On the Quantitative Hardness of CVP

Huck Bennett[*]
huckbennett@gmail.com

Alexander Golovnev[†‡]
alexgolovnev@gmail.com

Noah Stephens-Davidowitz[*§]
noahsd@gmail.com

Abstract

For odd integers $p \geq 1$ (and $p = \infty$), we show that the Closest Vector Problem in the ℓ_p norm (CVP_p) over rank n lattices cannot be solved in $2^{(1-\varepsilon)n}$ time for any constant $\varepsilon > 0$ unless the Strong Exponential Time Hypothesis (SETH) fails. We then extend this result to "almost all" values of $p \geq 1$, not including the even integers. This comes tantalizingly close to settling the quantitative time complexity of the important special case of CVP_2 (i.e., CVP in the Euclidean norm), for which a $2^{n+o(n)}$-time algorithm is known. In particular, our result applies for any $p = p(n) \neq 2$ that approaches 2 as $n \to \infty$.

We also show a similar SETH-hardness result for SVP_∞; hardness of approximating CVP_p to within some constant factor under the so-called Gap-ETH assumption; and other quantitative hardness results for CVP_p and CVPP_p for any $1 \leq p < \infty$ under different assumptions.

1   Introduction

A lattice $\mathcal{L}$ is the set of all integer combinations of linearly independent basis vectors $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n \in \mathbb{R}^d$,

$$\mathcal{L} = \mathcal{L}(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n) := \Big\{ \sum_{i=1}^{n} z_i \boldsymbol{b}_i \ : \ z_i \in \mathbb{Z} \Big\} .$$

We call n the *rank* of the lattice $\mathcal{L}$ and d the *dimension* or the *ambient dimension*.

The two most important computational problems on lattices are the Shortest Vector Problem (SVP) and the Closest Vector Problem (CVP). Given a basis for a lattice $\mathcal{L} \subset \mathbb{R}^d$, SVP asks us to compute the minimal length of a non-zero vector in $\mathcal{L}$, and CVP asks us to compute the distance from some target point $\boldsymbol{t} \in \mathbb{R}^d$ to the lattice. Typically, we define length and distance in terms of the ℓ_p norm for some $1 \leq p \leq \infty$, given by

$$\|\boldsymbol{x}\|_p := (|x_1|^p + |x_2|^p + \cdots + |x_d|^p)^{1/p}$$

for finite p and

$$\|\boldsymbol{x}\|_\infty := \max_{1 \leq i \leq d} |x_i| .$$

In particular, the ℓ_2 norm is the familiar Euclidean norm, and it is by far the best studied in this context. We write SVP_p and CVP_p for the respective problems in the ℓ_p norm. CVP is known to be at least as hard as SVP (in any norm, under an efficient reduction that preserves the rank and approximation factor) [GMSS99] and appears to be strictly harder.

Starting with the breakthrough work of Lenstra, Lenstra, and Lovász in 1982 [LLL82], algorithms for solving these problems in both their exact and approximate forms have found innumerable applications, including factoring polynomials over the rationals [LLL82], integer programming [Len83, Kan87, DPV11], cryptanalysis [Sha84, Odl90, JS98, NS01], etc. More recently, many cryptographic primitives have been constructed whose security is based on the *worst-case* hardness of these or closely related lattice problems [Ajt04, Reg09, GPV08, Pei08, Pei16]. Given the obvious importance of these problems, their complexity is quite well studied. Below, we survey some of these results. We focus on algorithms for the exact and near-exact problems since these are most relevant to our work and because the best known algorithms for the approximate variants of these problems typically use algorithms for the exact problems as subroutines [Sch87, GN08, MW16]. (Many of the results described below are also summarized in Table 1.)

1.1   Algorithms for SVP and CVP

The AKS algorithm and its descendants.   The current fastest known algorithms for solving SVP_p all use the celebrated randomized sieving technique due to Ajtai, Kumar, and Sivakumar [AKS01]. The original algorithm from [AKS01] was the first $2^{O(n)}$-time algorithm for SVP, and it worked for both $p = 2$ and $p = \infty$.

