta(1)}$ and $s = n^{-\Theta(1)}$.⁵ Details are given in Section 7.

Corollary 1.6. *Assuming the Sliding Scale Conjecture, γ -GapSVP $_p$ on n dimensional lattices is NP hard for all constants $p > 2$ (and for $p = \infty$), where $\gamma = n^c$ for a constant $c > 0$ that depends on p .*

Mukhopadhyay [Muk22] recently proved a weaker version of this result: she showed that the *Projection Games Conjecture* implies NP hardness of $n^{\Omega(1)}$ -GapSVP $_\infty$, that is, for the special case that $p = \infty$. This alternative conjecture posits that the Sliding Scale Conjecture holds even for projection games, which are a special type of PCP in which every constraint depends on $q = 2$ variables, and additionally every constraint is of a particular form [Mos12]. To get a sense of the relative strength of these two conjectures, observe that the Projection Games Conjecture implies the Sliding Scale Conjecture, but it is unknown whether the reverse implication holds. Indeed, the current best polynomial time reduction from SAT to a projection game achieves soundness $(\log n)^{-O(1)}$ [DS14], as opposed to the $2^{-(\log n)^{0.99}}$ soundness for general PCPs.

³Our Hadamard matrix gadgets bear topical resemblance to the “norm measuring rows” in Dinur’s NP hardness of approximation result for γ -GapSVP $_\infty$ [Din02], but it appears that we leverage properties of the gadgets not discussed in the SVP literature before. Additionally, Dinur’s hardness result for γ -GapSVP $_\infty$ relies on a novel characterization of NP in terms of a gap problem referred to as SSAT, which is very different from the regularized PCP we leverage.

⁴It’s also worth noting that we found our hardness of approximation result as a byproduct of research on new sources of hardness to build cryptographic primitives such as public key encryption and witness encryption; we expect the connection between complexity-theoretic problems and cryptographic problems to inspire even more results.

⁵It’s known that any non-degenerate PCP requires $|\Sigma| \geq (1/s)^{\Omega(1/q)}$, so when q is constant we cannot hope to achieve polynomially small soundness without a polynomially large alphabet.

2 Preliminaries

We use $[n]$ to denote the set $\{1, \dots, n\}$. All vectors are row vectors unless stated otherwise. Assume that all quantities are rounded integers as needed.

Regularized PCPs. Regularization is a classic tool for getting hardness of approximation results from the PCP theorem. The idea is to convert an arbitrary PCP into a new PCP where every variable in the proof is read with equal probability. More than three decades ago, Papadimitriou and Yannakakis showed that any PCP can be converted into a regularized PCP, but with the caveat that the final PCP has at best constant soundness [PY88]. Recently, Hirahara and Moshkovitz gave an elegant technique which allows the final PCP to have soundness that is polynomially related to that of the starting PCP [HM23]. For the convenience of the reader, we prove the following result in Appendix A.

Theorem 2.1 (Simplified version of theorem in [HM23]). *Assume that we have a PCP verifier \mathcal{V} which uses r random bits to make q queries to a proof over alphabet Σ and has soundness s . Assume further that $s \leq \min(1/(3q), 1/|\Sigma|^c)$, where c is a constant with $0 < c \leq 1$. Then we can construct a new PCP verifier \mathcal{V}' from \mathcal{V} deterministically in $|\mathcal{V}|^{O(1)}$ time, such that the following holds. \mathcal{V}' uses $r + O(\log(1/s))$ random bits to make $q^{O(1)}$ queries to a proof over the same alphabet Σ and has soundness $s^{\Theta(1)}$. Additionally, every proof symbol is read for exactly $d = (q/s)^{\Theta(1)}$ choices of randomness.*

Hölder’s Inequality. This inequality generalizes the Cauchy-Schwarz inequality.

Theorem 2.2 ([Höl89]). *Let $\mathbf{u}, \mathbf{v} \in \mathbb{Z}^n$, and let $p, q \geq 1$ satisfy $1/p + 1/q = 1$. Then*

$$\|\mathbf{u} \odot \mathbf{v}\|_1 \leq \|\mathbf{u}\|_p \|\mathbf{v}\|_q,$$

where \odot denotes the component-wise product.

We make extensive use of the following corollary, the proof of which is deferred to Appendix B.

Corollary 2.3. *For all $\mathbf{w} \in \mathbb{Z}^n$ and $p \geq 2$ (including $p = \infty$),*

$$\|\mathbf{w}\|_p \geq n^{1/p-1/2} \|\mathbf{w}\|_2.$$

3 Setting up the Gadgets

In this section, we give an overview of our tools and briefly discuss how we plan to use them to reduce from a regularized PCP to the shortest vector problem.

3.1 Viewing PCPs as CSPs

A helpful way to think of PCPs, and indeed the viewpoint we will adopt for our reduction, is to interpret them as showing that a certain constraint satisfaction problem (CSP) is hard to approximate. We think of every choice of randomness as corresponding to a distinct constraint in the CSP, and every position in the proof as a variable in the CSP. There are exactly $M = 2^r$ choices of randomness, so our CSP will have exactly M constraints. For every choice of randomness, there is an explicit list of assignments for the relevant proof symbols that would cause the verifier to accept; these become the satisfying assignments for the corresponding constraint in the CSP. To emphasize the CSP viewpoint, we refer to a PCP with $M = 2^r$ choices of randomness that reads a proof of length N as “a PCP having M constraints and N variables.”

Assuming the PCP is regular, every variable will appear in exactly d constraints and every constraint will depend on exactly q variables, so we have $Nd = Mq$.

If the SAT instance from which we constructed the PCP was satisfiable, then by definition there must exist a proof which makes the verifier accept with probability 1. From our CSP viewpoint, this corresponds to an assignment to the variables such that every constraint is satisfied. On the other hand, if the original SAT instance was not satisfiable, then for any purported proof the PCP verifier will accept with probability at most the soundness parameter s . From the CSP perspective, this means that any assignment to the variables will satisfy at most an s fraction of the constraints.

Associated Graphs and Matrices. Below, we define a graph which corresponds to the CSP. The graph will encode which constraints are incident to which variables, as well as which variable assignments satisfy each constraint.

Definition 3.1 (Label Extended Factor Graph). *Consider a PCP with N variables, M constraints, and alphabet Σ , where each constraint depends on q variables. Its label extended factor graph is the bipartite graph $([M] \times \Sigma^q, [N] \times \Sigma, E)$, where $((t, \alpha_1, \dots, \alpha_q), (x, \alpha)) \in E$ iff*

1. The (ordered) assignment $\alpha_1, \dots, \alpha_q$ is a satisfying assignment for constraint t ,
2. Constraint t contains the variable x , and
3. The (ordered) assignment $\alpha_1, \dots, \alpha_q$ indicates an assignment of α to the variable x .

In other words, each left vertex $(t, \alpha_1, \dots, \alpha_q)$ represents a *candidate* assignment for constraint t , and we list every such candidate. Similarly, each right vertex (x, α) represents a (variable, assignment) tuple. The left vertices $(t, \alpha_1, \dots, \alpha_q)$ which do not correspond to a satisfying assignment for constraint t will have no edges incident to them. (The reason we keep these vertices is simply to ensure the label extended factor graph has exactly $M|\Sigma|^q$ left vertices and $N|\Sigma|$ right vertices.)

We can also think of the label extended factor graph in terms of its indicator matrix:

Definition 3.2 (Indicator Matrix). *The indicator matrix $\mathbf{S} \in \{0, 1\}^{M|\Sigma|^q \times N|\Sigma|}$ for a given PCP has $\mathbf{S}_{(t, \alpha_1, \dots, \alpha_q), (x, \alpha)} = 1$ iff $((t, \alpha_1, \dots, \alpha_q), (x, \alpha))$ is an edge in the PCP's label extended factor graph.*

Below, we give some toy examples of indicator matrices.

1		1		
1		1		
	1		1	1
	1	1	1	
1				1
1			1	1
1	1			
1	1			
		1		1
		1	1	1
		1	1	1

The indicator matrix for a PCP that has $M = 8$ constraints, $N = 5$ variables, $q = 2$ variables read by each constraint, and $|\Sigma| = 2$. For ease of viewing, we omit rows that are entirely zero, and we only write the “1” entries. Notice that this PCP is satisfied by the assignment that corresponds to picking the left column for each variable.

1		1		
1		1		
	1		1	1
	1	1	1	
1				1
1			1	1
1	1			
1	1			
		1		1
		1	1	1
		1	1	1

The indicator matrix for a PCP with the same parameters as the one on the left, but which is unsatisfiable.

3.2 Matrix Gadgets

We form the set of basis vectors for our SVP problem by taking the indicator matrix \mathbf{S} for a regularized PCP, and then augmenting \mathbf{S} using different types of gadgets. (A modified copy of the matrix \mathbf{S} will directly appear within the final set of basis vectors.) In this section we introduce the gadgets themselves, with a particular focus on how they can help us ensure soundness of the reduction.

Reduced Vandermonde Matrices. The first gadget is a matrix where all entries are polynomially bounded integers, and every row-induced square submatrix has full rank.