In the $p = 2$ case, a sequence of works improved upon the constant in the exponent [NV08, PS09, MV10, LWXZ11], and the current fastest running time of an algorithm that provably solves SVP_2 exactly is $2^{n+o(n)}$ [ADRS15].[1] While progress has slowed, this seems unlikely to be the end of the story. Indeed, there are heuristic sieving algorithms that run in time $(3/2)^{n/2+o(n)}$ [NV08,

---

[1] The algorithm in [ADRS15] is quite a bit different than the other algorithms in this class, but it can still be thought of as a sieving algorithm.

WLTB11, Laa15, BDGL16], and there is some reason to believe that the provably correct [ADRS15] algorithm can be improved. In particular, there is a provably correct $2^{n/2+o(n)}$-time algorithm that approximates SVP_2 up to a small constant approximation factor [ADRS15].

A different line of work extended the randomized sieving approach of [AKS01] to obtain $2^{O(n)}$-time algorithms for SVP in additional norms. In particular, Blömer and Naewe extended it to all ℓ_p norms [BN09]. Subsequent work extended this further, first to arbitrary symmetric norms [AJ08] and then to the "near-symmetric norms" that arise in integer programming [Dad12].

Finally, a third line of work extended the [AKS01] approach to approximate CVP. Ajtai, Kumar, and Sivakumar themselves showed a $2^{O(n)}$-time algorithm for approximating CVP_2 to within any constant approximation factor strictly greater than one [AKS02]. Blömer and Naewe obtained the same result for all ℓ_p norms [BN09], and Dadush extended it further to arbitrary symmetric norms and again to "near-symmetric norms" [Dad12]. We stress, however, that none of these results apply to exact CVP, and indeed, there are some barriers to extending these algorithms to exact CVP. (See, e.g., [ADS15].)

Exact algorithms for CVP.   *Exact* CVP appears to be a much more subtle problem than exact SVP.[2] Indeed, progress on exact CVP has been much slower than the progress on exact SVP. Over a decade after [AKS01], Micciancio and Voulgaris presented the first $2^{O(n)}$-time algorithm for exact CVP_2 [MV13], using elegant new techniques built upon the approach of Sommer, Feder, and Shalvi [SFS09]. Specifically, they achieved a running time of $4^{n+o(n)}$, and subsequent work even showed a running time of $2^{n+o(n)}$ for CVP_2 with Preprocessing (in which the algorithm is allowed access to arbitrary advice that depends on the lattice but not the target vector; see Section 2.1) [BD15]. Later, [ADS15] showed a $2^{n+o(n)}$-time algorithm for CVP_2, so that the current best known asymptotic running time is actually the same for SVP_2 and CVP_2.

However, for $p \neq 2$, progress for exact CVP_p has been minimal. Indeed, the fastest known algorithms for exact CVP_p with $p \neq 2$ are still the $n^{O(n)}$-time enumeration algorithms first developed by Kannan in 1987 [Kan87, DPV11, MW15]. Both algorithms for exact CVP_2 mentioned in the previous paragraph use many special properties of the ℓ_2 norm, and it seems that substantial new ideas would be required to extend them to arbitrary ℓ_p norms.

1.2   Hardness of SVP and CVP

Van Emde Boas showed the NP-hardness of CVP_p for any p and SVP_∞ in 1981 [vEB81]. Extending this to SVP_p for finite p was a major open problem until it was proven (via a randomized reduction) for all $1 \leq p \leq \infty$ by Ajtai in 1998 [Ajt98]. There has since been much follow-up work, showing the hardness of these problems for progressively larger approximation factors, culminating in NP-hardness of approximating CVP_p up to a factor of $n^{c/\log \log n}$ for some constant $c > 0$ [ABSS93, DKRS03] and hardness of SVP_p with the same approximation factor under plausible complexity-theoretic assumptions [CN98, Mic01b, Kho05, HR12]. These results are nearly the best possible under plausible assumptions, since approximating either problem up to a factor of $\sqrt{n}$ is known to be in $NP \cap coNP$ [GG00, AR05, Pei08].

---

[2]In particular, there can be arbitrarily many lattice points that are approximate closest vectors, which makes sieving techniques seemingly useless for solving exact CVP. (See, e.g., [ADS15] for a discussion of this issue.) We note, however, that hardness results (including ours) tend to produce CVP instances with a bounded number of approximate closest vectors (e.g., $2^{O(n)}$).