Definition 3.3 (Reduced Vandermonde Matrix). *Let $a > b$ with a prime. We define an (a, b) reduced Vandermonde matrix $\mathbf{V} \in \mathbb{Z}^{(a-1) \times b}$ as $\mathbf{V}_{i,j} = i^{j-1} \pmod a$.*

We have the following property:

Lemma 3.4. *Every $b \times b$ submatrix of an (a, b) reduced Vandermonde matrix is of full rank.*

Proof. Well-known; see for example [HJ94]. The idea is that an $(a-1) \times b$ Vandermonde matrix over \mathbb{F}_a has every $b \times b$ submatrix being of full rank, and casting from \mathbb{F}_a to \mathbb{Z} cannot introduce new linear dependencies. \square

Thus, any linear combination of the rows which sums to zero must either (i) have all coefficients equal to zero, or (ii) have a large fraction of the coefficients being nonzero. More formally,

Corollary 3.5. *Let $\mathbf{v} \in \mathbb{Z}^{a-1}$ be a nonzero vector. If $\mathbf{v}\mathbf{V} = \mathbf{0}$, then $\|\mathbf{v}\|_0 > b$.*

Reduced Vandermonde matrices are used in two different parts of our construction:

1. We concatenate a single large reduced Vandermonde matrix (with its entries scaled up by a large multiplicative factor) to the matrix of basis vectors for our lattice. This enforces that any short nonzero lattice vector must correspond to a linear combination of basis vectors which entirely cancels out the Vandermonde component. With the right size parameters, we can ensure that the number of nonzero coefficients in any such linear combination is nearly the same as the number of constraints in the PCP.
2. We insert many relatively small reduced Vandermonde matrices (again with their entries scaled up by a large multiplicative factor) into the copy of the indicator matrix \mathbf{S} contained within the lattice basis vectors. The purpose of this transformation is a bit more subtle; roughly speaking, it enforces that any short lattice vector corresponds to a linear combination of basis vectors that is at least partially “self-consistent” with respect to the variable assignments. See Section 4 for details.

Hadamard Matrices. The second gadget is a matrix where every pair of rows is orthogonal (in the ℓ_2 norm), and all of the entries are integers of the same magnitude. These conditions will allow us to apply manipulations based on the relationship between the ℓ_2 norm and the ℓ_p norm with $p > 2$.

Definition 3.6 (Hadamard Matrix). *Define $\mathbf{H}_0 := [1]$, and inductively set*

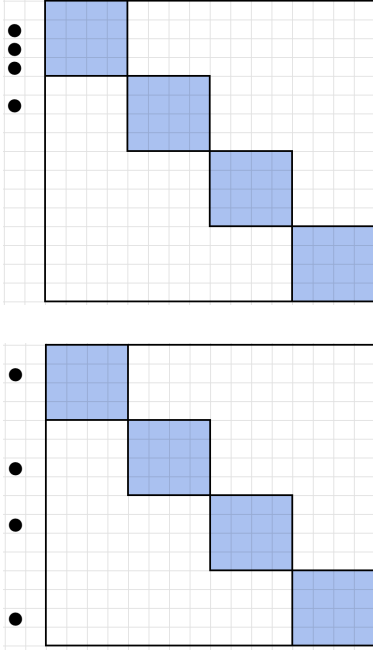
$$\mathbf{H}_{i+1} := \begin{bmatrix} \mathbf{H}_i & -\mathbf{H}_i \\ \mathbf{H}_i & \mathbf{H}_i \end{bmatrix}$$

For n a power of two, the $n \times n$ Hadamard matrix is $\mathbf{H}_{\log_2 n}$.

In our construction, we use a *block diagonal* set of Hadamard matrices concatenated to the set of basis vectors. To gain some intuition, let $\mathbf{Y} := \mathbf{I}_m \otimes \mathbf{H}_{\log_2 n}$ and consider the ℓ_p norm of linear combinations $\mathbf{v}\mathbf{Y}$.

1. For two vectors \mathbf{u}, \mathbf{v} with all entries coming from the set $\{-1, 0, 1\}$, if $\|\mathbf{u}\|_0 > \|\mathbf{v}\|_0$, then $\|\mathbf{u}\mathbf{Y}\|_2 > \|\mathbf{v}\mathbf{Y}\|_2$. This just follows from the orthogonality of the rows in \mathbf{Y} . In fact, we can make similar statements for ℓ_p with $p > 2$, assuming certain properties about \mathbf{u} and \mathbf{v} . With the right size parameters, the Hadamard gadgets will allow us to argue that any short nonzero vector in our final lattice must come from a linear combination of basis vectors where the number of nonzero coefficients is not too much larger than the number of constraints in the PCP. (See Section 4 for details.) This complements the effect of the reduced Vandermonde matrices discussed earlier.
2. In higher ℓ_p norms, \mathbf{Y} exhibits a certain “anti-concentration” property. Consider two vectors \mathbf{u} and \mathbf{v} with all entries coming from the set $\{-1, 0, 1\}$, such that $\|\mathbf{u}\|_0 = \|\mathbf{v}\|_0$. If the nonzero entries in \mathbf{u} correspond to a relatively small number of Hadamard blocks, but the nonzero entries in \mathbf{v} are relatively spread out across the blocks, then we will have $\|\mathbf{u}\mathbf{Y}\|_p \gg \|\mathbf{v}\mathbf{Y}\|_p$ for all $p > 2$. (A concrete analysis of this property is given in Lemma 6.5.) This is useful because we can force short lattice vectors to come from a linear combination of basis vectors that is spread across many of the different PCP constraints.

Below, we give an illustration showing the anti-concentration property.



(Top) The linear combination corresponding to \mathbf{u} , with $\|\mathbf{u}\|_0 = 4$. (Bottom) The linear combination corresponding to \mathbf{v} , with $\|\mathbf{v}\|_0 = 4$. In each diagram, the rows of the matrix \mathbf{Y} participating in the linear combination are indicated by black dots.

We have that $\mathbf{u}\mathbf{Y}$ is a vector containing mostly zeros, but the nonzero entries will (on average) be much greater than 1. In higher ℓ_p norms, the length contribution from these entries is amplified significantly. In contrast, $\mathbf{v}\mathbf{Y}$ is a vector whose entries are all either -1 or 1 ; while the number of nonzero entries in $\mathbf{v}\mathbf{Y}$ is much larger than in $\mathbf{u}\mathbf{Y}$, all of them are small, so ℓ_p length amplification doesn't occur.

4 Reducing from a PCP Instance to an SVP Instance

We start with a regularized PCP that has M constraints and N variables over alphabet Σ , where each constraint depends on exactly q variables, each variable appears in exactly d constraints, and the soundness parameter is s . Assume that $q \leq O(\log M)$, that $(1/s)^{1/q} \geq (\log M)^{\omega(1)}$, and that the alphabet size $|\Sigma|$ is a power of two. Two additional properties hold automatically:

1. q must be at least 2. Otherwise, it is trivial to find a maximally satisfying assignment for the PCP.
2. d must be at least $1/(sq) \geq (\log M)^{\omega(1)}$. This is because we can always find a set of $M/(dq)$ variable-disjoint constraints, and these are mutually satisfied by taking a locally satisfying assignment for each. If $d < 1/(sq)$, then the number of satisfied constraints is strictly more than $M/(1/s) = Ms$, violating soundness.

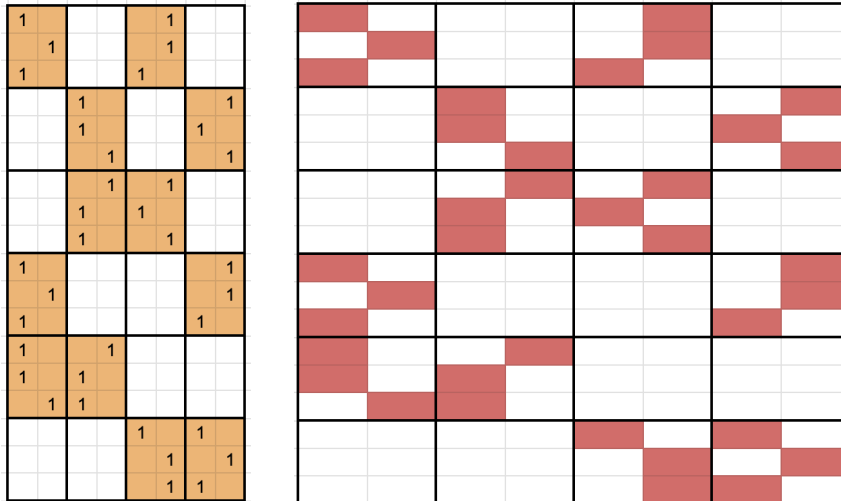
Throughout, fix a prime $a = \Theta((M|\Sigma|^q)^2)$.

We construct the matrix of basis vectors for a lattice by applying two procedures, each of which introduces a distinct set of gadgets. (The construction is agnostic to the specific choice of ℓ_p norm so long as $p > 2$.) We encourage the reader to adopt the following perspective: Each individual gadget converts many of the previously “short problematic” vectors into vectors with a large ℓ_p norm. When all gadgets are considered together, we show that we can map all of the remaining short vectors to assignments for the PCP which violate soundness. We will be able to go in the other direction, too: a satisfying assignment for the PCP can be mapped to a nonzero vector with short ℓ_p norm in the final lattice.

Manipulating the Indicator Matrix. Our first procedure constructs a matrix \mathbf{P} based on the indicator matrix for the PCP by inserting reduced Vandermonde matrices in place of each nonzero entry.

1. Let $\mathbf{S} \in \{0, 1\}^{M|\Sigma|^q \times N|\Sigma|}$ be the indicator matrix for the PCP. Assume that $|\Sigma|$ is a power of two.
2. Construct matrix $\mathbf{P} \in \mathbb{Z}^{M|\Sigma|^q \times Mq|\Sigma|/(\log^3 M)}$ as follows. Start with an empty matrix \mathbf{P} , and for each column \mathbf{c} of \mathbf{S} :
 - (a) Let $\mathbf{V} \in \mathbb{Z}^{(a-1) \times d/(\log^3 M)}$ be an $(a, d/(\log^3 M))$ reduced Vandermonde matrix.
 - (b) Define a matrix $\mathbf{V}' \in \mathbb{Z}^{M|\Sigma|^q \times d/(\log^3 M)}$ as follows:
 - i. If $\mathbf{c}_{(t, \alpha_1, \dots, \alpha_q)} = 0$, then $\mathbf{V}'_{(t, \alpha_1, \dots, \alpha_q)} = \mathbf{0}^{d/(\log^3 M)}$.
 - ii. If $\mathbf{c}_{(t, \alpha_1, \dots, \alpha_q)} = 1$, we set $\mathbf{V}'_{(t, \alpha_1, \dots, \alpha_q)}$ to be a distinct row⁶ of \mathbf{V} , scaled up by a multiplicative factor of $(M|\Sigma|^q/s)^2$.
 - (c) Insert \mathbf{V}' as a new set of columns in \mathbf{P} .

Recall that because the PCP is regular, we must have $Nd = Mq$. So the final width of \mathbf{P} is indeed $N|\Sigma| \cdot d/(\log^3 M) = Mq|\Sigma|/(\log^3 M)$. The construction is well-defined because $d \geq (\log M)^{\omega(1)}$, meaning $d/(\log^3 M) \geq (\log M)^{\omega(1)}$.



(Left) The indicator matrix \mathbf{S} for a PCP with $N = 4$, $M = 6$, $|\Sigma| = 2$, and $q = 2$. As before, we omit zero rows and only write the “1” entries. (Right) The corresponding matrix \mathbf{P} . Every red block is a row taken from a reduced Vandermonde matrix.

⁶We only need $\mathbf{V}'_{(t, \alpha_1, \dots, \alpha_q)}$ to be distinct with respect to the other \mathbf{V}' rows in the current matrix \mathbf{V}' , not with respect to the rows in matrices \mathbf{V}' constructed previously.

Why is this useful? Consider an arbitrary vector $\mathbf{v} \in \mathbb{Z}^{M|\Sigma|^q}$, and let \mathcal{S} be the set of rows from \mathbf{P} having a nonzero coefficient in the linear combination $\mathbf{v}\mathbf{P}$. Because every Vandermonde entry is scaled up by a “very large” multiplicative factor, we have that $\mathbf{v}\mathbf{P}$ is either zero or has a large ℓ_p norm, for any choice of $p \geq 1$. In the case that $\mathbf{v}\mathbf{P}$ is zero, an application of Corollary 3.5 on the reduced Vandermonde gadgets shows that a certain gap property holds: *every* set of columns corresponding to a (variable, assignment) tuple in \mathbf{P} satisfies one of the following properties.


1. None of the basis vectors in \mathcal{S} have a nonzero entry in that set of columns.
2. At least $d/(\log^3 M)$ of the basis vectors in \mathcal{S} have nonzero entries in that set of columns.

We show in Lemma 6.9 that this gap property ensures any short lattice vector resulting from a linear combination of (say) M basis vectors cannot have those basis vectors “indicating” more than $f \cdot N$ (variable, assignment) tuples, for some small slack factor f . The conversion from a bound in terms of M to a bound in terms of N makes critical use of the fact that the PCP is regular.

Appending More Gadgets. In the second procedure, we append Hadamard matrices and another reduced Vandermonde matrix to \mathbf{P} . The total number of basis vectors does not change; instead, the width of each basis vector increases significantly.

1. Let $\mathbf{X} \in \mathbb{Z}^{M|\Sigma|^q \times M/(\log^3 M)}$ be the first $M|\Sigma|^q$ rows of an $(a, M/(\log^3 M))$ reduced Vandermonde matrix, but where each entry is scaled up by a multiplicative factor of $(M|\Sigma|^q/s)^2$. As before, this is well-defined because $d \geq (\log M)^{\omega(1)}$, meaning $d/(\log^3 M) \geq (\log M)^{\omega(1)}$.
2. Let $\mathbf{H} \in \{\pm 1\}^{|\Sigma|^q \times |\Sigma|^q}$ be a Hadamard matrix. Define $\mathbf{Y} = \mathbf{I}_M \otimes \mathbf{H}$, where \mathbf{I}_M is the $M \times M$ identity matrix and \otimes is the Kronecker product. Each copy of the Hadamard matrix should align with a set of rows in \mathbf{P} representing the candidate assignments for a single constraint.
3. Define $\mathbf{G} := [\mathbf{P} \parallel \mathbf{X} \parallel \mathbf{Y}]$, and then delete all rows of \mathbf{G} which are entirely zero when restricted to the submatrix \mathbf{P} . All remaining rows in \mathbf{G} are those which correspond to a *satisfying assignment* for some constraint. The height M' of \mathbf{G} is upper bounded as $M' \leq M|\Sigma|^q$, and the width N' of \mathbf{G} is exactly $N' := Mq|\Sigma|/(\log^3 M) + M/(\log^3 M) + M|\Sigma|^q$. Because $2 \leq q \leq O(\log M)$, we have $N' = \Theta(M|\Sigma|^q)$.

Below, we give an illustration of the matrix \mathbf{G} . Reduced Vandermonde gadgets are marked in red, and Hadamard gadgets are marked in blue. Because of the row deletion step applied to \mathbf{G} , all of the Hadamard matrices (which were originally square) are now wide. Asymptotically, the main width contribution is from the Hadamards.



$\gamma = (\log M)^{\omega(1)}$. Because $M = 2^r$ (recall that r is the number of random bits used by the verifier), we can rewrite these inequalities as $2 \leq q \leq O(r)$ and $(1/s)^{1/q} \geq r^{\omega(1)}$. It's clear from Section 4 that the construction proceeds deterministically in $(2^r |\Sigma|^q)^{O(1)}$ time. We claim:

Claim 4.1 (Completeness). *If the PCP has a satisfying assignment, then assuming M is sufficiently large, there exists a nonzero vector $\mathbf{w} \in \mathcal{L}$ with $\|\mathbf{w}\|_p \leq (N')^{1/p}$.*⁸

Claim 4.2 (Soundness). *If the PCP does not have a satisfying assignment, then assuming M is sufficiently large, every nonzero vector $\mathbf{w} \in \mathcal{L}$ has $\|\mathbf{w}\|_p > \gamma(N')^{1/p}$.*

From these claims we deduce the main theorem.

Theorem 4.3. *Suppose there exists a PCP for SAT instances of size n that uses r random bits to read $q \leq O(r)$ locations in a proof over alphabet Σ and has soundness parameter s . Suppose further that every proof location is read with equal probability, that $(1/s)^{1/q} \geq r^{\omega(1)}$, and that the PCP can be constructed deterministically in time $T(n)$. Then for all constants $p > 2$ (and for $p = \infty$), there exists a $T(n) + (2^r |\Sigma|^q)^{O(1)}$ time deterministic reduction from SAT instances of size n to instances of γ -GapSVP _{p} on lattices of dimension $O(2^r |\Sigma|^q)$, where $\gamma = (1/s)^{(1/2-1/p)/(25q)}$.*

5 Proof of Claim 4.1 (Completeness)

Here we show that a satisfying assignment for the PCP maps to a short lattice vector. We will use the following proposition to help zero out the reduced Vandermonde gadgets, which ensures that our lattice vector will not have any large-magnitude entries.

Proposition 5.1. *Let $\mathbf{M} \in \mathbb{Z}^{n \times O(n/\log^2 n)}$ be a matrix with entries in $\{-n^{O(1)}, \dots, n^{O(1)}\}$. Then assuming n is sufficiently large, there exists a nonzero vector $\mathbf{v} \in \{-1, 0, 1\}^n$ with $\mathbf{v}\mathbf{M} = \mathbf{0}$.*

Proof. There are 2^n vectors $\mathbf{w} \in \{0, 1\}^n$, and there are $n^{O(n/\log^2 n)} = 2^{O(n/\log n)}$ possible values for $\mathbf{w}\mathbf{M}$. By the pigeonhole principle, there must exist two distinct vectors $\mathbf{w}', \mathbf{w}''$ with $\mathbf{w}'\mathbf{M} = \mathbf{w}''\mathbf{M}$. Take $\mathbf{v} = \mathbf{w}' - \mathbf{w}''$. \square

Now let σ be a satisfying assignment for the PCP, and let \mathbf{G}' be the submatrix of \mathbf{G} induced by the M rows corresponding to σ . (\mathbf{G}' consists of exactly one row per PCP constraint.) We can write $\mathbf{G}' = [\mathbf{P}' \parallel \mathbf{X}' \parallel \mathbf{Y}']$, where $\mathbf{P}', \mathbf{X}', \mathbf{Y}'$ are the corresponding row induced submatrices of $\mathbf{P}, \mathbf{X}, \mathbf{Y}$.

Observe that the nonzero entries in \mathbf{G}' only touch the columns of \mathbf{P} corresponding to one consistent assignment for each variable, which makes for a total of $Mq/(\log^3 M)$ columns. The submatrix \mathbf{X} has width $M/(\log^3 M)$, so the nonzero entries in \mathbf{G}' in fact touch $Mq/(\log^3 M) + M/(\log^3 M) = O(M/(\log^2 M))$ columns across both \mathbf{P} and \mathbf{X} , where we used that $q \leq O(\log M)$. Thus we can apply Proposition 5.1 to guarantee the existence of a nonzero vector $\mathbf{v} \in \{-1, 0, 1\}^M$ with $\mathbf{v}\mathbf{P}' = \mathbf{0}$ and $\mathbf{v}\mathbf{X}' = \mathbf{0}$.