However, such results only rule out the possibility of polynomial-time algorithms (under reasonable complexity-theoretic assumptions). They say very little about the *quantitative* hardness of these problems for a fixed lattice rank n.[3]

This state of affairs is quite frustrating for two reasons. First, in the specific case of CVP_2, algorithmic progress has reached an apparent barrier. In particular, both known techniques for solving exact CVP_2 in singly exponential time are fundamentally unable to produce algorithms whose running time is asymptotically better than the current best of $2^{n+o(n)}$ [MV13, ADS15].[4] Second, some lattice-based cryptographic constructions are close to deployment [ADPS16, BCD+16, NIS16]. In order to be practically secure, these constructions require the quantitative hardness of certain lattice problems, and so their designers rely on quantitative hardness assumptions [APS15]. If, for example, there existed a $2^{n/20}$-time algorithm for SVP_p or CVP_p, then these cryptographic schemes would be insecure in practice.

We therefore move in a different direction. Rather than trying to extend non-quantitative hardness results to larger approximation factors, we show quantitative hardness results for exact (or nearly exact) problems. To do this, we use the tools of *fine-grained complexity*.

1.3 Fine-grained complexity

Impagliazzo and Paturi [IP99] introduced the *Exponential Time Hypothesis* (ETH) and the *Strong Exponential Time Hypothesis* (SETH) to help understand the precise hardness of k-SAT. Informally, ETH asserts that 3-SAT takes $2^{\Omega(n)}$-time to solve in the worst case, and SETH asserts that k-SAT takes essentially 2^n-time to solve for unbounded k. I.e., SETH asserts that brute-force search is essentially optimal for solving k-SAT for large k.

Recently, the study of fine-grained complexity has leveraged ETH, SETH, and several other assumptions to prove quantitative hardness results about a wide range of problems. These include both problems in P (see, e.g., [CLR+14, BI15, ABW15] and the survey by Vassilevska Williams [Wil15]), and NP-hard problems (see, e.g., [PW10, CDL+12, CFK+15]). Although these results are all conditional, they help to explain *why* making further algorithmic progress on these problems is difficult—and suggest that it might be impossible. Namely, any non-trivial algorithmic improvement would disprove a very well-studied hypothesis.

One proves quantitative hardness results using *fine-grained* reductions (see [Wil15] for a formal definition). For example, there is an efficient mapping from k-SAT formulas on n variables to Hitting Set instances with universes of n elements [CDL+12]. This reduction is fine-grained in the sense that for any constant $\varepsilon > 0$, a $2^{(1-\varepsilon)n}$-time algorithm for Hitting Set implies a $2^{(1-\varepsilon)n}$-time algorithm for k-SAT, breaking SETH.

Despite extensive effort, no faster-than-2^n-time algorithm for k-SAT with unbounded k has been found. Nevertheless, there is no consensus on whether SETH is true or not, and recently, Williams [Wil16] refuted a very strong variant of SETH. This makes it desirable to base quantitative hardness results on weaker assumptions when possible, and indeed our main result holds even

---

[3] One can derive certain quantitative hardness results from known hardness proofs, but in most cases the resulting lower bounds are quite weak. The only true quantitative hardness results known prior to this work were a folklore ETH-hardness result for CVP and an unpublished result due to Samuel Yeom, showing that CVP cannot be solved in time $2^{10^{-4}n}$ under plausible complexity-theoretic assumptions [Vai15]. (In Section 6.2, we present a similar proof of a stronger statement.)

[4] Both techniques require short vectors in each of the 2^n cosets of $\mathcal{L}$ mod $2\mathcal{L}$ (though for apparently different reasons).

Problem	Upper Bound	Lower Bounds				Notes
		SETH	Max-2-SAT	ETH	Gap-ETH	
CVP_p	$n^{O(n)}\ (2^{O(n)})$	2^n	$2^{\omega n/3}$	$2^{\Omega(n)}$	$2^{\Omega(n)}*$	"almost all" $p \notin 2\mathbb{Z}$
CVP_2	2^n	—	$2^{\omega n/3}$	$2^{\Omega(n)}$	$2^{\Omega(n)}*$	
$\mathrm{CVP}_\infty/\mathrm{SVP}_\infty$	$2^{O(n)}$	$2^{n}*$	—	$2^{\Omega(n)}$	$2^{\Omega(n)}*$	
CVPP_p	$n^{O(n)}\ (2^{O(n)})$	—	$2^{\Omega(\sqrt{n})}$	$2^{\Omega(\sqrt{n})}$	—	

Table 1: Summary of known quantitative upper and lower bounds, with new results in blue. Upper bounds in parentheses hold for any constant approximation factor strictly greater than one, and lower bounds with a $*$ apply for some constant approximation factor strictly greater than one. ω is the matrix multiplication exponent, satisfying $2 \leq \omega < 2.373$. We have suppressed smaller factors.

assuming a weaker variant of SETH based on the hardness of Weighted Max-k-SAT (except for the case of $p = \infty$).