The above implies that all of the nonzero entries in $\mathbf{v}\mathbf{G}'$ will come from $\mathbf{v}\mathbf{Y}'$. Due to the structure of \mathbf{Y} and the fact that \mathbf{G}' consists of exactly one row per PCP constraint, every column of \mathbf{Y}' has exactly one nonzero entry. This implies $\|\mathbf{v}\mathbf{Y}'\|_\infty = 1$ and thus $\|\mathbf{v}\mathbf{G}'\|_\infty = 1$. By padding \mathbf{v} with zeros we get a vector \mathbf{v}' such that $\|\mathbf{v}'\mathbf{G}\|_\infty = 1$. Since the width of \mathbf{G} is defined as N' , this implies that $\|\mathbf{v}'\mathbf{G}\|_p \leq (N')^{1/p}$, and we can take $\mathbf{w} = \mathbf{v}'\mathbf{G}$ as our short lattice vector.

⁸In the case of $p = \infty$, we have $(N')^{1/p} = 1$.

6 Proof of Claim 4.2 (Soundness)

Here we prove that, if the PCP does not have a satisfying assignment, every nonzero vector $\mathbf{w} \in \mathcal{L}$ will have a relatively large ℓ_p norm. Most of the work is dedicated to ruling out certain nonzero lattice points $\mathbf{w} \in \mathcal{L}$, based on the structure of the corresponding vectors \mathbf{v} with $\mathbf{v}\mathbf{G} = \mathbf{w}$.

The first lemma⁹ shows that if \mathbf{v} has too few nonzero entries, then $\|\mathbf{v}\mathbf{G}\|_p$ must be large. We make use of the reduced Vandermonde gadget \mathbf{X} .

Lemma 6.1. *Assume that $\mathbf{v} \in \mathbb{Z}^{M'}$ is a nonzero vector with $\|\mathbf{v}\|_0 \leq M/(\log^3 M)$. Then assuming M is sufficiently large, $\|\mathbf{v}\mathbf{G}\|_p \geq (M|\Sigma|^q/s)^2 > \gamma(N')^{1/p}$.*

Proof. Observe that $\|\mathbf{v}\mathbf{G}\|_p \geq \|\mathbf{v}\mathbf{X}\|_p$, because \mathbf{X} is a column-induced submatrix of \mathbf{G} . Recall that \mathbf{X} consists of a subset of rows taken from an $(a, M/(\log^3 M))$ reduced Vandermonde matrix, where every entry is multiplied by $(M|\Sigma|^q/s)^2$. By Corollary 3.5, we know that $\mathbf{v}\mathbf{X}$ must have at least one nonzero coordinate, and the smallest nonzero value possible is $(M|\Sigma|^q/s)^2$. This gives $\|\mathbf{v}\mathbf{G}\|_p \geq (M|\Sigma|^q/s)^2$, and since $q \geq 2$, $s \leq 1$, and $p > 2$, we have $(M|\Sigma|^q/s)^2 > \gamma N' = \Theta((1/s)^{(1/2-1/p)/(25q)} M|\Sigma|^q)$ when M is sufficiently large. \square

The next lemma¹⁰ gives an upper bound on $\|\mathbf{v}\|_0$. Here, we make use of the block-diagonal Hadamard gadgets in submatrix \mathbf{Y} .

Lemma 6.2. *Suppose $\mathbf{v} \in \mathbb{Z}^{M'}$ is a nonzero vector with $\|\mathbf{v}\|_0 \geq M \cdot \gamma^3$. Then assuming M is sufficiently large, $\|\mathbf{v}\mathbf{G}\|_p > \gamma(N')^{1/p}$.*

Proof. Because \mathbf{Y} is a column-induced submatrix of \mathbf{G} , we have $\|\mathbf{v}\mathbf{G}\|_p \geq \|\mathbf{v}\mathbf{Y}\|_p$. Since each row of \mathbf{Y} is a vector in $\{\pm 1\}^{|\Sigma|^q}$, and every pair of rows is orthogonal (in the ℓ_2 norm), we have that

$$\|\mathbf{v}\mathbf{Y}\|_2 = \Omega\left(\sqrt{\|\mathbf{v}\|_0 |\Sigma|^q}\right) = \Omega\left(\sqrt{M \cdot \gamma^3 \cdot |\Sigma|^q}\right) = \Omega(\gamma^{3/2} \cdot \sqrt{N'}),$$

where the last equality holds because $N' = \Theta(M|\Sigma|^q)$.

Because $p > 2$, an application of Corollary 2.3 gives

$$\begin{aligned} \|\mathbf{v}\mathbf{Y}\|_p &\geq (N')^{1/p-1/2} \cdot \|\mathbf{v}\mathbf{Y}\|_2 \\ &= \Omega((N')^{1/p-1/2} \cdot \gamma^{3/2} \cdot \sqrt{N'}) \\ &= \Omega(\gamma^{3/2} \cdot (N')^{1/p}). \end{aligned}$$

Since $\gamma \geq (\log M)^{\omega(1)}$, we have $\Omega(\gamma^{3/2}) > \gamma$ for sufficiently large M , in which case $\|\mathbf{v}\mathbf{G}\|_p > \gamma(N')^{1/p}$. \square

Later, we actually use the contrapositive of Lemmas 6.1 and 6.2. In particular, we know that if $\|\mathbf{v}\mathbf{G}\|_p \leq \gamma(N')^{1/p}$ and \mathbf{v} is nonzero, then $M/(\log^3 M) < \|\mathbf{v}\|_0 < M \cdot \gamma^3$.

In the remaining lemmas, we need to formalize the correspondence between nonzero entries in the vector \mathbf{v} and different components of \mathbf{G} and of the PCP. Our first definition maps from nonzero entries of \mathbf{v} to rows of \mathbf{G} . (Recall that each row of \mathbf{G} is a basis vector of the lattice.)

Definition 6.3 (Indicated Basis Vectors). *A row vector \mathbf{r} from \mathbf{G} is said to be indicated by \mathbf{v} if the linear combination $\mathbf{v}\mathbf{G}$ assigns a nonzero coefficient to \mathbf{r} .*

⁹This lemma is also true for the case that $p = 2$.

¹⁰This lemma is also true for the case that $p = 2$.

The next definition formalizes the correspondence between \mathbf{v} and different PCP constraints.

Definition 6.4 (Indicated Constraint). *We say that \mathbf{v} indicates constraint t if \mathbf{v} indicates at least one row vector representing an assignment to constraint t .*

In the following we use the Hadamard matrices again, this time exploiting (in a formal manner) the “anti-concentration” property.¹¹

Lemma 6.5. *Let $\mathbf{v} \in \mathbb{Z}^{M'}$ be a nonzero vector, and assume that \mathbf{v} indicates at most $M/\gamma^{2/(1/2-1/p)}$ distinct constraints of the PCP. Then assuming M is sufficiently large, $\|\mathbf{v}\mathbf{G}\|_p > \gamma(N')^{1/p}$.*

Proof. We can assume that $\|\mathbf{v}\|_0 \geq M/(\log^3 M)$; otherwise by Lemma 6.1, we already have $\|\mathbf{v}\mathbf{G}\|_p \geq (M|\Sigma|^q/s)^2 > \gamma(N')^{1/p}$. As before, we can write $\|\mathbf{v}\mathbf{G}\|_p \geq \|\mathbf{v}\mathbf{Y}\|_p$.

We assumed that \mathbf{v} indicates at most $M/\gamma^{2/(1/2-1/p)}$ distinct constraints of the PCP. So by the construction of \mathbf{Y} , the basis vectors indicated by \mathbf{v} have nonzero entries in at most $(M/\gamma^{2/(1/2-1/p)}) \cdot |\Sigma|^q = \Theta(N'/\gamma^{2/(1/2-1/p)})$ columns of \mathbf{Y} . Thus we can find a column-induced submatrix $\hat{\mathbf{Y}}$ of \mathbf{Y} such that the width of $\hat{\mathbf{Y}}$ is at most $\Theta(N'/\gamma^{2/(1/2-1/p)})$, and $\|\mathbf{v}\mathbf{Y}\|_p = \|\mathbf{v}\hat{\mathbf{Y}}\|_p$. Since each row of $\hat{\mathbf{Y}}$ is a vector in $\{\pm 1\}^{|\Sigma|^q}$, and every pair of rows is orthogonal (in the ℓ_2 norm), we have that

$$\|\mathbf{v}\hat{\mathbf{Y}}\|_2 = \Omega\left(\sqrt{\|\mathbf{v}\|_0|\Sigma|^q}\right) = \Omega\left(\sqrt{\frac{M|\Sigma|^q}{\log^3 M}}\right) = \Omega(\sqrt{N'}/\log^{3/2} M),$$

since $N' = \Theta(M|\Sigma|^q)$.

Using our upper bound on the width of $\hat{\mathbf{Y}}$, and the fact that $p > 2$, an application of Corollary 2.3 gives

$$\begin{aligned} \|\mathbf{v}\hat{\mathbf{Y}}\|_p &\geq (N'/\gamma^{2/(1/2-1/p)})^{1/p-1/2} \cdot \|\mathbf{v}\hat{\mathbf{Y}}\|_2 \\ &= (N')^{1/p-1/2} \cdot \gamma^2 \cdot \|\mathbf{v}\hat{\mathbf{Y}}\|_2 \\ &\geq \Omega((N')^{1/p-1/2} \cdot \gamma^2 \cdot \sqrt{N'}/\log^{3/2} M) \\ &= \Omega((N')^{1/p} \cdot \gamma^2 / \log^{3/2} M) \\ &> (N')^{1/p} \cdot \gamma \end{aligned}$$

where we used that, because $\gamma = (\log M)^{\omega(1)}$, we have $\gamma^2 / \log^{3/2} M > \gamma$ when M is sufficiently large.

To summarize, we have shown that $\|\mathbf{v}\mathbf{G}\|_p \geq \|\mathbf{v}\mathbf{Y}\|_p = \|\mathbf{v}\hat{\mathbf{Y}}\|_p > (N')^{1/p} \cdot \gamma$, completing the proof. \square

Now our plan is to use the reduced Vandermonde matrices inserted within submatrix \mathbf{P} to argue that any nonzero vector \mathbf{v} with $\mathbf{v}\mathbf{G}$ being short must indicate a small number of (variable, assignment) tuples in the original PCP. First we make the notion of “indicated assignments” more concrete, and define a notion of multiplicity for assignments.

Definition 6.6 (Indicated Assignments). *Let $\mathbf{v} \in \mathbb{Z}^{M'}$ be any vector. Consider the row-induced submatrix $\hat{\mathbf{P}}$ of \mathbf{P} indicated by the nonzero entries in \mathbf{v} . (In other words, this is the submatrix of \mathbf{P} composed of all rows that are assigned a nonzero coefficient in the linear combination $\mathbf{v}\mathbf{P}$.) Every (variable, assignment) tuple (x, α) corresponds to a set of columns in \mathbf{P} , and we say that a tuple (x, α) is indicated by \mathbf{v} iff $\hat{\mathbf{P}}$ has a nonzero entry in that set of columns.*

¹¹This property only holds for p norms with $p > 2$.

Note that the set of all indicated assignments may be larger than the number of variables in the PCP. We now define a notion of *multiplicity*, which counts how many times \mathbf{v} indicates the same assignment for the same variable.

Definition 6.7 (Multiplicity). *We say that a tuple (x, α) has multiplicity h with respect to \mathbf{v} if there exist exactly h distinct nonzero entries of \mathbf{v} that indicate (x, α) .*

We also need a definition which counts the number of *different* assignments for the same variable.

Definition 6.8 (Distinct Assignment Count). *We say that a variable x has distinct assignment count h with respect to \mathbf{v} if there exist exactly h distinct values α such that \mathbf{v} indicates (x, α) .*

Below we demonstrate that, if $\mathbf{v}\mathbf{G}$ has short ℓ_p norm, then \mathbf{v} cannot indicate too many distinct (variable, assignment) tuples.

Lemma 6.9. *Let $\mathbf{v} \in \mathbb{Z}^{M'}$ be a nonzero vector, and assume that \mathbf{v} indicates at least $N \cdot \gamma^4$ distinct (variable, assignment) tuples. Then assuming M is sufficiently large, we have $\|\mathbf{v}\mathbf{G}\|_p \geq (M|\Sigma|^{q/s})^2 > \gamma(N')^{1/p}$.*

Proof. We know by Lemma 6.2 that if $\|\mathbf{v}\|_0 \geq M \cdot \gamma^3$, we automatically have $\|\mathbf{v}\mathbf{G}\|_p \geq (M|\Sigma|^{q/s})^2 > \gamma(N')^{1/p}$. So assume otherwise.

By construction, each nonzero entry of \mathbf{v} will indicate exactly q (variable, assignment) tuples. Thus, allowing for duplicates, \mathbf{v} indicates at most $Mq \cdot \gamma^3$ (variable, assignment) tuples. By the averaging principle and the assumption that \mathbf{v} indicates at least $N \cdot \gamma^4$ distinct (variable, assignment) tuples, there exists (x, σ) with multiplicity at least one and at most

$$\frac{Mq \cdot \gamma^3}{N \cdot \gamma^4} = d/\gamma,$$

where we used that $Nd = Mq$.

Now consider the column-induced submatrix $\hat{\mathbf{P}}$ of \mathbf{P} corresponding to (x, σ) . We know that all nonzero rows of $\hat{\mathbf{P}}$ are distinct rows of a width $d/(\log^3 M)$ reduced Vandermonde matrix, where all entries are scaled up by a factor of $(M|\Sigma|^{q/s})^2$. By the argument above, \mathbf{v} corresponds to a linear combination of at least one nonzero row from $\hat{\mathbf{P}}$ and at most d/γ nonzero rows from $\hat{\mathbf{P}}$. Using that $\gamma = (\log M)^{\omega(1)}$, we have that $d/\gamma \leq d/(\log M)^{\omega(1)} < d/\log^3 M$ for sufficiently large M . Therefore by Corollary 3.5, we have $\|\mathbf{v}\hat{\mathbf{P}}\|_p \geq (M|\Sigma|^{q/s})^2 > \gamma(N')^{1/p}$. Because $\|\mathbf{v}\mathbf{G}\|_p \geq \|\mathbf{v}\hat{\mathbf{P}}\|_p$, the lemma follows. \square

The final step is to combine all of the restrictions on \mathbf{v} to show that the remaining candidate short vectors must implicate soundness of the PCP.

Lemma 6.10. *Assume that the PCP from which we constructed \mathbf{G} was not satisfiable. Then for every nonzero vector $\mathbf{v} \in \mathbb{Z}^{M'}$, it holds that $\|\mathbf{v}\mathbf{G}\|_p > \gamma(N')^{1/p}$, assuming M is sufficiently large.*

Proof. Assume for contradiction that the PCP was not satisfiable, and there does exist a nonzero vector $\mathbf{v} \in \mathbb{Z}^{M'}$ with $\|\mathbf{v}\mathbf{G}\|_p \leq \gamma(N')^{1/p}$. By Lemmas 6.5 and 6.9 respectively, we know:

1. \mathbf{v} indicates at least $M/\gamma^{2/(1/2-1/p)}$ distinct constraints in the PCP.
2. \mathbf{v} indicates at most $N \cdot \gamma^4$ distinct (variable, assignment) tuples.

Let $\mathcal{X}_{\text{high}}$ be the set of all variables x whose distinct assignment count with respect to \mathbf{v} is at least $\gamma^{7+2/(1/2-1/p)}$. Our first step is to upper bound $|\mathcal{X}_{\text{high}}|$.

Claim 6.11. $|\mathcal{X}_{\text{high}}| \leq N/\gamma^{2+2/(1/2-1/p)}$.

Proof. Suppose not, then we would need for \mathbf{v} to indicate at least

$$\gamma^{7+2/(1/2-1/p)} \cdot N/\gamma^{2+2/(1/2-1/p)} = N \cdot \gamma^5$$

distinct (variable, assignment) tuples, which violates the upper bound from Lemma 6.9 because $\gamma = (\log M)^{\omega(1)} > 1$ for sufficiently large M . \square

Now we quantify the number of constraints that $\mathcal{X}_{\text{high}}$ can interact with. Define $\mathcal{T}_{\text{high}}$ as the set of all constraints in the PCP which are incident to at least one variable from $\mathcal{X}_{\text{high}}$.

Claim 6.12. $|\mathcal{T}_{\text{high}}| \leq M/\gamma^{1+2/(1/2-1/p)}$.

Proof. Because every variable appears in exactly d constraints, we have

$$\begin{aligned} |\mathcal{T}_{\text{high}}| &\leq d|\mathcal{X}_{\text{high}}| \\ &\leq Nd/\gamma^{2+2/(1/2-1/p)} \\ &= Mq/\gamma^{2+2/(1/2-1/p)} \\ &\leq M/\gamma^{1+2/(1/2-1/p)}, \end{aligned}$$

where we used that $Nd = Mq$ and $q = O(\log M) < \gamma = (\log M)^{\omega(1)}$ for sufficiently large M . \square

Let $\mathcal{T}_{\mathbf{v}}$ be the set of constraints indicated by \mathbf{v} , and define $\mathcal{T}_{\mathbf{v},\text{low}}$ as $\mathcal{T}_{\mathbf{v}}$ with the constraints in $\mathcal{T}_{\text{high}}$ removed. We now argue that a significant fraction of the constraints must be left over.

Claim 6.13. $|\mathcal{T}_{\mathbf{v},\text{low}}| \geq M/\gamma^{1+2/(1/2-1/p)}$.

Proof. Recall that we must have $|\mathcal{T}_{\mathbf{v}}| \geq M/\gamma^{2/(1/2-1/p)}$ by Lemma 6.5. We have

$$\begin{aligned} |\mathcal{T}_{\mathbf{v},\text{low}}| &\geq |\mathcal{T}_{\mathbf{v}}| - |\mathcal{T}_{\text{high}}| \\ &\geq M/\gamma^{2/(1/2-1/p)} - M/\gamma^{1+2/(1/2-1/p)} \\ &\geq 2 \cdot M/\gamma^{1+2/(1/2-1/p)} - M/\gamma^{1+2/(1/2-1/p)} \\ &\geq M/\gamma^{1+2/(1/2-1/p)} \end{aligned}$$

where we used that $\gamma \geq 2$ when M is sufficiently large. \square

Now we derive a contradiction by implicating the soundness of the PCP. For each variable x in the PCP, pick a tuple (x, σ) uniformly at random from the set of all tuples indicated by \mathbf{v} that involve x . (If \mathbf{v} has no tuples for the variable x , assign x to an arbitrary alphabet symbol.) Denote this random assignment as \mathcal{R} .