1.4 Our contribution

We now enumerate our results. See also Table 1.

SETH-hardness of CVP_p**.**   Our main result is the SETH-hardness of CVP_p for any odd integer $p \geq 1$ and $p = \infty$ (and SVP_∞). Formally, we prove the following. (See Sections 3 and 4 for finite p and Section 6.3 for $p = \infty$.)

Theorem 1.1. *For any constant integer $k \geq 2$ and any odd integer $p \geq 1$ or $p = \infty$, there is an efficient reduction from k-SAT with n variables and m clauses to CVP_p (or SVP_∞) on a lattice of rank n (with ambient dimension $n + O(m)$).*

   In particular, there is no $2^{(1-\varepsilon)n}$-time algorithm for CVP_p for any odd integer $p \geq 1$ or $p = \infty$ (or SVP_∞) and any constant $\varepsilon > 0$ unless SETH is false.

   Unfortunately, we are unable to extend this result to even integers p, and in particular, to the important special case of $p = 2$. In fact, this is inherent, as we show that our approach necessarily fails for even integers $p \leq k - 1$. In spite of this, we actually prove the following result that generalizes Theorem 1.1 to "almost all" $p \geq 1$ (including non-integer p).

Theorem 1.2. *For any constant integer $k \geq 2$, there is an efficient reduction from k-SAT with n variables and m clauses to CVP_p on a lattice of rank n (with ambient dimension $n + O(m)$) for any $p \geq 1$ such that*

   1. p is an odd integer or $p = \infty$;

   2. $p \notin S_k$, where S_k is some finite set (containing all even integers $p \leq k - 1$); or

   3. $p = p_0 + \delta(n)$ for any $p_0 \geq 1$ and any $\delta(n) \neq 0$ that converges to zero as $n \to \infty$.

   In particular, if SETH holds then for any constant $\varepsilon > 0$, there is no $2^{(1-\varepsilon)n}$-time algorithm for CVP_p for any $p \geq 1$ such that

   1. p is an odd integer or $p = \infty$;

4

2. $p \notin S_k$ for some sufficiently large k (depending on ε); or

3. $p = p_0 + \delta(n)$.

Notice that this lower bound (Theorem 1.2) comes tantalizingly close to resolving the quantitative complexity of CVP_2. In particular, we obtain a 2^n-time lower bound on $CVP_{2+\delta}$ for any $0 \neq \delta(n) = o(1)$, and the fastest algorithm for CVP_2 run in time $2^{n+o(n)}$. But, formally, Theorems 1.1 and 1.2 say nothing about CVP_2. (Indeed, there is at least some reason to believe that CVP_2 is easier than CVP_p for $p \neq 2$ [RR06].)

We note that our reductions actually work for Weighted Max-k-SAT for all finite $p \neq \infty$, so that our hardness result holds under a weaker assumption than SETH, namely, the corresponding hypothesis for Weighted Max-k-SAT.

Finally, we note that in the special case of $p = \infty$, our reduction works even for approximate CVP_∞, or even approximate SVP_∞, with an approximation factor of $\gamma := 1 + 2/(k - 1)$. In particular, γ is constant for fixed k. This implies that for every constant $\varepsilon > 0$, there is a constant $\gamma_\varepsilon > 1$ such that no $2^{(1-\varepsilon)n}$-time algorithm approximates SVP_∞ or CVP_∞ to within a factor of γ_ε unless SETH fails.