Claim 6.14. *The expected number of constraints satisfied by assignment \mathcal{R} is at least $|\mathcal{T}_{\mathbf{v},\text{low}}|/\gamma^{7q+2q/(1/2-1/p)}$.*

Proof. Observe that each constraint t in $\mathcal{T}_{\mathbf{v},\text{low}}$ is by definition *not* incident to a variable in $\mathcal{X}_{\text{high}}$, meaning that each of its variables has at most $\gamma^{7+2/(1/2-1/p)}$ distinct assignments indicated by \mathbf{v} . (By definition, each incident variable must also have at least one assignment indicated by \mathbf{v} .) Now even if \mathbf{v} only indicates one basis vector corresponding to t , the probability that the exact matching assignment is selected is $1/\gamma^{7q+2q/(1/2-1/p)}$. (Keep in mind that, by construction of \mathbf{G} , every row that \mathbf{v} can pick for a constraint will correspond to a *satisfying* assignment.) By linearity of expectation, the expected number of satisfied constraints is at least $|\mathcal{T}_{\mathbf{v},\text{low}}|/\gamma^{7q+2q/(1/2-1/p)}$. \square

Using our lower bound on $|\mathcal{T}_{\mathbf{v},\text{low}}|$, we can re-write this expectation as

$$\begin{aligned}
|\mathcal{T}_{\mathbf{v},\text{low}}|/\gamma^{7q+2q/(1/2-1/p)} &\geq M/(\gamma^{1+2/(1/2-1/p)} \cdot \gamma^{7q+2q/(1/2-1/p)}) \\
&\geq M/(\gamma^{7q+2q/(1/2-1/p)} \cdot \gamma^{7q+2q/(1/2-1/p)}) \\
&\geq M/\gamma^{20q/(1/2-1/p)} \\
&= M/(1/s)^{((1/2-1/p)/(25q)) \cdot (20q/(1/2-1/p))} \\
&= M \cdot s^{((1/2-1/p)/(25q)) \cdot (20q/(1/2-1/p))} \\
&= M \cdot s^{4/5},
\end{aligned}$$

where we used that $\gamma = (1/s)^{(1/2-1/p)/(25q)}$. By picking the best choice of randomness, we get a deterministic assignment which violates soundness. \square



An example PCP with $N = 9$, $M = 18$, $q = 3$, $d = 6$. Every column represents a variable, and every row represents a constraint. In this example, we have two variables in $\mathcal{X}_{\text{high}}$; their columns are marked in gray in $\mathcal{X}_{\text{high}}$. We also mark every constraint in $\mathcal{T}_{\text{high}}$ using gray. We can afford to “throw away” all constraints in $\mathcal{T}_{\text{high}}$ because $|\mathcal{T}_{\text{high}}| \ll |\mathcal{T}_{\mathbf{v}}|$. (This effect becomes more pronounced as the PCP gets larger.)

7 Putting it all Together

In this section we derive Corollaries 1.4 and 1.6. Recall the PCP theorem of [DFK⁺99], the Sliding Scale Conjecture [BGLR93], and the regularization technique of Hirahara and Moshkovitz [HM23].

Theorem 1.2 (Restated, from [DFK⁺99]). *For every constant $\varepsilon > 0$, there exists a PCP for SAT instances of size n that uses $r = O(\log n)$ random bits to make $q = O(1)$ queries to a proof over alphabet Σ , where the alphabet size is $|\Sigma| = 2^{\Theta(\log^{1-\varepsilon} n)}$ and the soundness parameter is $s = 2^{-\Theta(\log^{1-\varepsilon} n)}$. The PCP can be constructed deterministically in $n^{O(1)}$ time.*

Conjecture 7.1 (Sliding Scale Conjecture, from [BGLR93]). *There exists a PCP for SAT instances of size n that uses $r = O(\log n)$ random bits to make $q = O(1)$ queries to a proof over alphabet Σ , where the alphabet size is $|\Sigma| = n^{\Theta(1)}$ and the soundness parameter is $s = n^{-\Theta(1)}$. The PCP can be constructed deterministically in $n^{O(1)}$ time.*

Theorem 2.1 (Restated, from [HM23]). *Assume that we have a PCP verifier \mathcal{V} which uses r random bits to make q queries to a proof over alphabet Σ and has soundness s . Assume further that $s \leq \min(1/(3q), 1/|\Sigma|^c)$, where c is a constant with $0 < c \leq 1$. Then we can construct a new PCP verifier \mathcal{V}' from \mathcal{V} deterministically*

in $|\mathcal{V}|^{O(1)}$ time, such that the following holds. \mathcal{V}' uses $r + O(\log(1/s))$ random bits to make $q^{O(1)}$ queries to a proof over the same alphabet Σ and has soundness $s^{\Theta(1)}$. Additionally, every proof symbol is read for exactly $d = (q/s)^{\Theta(1)}$ choices of randomness.

Now recall our main theorem:

Theorem 4.3 (Restated). *Suppose there exists a PCP for SAT instances of size n that uses r random bits to read $q \leq O(r)$ locations in a proof over alphabet Σ and has soundness parameter s . Suppose further that every proof location is read with equal probability, that $(1/s)^{1/q} \geq r^{\omega(1)}$, and that the PCP can be constructed deterministically in time $T(n)$. Then for all constants $p > 2$ (and for $p = \infty$), there exists a $T(n) + (2^r |\Sigma|^q)^{O(1)}$ time deterministic reduction from SAT instances of size n to instances of γ -GapSVP $_p$ on lattices of dimension $O(2^r |\Sigma|^q)$, where $\gamma = (1/s)^{(1/2-1/p)/(25q)}$.*

We first prove the unconditional NP hardness result.

Corollary 1.4 (Restated). *For all constants $\varepsilon > 0$ and $p > 2$, γ -GapSVP $_p$ on n dimensional lattices is NP hard, where $\gamma = 2^{\log^{1-\varepsilon} n}$.*

Proof. Set $\varepsilon' = \varepsilon/2$. The PCP from Theorem 1.2 meets the preconditions of Theorem 2.1, so together they imply that there exists a PCP for SAT instances of size n that uses $r = O(\log n)$ random bits to read $q = O(1)$ locations in a proof over alphabet Σ , where the alphabet size is $|\Sigma| = 2^{\Theta(\log^{1-\varepsilon'} n)}$, the soundness parameter is $s = 2^{-\Theta(\log^{1-\varepsilon'} n)}$, and every proof symbol is read for exactly $d = 2^{\Theta(\log^{1-\varepsilon'} n)}$ choices of randomness. Clearly $q \leq O(r)$ in this case, and $(1/s)^{1/q} = 2^{\Theta(\log^{1-\varepsilon'} n)} \geq r^{\omega(1)} = (\log n)^{\omega(1)}$. We also have that $2^r |\Sigma|^q = n^{O(1)}$. Putting all of this together, we can apply Theorem 4.3 to show that γ -GapSVP $_p$ is NP hard on lattices of dimension $n^{O(1)}$, where $\gamma = 2^{\Theta(\log^{1-\varepsilon'} n)}$. Rescaling the dimension of the lattice and using ε in place of ε' yields the corollary. \square

The proof of the conditional NP hardness result is nearly the same.

Corollary 1.6 (Restated). *Assuming the Sliding Scale Conjecture, γ -GapSVP $_p$ on n dimensional lattices is NP hard for all constants $p > 2$ (and for $p = \infty$), where $\gamma = n^c$ for a constant $c > 0$ that depends on p .*

Proof. The PCP from Conjecture 7.1 meets the preconditions of Theorem 2.1, so together they imply that there exists a PCP for SAT instances of size n that uses $r = O(\log n)$ random bits to read $q = O(1)$ locations in a proof over alphabet Σ , where the alphabet size is $|\Sigma| = n^{\Theta(1)}$, the soundness parameter is $s = n^{-\Theta(1)}$, and every proof symbol is read for exactly $d = n^{\Theta(1)}$ choices of randomness. As before, $q \leq O(r)$, and $(1/s)^{1/q} = n^{\Theta(1)} \geq r^{\omega(1)} = (\log n)^{\omega(1)}$, and $2^r |\Sigma|^q = n^{O(1)}$. Putting all of this together, we can apply Theorem 4.3 to show that γ -GapSVP $_p$ on lattices of dimension $n^{O(1)}$ is NP hard for all constants $p > 2$, where $\gamma = n^{\Theta(1)}$. Rescaling the dimension of the lattice yields the corollary. \square

8 Acknowledgments

This research was supported in part from a Simons Investigator Award, DARPA SIEVE award, NTT Research, NSF grant 2333935, BSF grant 2022370, a Xerox Faculty Research Award, a Google Faculty Research Award, an Okawa Foundation Research Grant, and the Symantec Chair of Computer Science. This material is based upon work supported by the Defense Advanced Research Projects Agency through Award HR00112020024.

9 References

- [ABSS97] Sanjeev Arora, László Babai, Jacques Stern, and Z Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, 1997. ²
- [ACK⁺21] Divesh Aggarwal, Yanlin Chen, Rajendra Kumar, Zeyong Li, and Noah Stephens-Davidowitz. Dimension-preserving reductions between svp and cvp in different p-norms. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2444–2462. SIAM, 2021. ¹
- [Ajt98] Miklós Ajtai. The shortest vector problem in \mathbb{Z}^2 is np-hard for randomized reductions. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 10–19, 1998. ¹
- [AKS02] Miklós Ajtai, Ravi Kumar, and Dandapani Sivakumar. Sampling short lattice vectors and the closest lattice vector problem. In *Proceedings 17th IEEE annual conference on computational complexity*, pages 53–57. IEEE, 2002. ¹
- [ASD18] Divesh Aggarwal and Noah Stephens-Davidowitz. (gap/s) hardness of svp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 228–238, 2018. ¹
- [Ben23] Huck Bennett. The complexity of the shortest vector problem. *ACM SIGACT News*, 54(1):37–61, 2023. ¹
- [BGLR93] Mihir Bellare, Shafi Goldwasser, Carsten Lund, and Alexander Russell. Efficient probabilistically checkable proofs and applications to approximations. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 294–304, 1993. ^{3, 16}
- [BGLR24] Vijay Bhattiprolu, Venkatesan Guruswami, Euiwoong Lee, and Xuandi Ren. Inapproximability of finding sparse vectors in codes, subspaces, and lattices. *arXiv preprint arXiv:2410.02636*, 2024. ¹
- [BN09] Johannes Blömer and Stefanie Naewe. Sampling methods for shortest vectors, closest vectors and successive minima. *Theoretical Computer Science*, 410(18):1648–1665, 2009. ¹
- [BP22] Huck Bennett and Chris Peikert. Hardness of the (approximate) shortest vector problem: A simple proof via reed-solomon codes. *arXiv preprint arXiv:2202.07736*, 2022. ^{1, 2}
- [CN98] Jin-Yi Cai and Ajay Nerurkar. Approximating the svp to within a factor $(1-1/\dim/\sup/\text{spl } \epsilon/\text{spl } \epsilon)$ is np-hard under randomized conditions. In *Proceedings. Thirteenth Annual IEEE Conference on Computational Complexity (Formerly: Structure in Complexity Theory Conference)(Cat. No. 98CB36247)*, pages 46–55. IEEE, 1998. ¹
- [dB02] Rudi de Buda. Some optimal codes have structure. *IEEE Journal on Selected Areas in Communications*, 7(6):893–899, 2002. ¹
- [DFK⁺99] Irit Dinur, Eldar Fischer, Guy Kindler, Ran Raz, and Shmuel Safra. Pcp characterizations of np: Towards a polynomially-small error-probability. In *Proceedings of the thirty-first annual ACM symposium on Theory of Computing*, pages 29–40, 1999. ^{2, 16}

- [DHK15] Irit Dinur, Prahladh Harsha, and Guy Kindler. Polynomially low error pcps with polyloglog n queries via modular composition. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pages 267–276, 2015. ³
- [Din02] Irit Dinur. Approximating svp_∞ to within almost-polynomial factors is np-hard. *Theoretical Computer Science*, 285(1):55–71, 2002. ^{1,3}
- [DS14] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 624–633, 2014. ³
- [EV22] Friedrich Eisenbrand and Moritz Venzin. Approximate cvpp in time $20.802 n$. *Journal of Computer and System Sciences*, 124:129–139, 2022. ¹
- [For88] G David Forney. Coset codes. i. introduction and geometrical classification. *IEEE Transactions on Information Theory*, 34(5):1123–1151, 1988. ¹
- [FT87] András Frank and Éva Tardos. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica*, 7(1):49–65, 1987. ¹
- [HJ94] Roger A Horn and Charles R Johnson. *Topics in matrix analysis*. Cambridge university press, 1994. ⁶
- [HM23] Shuichi Hirahara and Dana Moshkovitz. Regularization of low error pcps and an application to mcspp. In *34th International Symposium on Algorithms and Computation (ISAAC 2023)*, pages 39–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023. ^{2,3,4,16,21}
- [Höl89] Otto Hölder. Ueber einen mittelwerthabsatz. *Nachrichten von der Königl. Gesellschaft der Wissenschaften und der Georg-Augusts-Universität zu Göttingen*, 1889:38–47, 1889. ⁴
- [HR07] Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 469–477, 2007. ¹
- [Kan87] Ravi Kannan. Minkowski’s convex body theorem and integer programming. *Mathematics of operations research*, 12(3):415–440, 1987. ¹
- [Kho03] Subhash Khot. Hardness of approximating the shortest vector problem in high l_p norms. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 290, 2003. ¹
- [Kho05] Subhash Khot. Hardness of approximating the shortest vector problem in lattices. *Journal of the ACM (JACM)*, 52(5):789–808, 2005. ¹
- [LJ83] Hendrik W Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of operations research*, 8(4):538–548, 1983. ¹
- [LLL82] AK Lenstra, HW Lenstra, and L Lovász. Factoring polynomials with rational coefficients. *Math. ann*, 261(4):515–534, 1982. ¹
- [Mic01] Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM journal on Computing*, 30(6):2008–2035, 2001. ¹
- [Mic12] Daniele Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Theory of Computing*, 8(1):487–512, 2012. ¹

- [Mos12] Dana Moshkovitz. The projection games conjecture and the np-hardness of $\ln n$ -approximating set-cover. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 276–287. Springer, 2012. ³
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM journal on computing*, 37(1):267–302, 2007. ¹
- [MR08] Dana Moshkovitz and Ran Raz. Two-query pcp with subconstant error. *Journal of the ACM (JACM)*, 57(5):1–29, 2008. ²¹
- [Muk22] Priyanka Mukhopadhyay. The projection games conjecture and the hardness of approximation of super-sat and related problems. *Journal of Computer and System Sciences*, 123:186–201, 2022. ³
- [NS01] Phong Q Nguyen and Jacques Stern. The two faces of lattices in cryptology. In *International Cryptography and Lattices Conference*, pages 146–180. Springer, 2001. ¹
- [P⁺16] Chris Peikert et al. A decade of lattice cryptography. *Foundations and trends® in theoretical computer science*, 10(4):283–424, 2016. ¹
- [PY88] Christos Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 229–234, 1988. ⁴
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009. ¹
- [RR06] Oded Regev and Ricky Rosen. Lattice problems and norm embeddings. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of Computing*, pages 447–456, 2006. ¹
- [Sch98] Alexander Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, 1998. ¹
- [V⁺12] Salil P Vadhan et al. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012. ²¹
- [vEB81] Peter van Emde Boas. Another np-complete problem and the complexity of computing short vectors in a lattice. *Technical Report, Department of Mathematics, University of Amsterdam*, 1981. ¹
- [Zam14] Ram Zamir. *Lattice coding for signals and networks: A structured coding approach to quantization, modulation, and multiuser information theory*. Cambridge University Press, 2014. ¹

A PCP Regularization

For the convenience of the reader, in this section we prove the PCP Regularization theorem due to Hirahara and Moshkovitz [HM23], which we made use of to derive our main result:

Theorem 2.1 (Restated). Assume that we have a PCP verifier \mathcal{V} which uses r random bits to make q queries to a proof over alphabet Σ and has soundness s . Assume further that $s \leq \min(1/(3q), 1/|\Sigma|^c)$, where c is a constant with $0 < c \leq 1$. Then we can construct a new PCP verifier \mathcal{V}' from \mathcal{V} deterministically in $|\mathcal{V}|^{O(1)}$ time, such that the following holds. \mathcal{V}' uses $r + O(\log(1/s))$ random bits to make $q^{O(1)}$ queries to a proof over the same alphabet Σ and has soundness $s^{\Theta(1)}$. Additionally, every proof symbol is read for exactly $d = (q/s)^{\Theta(1)}$ choices of randomness.

As in the body of our paper, we will adopt a CSP perspective (see Section 3.1), but the proof below corresponds exactly to the proof of [HM23]. At a high level the plan is to construct \mathcal{V}' by duplicating constraints and variables in a careful manner, and then connect constraints to the variable duplicates in accordance with an *explicit disperser*, which is defined as follows.

Definition A.1 ((δ, β) -Disperser). A (δ, β) -disperser is a bi-regular bipartite graph $G = (U, V, E)$ such that for all $V' \subset V$ of size at most $\beta|V|$, at most $\delta|U|$ of the vertices $u \in U$ have all of their neighbors contained within V' .

Hirahara and Moshkovitz's starting point for constructing an explicit disperser is the following theorem, which appears in a paper by Moshkovitz and Raz [MR08]:

Lemma A.2 ([MR08]). There is a constant $\kappa < 1$ and a function $T(\Delta) = \Theta(\Delta)$ such that, given two natural numbers n and Δ , one can find in time $\text{poly}(n\Delta)$ an undirected multigraph $G = (V, E)$ where $|V| = n$, every vertex is of the same degree $T(\Delta)$, and the adjacency matrix for G has its second largest eigenvalue being of magnitude at most $(T(\Delta))^\kappa$.

It's well known [V⁺12, HM23] that we can construct explicit dispersers by taking the set of all walks of a certain length in graphs where the second largest eigenvalue is of small magnitude. Using the graph from Lemma A.2 in such a construction, we can additionally ensure that the degree for both partitions of the expander is fixed independently of the partition sizes.

Corollary A.3 (Simplified version of corollary in [HM23]). There exists a function $f(w, \beta) = w \cdot (1/\beta)^{O(w)}$ such that the following holds. For all parameters $w \geq 1$, $\beta \in (0, 1)$, and $A \geq (1/\beta)^{q+1}$ there exists an explicit, $\text{poly}(A)$ time construction of a $(3\beta^w, \beta)$ -disperser $G = ([A], [B], E)$, where $B = \frac{wA}{f(w, \beta)}$. Every vertex in the partition indexed by $[A]$ has degree w and every vertex in the partition indexed by $[B]$ has degree $f(w, \beta)$.

Remark A.4. The corollary holds for $w = 1$ because in this case a matching between the two partitions constitutes a $(3\beta, \beta)$ -disperser. We also note that Hirahara and Moshkovitz use a slightly stronger version of the corollary, which guarantees an efficient construction of an $(e\beta^w, \beta)$ -disperser, where e is Euler's number. We use the constant 3 instead of e for ease of notation.

Now we demonstrate how to prove the PCP Regularization theorem by using an explicit disperser.

Proof of Theorem 2.1. The reduction from a PCP verifier \mathcal{V} to a regularized PCP verifier \mathcal{V}' goes as follows:

1. Duplicate each constraint in \mathcal{V} exactly $1/s$ times, which means that each variable appears in at least $1/s$ constraints. This does not affect completeness or soundness. The number of constraints increases from 2^r to $2^r/s$.

2. Do the following for each variable x in the PCP verifier

- (a) Let $d(x)$ be the number of constraints that depend on x , and invoke Corollary A.3 where $w = \lceil 6q/c \rceil$, $\beta = s^{1/(2q)}$, and $A = d(x)$. (Recall that q is the number of queries made by the PCP verifier \mathcal{V} and $0 < c \leq 1$ is a constant such that $s \leq 1/|\Sigma|^c$.) This gives us an explicit disperser $G_x = ([A], [B], E)$, where the degree of each vertex in the $[B]$ partition is exactly $f(w, \beta)$.
- (b) Create B duplicates x_1, \dots, x_B of the variable x .
- (c) Arbitrarily order the $A = d(x)$ constraints that depend on the original (un-duplicated) variable x .
- (d) For all $i \in [d(x)]$, modify the i th constraint that depends on x to instead depend on the $\lceil 6q/c \rceil$ copies of x that correspond with the $\lceil 6q/c \rceil$ neighbors of the i th vertex in partition $[A]$ in the graph G_x . The modified constraint now performs its original test, in addition to ensuring that all incident copies of variable x are equal to each other.

Observe that the construction runs in time $|\mathcal{V}|^{O(1)}$, and the new PCP verifier \mathcal{V}' immediately satisfies most of the conditions in Theorem 2.1. In particular, \mathcal{V}' uses $\log(2^r/s) = r + O(\log(1/s))$ bits of randomness, makes $q \cdot \lceil 6q/c \rceil = O(q^2)$ queries to a proof over the same alphabet as before, and every proof symbol is read for exactly $f(w, \beta) = q \cdot (1/s^{1/(2q)})^{O(q)} = (q/s)^{O(1)}$ choices of randomness. Completeness is also immediate: If \mathcal{V} has a satisfying assignment, then we can simply assign each copy of each variable in accordance with the assignment for \mathcal{V} , and every constraint of \mathcal{V}' is satisfied.

All that remains is to demonstrate soundness.

Claim A.5. *Suppose that the original PCP verifier \mathcal{V} was not satisfiable, i.e. every assignment to the variables satisfies at most an s fraction of the constraints. Then every assignment to the variables in \mathcal{V}' satisfies at most a $2\sqrt{s}$ fraction of the constraints.*

Proof. Suppose for contradiction that there exists an assignment to the variables of \mathcal{V}' that satisfies more than a $2\sqrt{s}$ fraction of the constraints. For each original variable x , consider the corresponding disperser $G_x = ([A], [B], E)$ and examine the set of variable copies x_1, \dots, x_B . Let $\Sigma_{x,\text{high}}$ be the $(1/s)^{1/(2q)}$ most popular assignments from Σ given to the copies of variable x , and let $\Sigma_{x,\text{low}} = \Sigma \setminus \Sigma_{x,\text{high}}$. We say that a constraint of \mathcal{V}' is bad if it is satisfied while simultaneously being incident to a copy of at least one variable x that was assigned to a symbol in $\Sigma_{x,\text{low}}$.

We first argue that, regardless of whether the original PCP verifier \mathcal{V} was satisfiable, at most an s fraction of the constraints in \mathcal{V}' are bad. First consider any variable x and any $\alpha \in \Sigma_{x,\text{low}}$. Notice that by definition of $\Sigma_{x,\text{low}}$, at most an $s^{1/(2q)}$ fraction of the copies of x can be assigned to α . This value of $s^{1/(2q)}$ is the parameter β we chose when constructing our explicit dispersers. So by the definition of disperser, we know that at most a $3\beta^w$ fraction of $\mathcal{X}_{x,\alpha} = \{\text{the constraints in } \mathcal{V}' \text{ that are incident to a copy of } x\}$ will depend only on copies assigned to α . These are the only constraints incident to at least one copy assigned to α that could possibly be satisfied, because our constraints check equality between all of their incident copies. Phrased more explicitly, at most a

$$3\beta^w \leq 3(s^{1/(2q)})^{\lceil 6q/c \rceil} \leq 3s^{3/c}$$

fraction of the constraints in \mathcal{V}' that are incident to a copy of x will simultaneously (i) be satisfied and (ii) have the copy assigned to α . Union bounding over the at most $|\Sigma|$ choices for $\alpha \in \Sigma_{x,\text{low}}$, and using the assumptions that $s \leq \min(1/(3q), 1/|\Sigma|^c)$ and $0 < c \leq 1$, at most a

$$3s^{3/c} \cdot |\Sigma| \leq 3s^{2/c} \leq 3s^2 \leq s/q$$

fraction of the constraints in \mathcal{V}' incident to a copy of x will simultaneously (i) be satisfied and (ii) have the copy assigned to any value in $\Sigma_{x,\text{low}}$. Because each constraint is incident to copies corresponding to q different original variables, and each set of copies could cause the constraint to be bad, a final union bound over the arity q gives the desired upper bound, that at most an s fraction of the constraints in \mathcal{V}' are bad.

Because we assumed that we have an assignment which satisfies more than a $2\sqrt{s}$ fraction of the constraints in \mathcal{V}' , there must be more than a $2\sqrt{s} - s \geq \sqrt{s}$ fraction of constraints in \mathcal{V}' that are satisfied and not bad. In other words, more than a \sqrt{s} fraction of the constraints in \mathcal{V}' are satisfied while simultaneously being incident only to copies of variables x that are assigned to symbols in $\Sigma_{x,\text{high}}$. By definition, for all variables x we have that $|\Sigma_{x,\text{high}}| \leq (1/s)^{1/(2q)}$. If we randomly pick a value in $\Sigma_{x,\text{high}}$ for each variable x and take this as an assignment for the original PCP verifier \mathcal{V} , we expect to satisfy strictly more than a $\sqrt{s} \cdot (s^{1/(2q)})^q = \sqrt{s} \cdot s^{1/2} = s$ fraction of the constraints in \mathcal{V} . Taking the best choice of randomness gives a deterministic assignment which contradicts the soundness of \mathcal{V} . \square

With both completeness and soundness proven, the theorem follows. \square

B Proof of Corollary 2.3

Recall Hölder's Inequality.

Theorem 2.2 (Restated). *Let $\mathbf{u}, \mathbf{v} \in \mathbb{Z}^n$, and let $p, q \geq 1$ be parameters satisfying $1/p + 1/q = 1$. Then*

$$\|\mathbf{u} \odot \mathbf{v}\|_1 \leq \|\mathbf{u}\|_p \|\mathbf{v}\|_q,$$

where \odot denotes the component-wise product.

We now prove Corollary 2.3

Corollary 2.3 (Restated). *For all $\mathbf{w} \in \mathbb{Z}^n$ and $p \geq 2$ (including $p = \infty$),*

$$\|\mathbf{w}\|_p \geq n^{1/p-1/2} \|\mathbf{w}\|_2.$$

Proof. We start with the special cases $p = 2$ and $p = \infty$. Notice that when $p = 2$, we have $n^{1/p-1/2} = 1$, so the inequality holds immediately. Now consider $p = \infty$, and assume for contradiction that $\|\mathbf{w}\|_\infty < n^{-1/2} \|\mathbf{w}\|_2$. Then for all $i \in [n]$, we have $|w_i| < n^{-1/2} \|\mathbf{w}\|_2$, and the ℓ_2 norm is upper bounded as

$$\begin{aligned} \|\mathbf{w}\|_2 &< \left(\sum_{i=1}^n n^{-1} (\|\mathbf{w}\|_2)^2 \right)^{1/2} \\ &= ((\|\mathbf{w}\|_2)^2)^{1/2} \\ &= \|\mathbf{w}\|_2. \end{aligned}$$

Clearly $\|\mathbf{w}\|_2 < \|\mathbf{w}\|_2$ is a contradiction.

Now consider the case that $p > 2$ is finite. Set $\mathbf{v} \in \mathbb{Z}^n$ to be $\mathbf{w} \odot \mathbf{w}$, i.e. we square each component. Define $\mathbf{u} = \mathbf{1}^n$, and set $a = p/(p-2) > 1$ and $b = p/2 > 1$. Because $1/a + 1/b = (p-2)/p + 2/p = 1$, an application of Theorem 2.2 gives

$$\|\mathbf{u} \odot \mathbf{v}\|_1 \leq \|\mathbf{u}\|_a \|\mathbf{v}\|_b.$$

By the definition of p norm we can write this as

$$\sum_{i=1}^n |\mathbf{u}_i \mathbf{v}_i| \leq \left(\sum_{i=1}^n |\mathbf{u}_i^a| \right)^{1/a} \left(\sum_{i=1}^n |\mathbf{v}_i^b| \right)^{1/b}.$$

Now because \mathbf{v} is the component-wise square of \mathbf{w} , and \mathbf{u} is the all-ones vector, we have

$$\begin{aligned} \sum_{i=1}^n |\mathbf{u}_i \mathbf{v}_i| &\leq \left(\sum_{i=1}^n |\mathbf{u}_i^a| \right)^{1/a} \left(\sum_{i=1}^n |\mathbf{v}_i^b| \right)^{1/b} \\ \sum_{i=1}^n |\mathbf{w}_i^2| &\leq \left(\sum_{i=1}^n 1 \right)^{1/a} \left(\sum_{i=1}^n |\mathbf{w}_i^{2b}| \right)^{1/b} \\ \sum_{i=1}^n |\mathbf{w}_i^2| &\leq n^{(p-2)/p} \left(\sum_{i=1}^n |\mathbf{w}_i^p| \right)^{2/p} \\ (\|\mathbf{w}\|_2)^2 &\leq n^{(p-2)/p} (\|\mathbf{w}\|_p)^2. \end{aligned}$$

Re-arranging and taking square roots yields

$$\begin{aligned} n^{(p-2)/p} (\|\mathbf{w}\|_p)^2 &\geq (\|\mathbf{w}\|_2)^2 \\ (\|\mathbf{w}\|_p)^2 &\geq n^{(2-p)/p} (\|\mathbf{w}\|_2)^2 \\ \|\mathbf{w}\|_p &\geq n^{1/p-1/2} \|\mathbf{w}\|_2. \end{aligned}$$

□