Quantitative hardness of approximate CVP.   As we discussed above, many $2^{O(n)}$-time algorithms for CVP_p only work for γ-approximate CVP_p for constant approximation factors $\gamma > 1$. However, the reduction described above only works for *exact* CVP_p (except when $p = \infty$).[5]

So, it would be preferable to show hardness for some constant approximation factor $\gamma > 1$. One way to show such a hardness result is via a fine-grained reduction from the problem of approximating Max-k-SAT to within a constant factor. Indeed, in the $k = 2$ case, we show that such a reduction exists, so that there is no $2^{o(n)}$-time algorithm for approximating CVP_p to within some constant factor unless a $2^{o(n)}$-time algorithm exists for approximating Max-2-SAT. We also note that a $2^{o(n)}$-time algorithm for approximating Max-2-SAT to within a constant factor would imply one for Max-3-SAT as well. (See Proposition 2.12.)

We present this result informally here (without worrying about specific parameters and the exact definition of approximate Max-2-SAT). See Section 5 for the formal statement.

Theorem 1.3. *There is an efficient reduction from approximating Max-2-SAT with n variables and m clauses to within a constant factor to approximating CVP_p to within a constant factor on a lattice of rank n (with ambient dimension $n + O(m)$) for any finite $p \geq 1$.*

Quantitative hardness of CVP **with Preprocessing.**   CVP with Preprocessing (CVPP) is the variant of CVP in which we are allowed arbitrary advice that depends on the lattice, but not the target vector. CVPP and its variants have potential applications in both cryptography (e.g., [GPV08]) and cryptanalysis. And, an algorithm for $CVPP_2$ is used as a subroutine in the celebrated Micciancio-Voulgaris algorithm for CVP_2 [MV13, BD15]. The complexity of $CVPP_p$ is well studied, with both hardness of approximation results [Mic01a, FM04, Reg04, AKKV11, KPV14], and efficient approximation algorithms [AR05, DRS14].

We prove the following quantitative hardness result for $CVPP_p$. (See Section 6.1.)

---

[5]One can likely show that our "exact" reductions actually work for γ-approximate CVP_p with some approximation factor $\gamma = 1 + o(1)$. But, this is not very interesting because standard techniques for "boosting" the approximation factor are useless for us. (They increase the rank far too much.)

5

Theorem 1.4. *For any $1 \leq p < \infty$, there is no $2^{o(\sqrt{n})}$-time algorithm for* CVPP *unless there is a (non-uniform) $2^{o(n)}$-time algorithm for Max-2-SAT. In particular, no such algorithm exists unless (non-uniform) ETH fails.*

Additional quantitative hardness results for CVP_p**.**   We also observe the following weaker hardness result for CVP_p for any $1 \leq p < \infty$ based on different assumptions. The ETH-hardness of CVP_p was already known in folklore, and even written down by Samuel Yeom in unpublished work [Vai15]. We present a slightly stronger theorem than what was previously known, showing a reduction from Max-2-SAT on n variables to CVP_p on a lattice of rank n. (Prior to this work, we were only aware of reductions from 3-SAT on n variables to CVP_p on a lattice of rank Cn for some very large constant $C > 1000$.)

Theorem 1.5. *For any $1 \leq p < \infty$, there is an efficient reduction from Max-2-SAT with n variables to CVP_p on a lattice of rank n (and dimension $n + m$, where m is the number of clauses).*
    In particular, for any constant $c > 0$, there is no $(\mathrm{poly}(n) \cdot 2^{cn})$-time algorithm for CVP_p unless there is a 2^{cn}-time algorithm for Max-2-SAT, and there is no $2^{o(n)}$-time algorithm for CVP_p unless ETH fails.

    The fastest known algorithm for the Max-2-SAT problem is the $\mathrm{poly}(n) \cdot 2^{\omega n/3}$-time algorithm due to Williams [Wil05], where $2 \leq \omega < 2.373$ is the matrix multiplication exponent [Wil12, LG14]. This implies that a faster than $2^{\omega n/3}$-time algorithm for CVP_p (and CVP_2 in particular) would yield a faster algorithm for Max-2-SAT. (See, e.g., [Woe08] Open Problem 4.7 and the preceding discussion.)

1.5   Techniques

Max-2-SAT.   We first show a straightforward reduction from Max-2-SAT to CVP_p for any $1 \leq p < \infty$. I.e., we prove Theorem 1.5. This simple reduction will introduce some of the high-level ideas needed for our more difficult reductions.
    Given a Max-2-SAT instance Φ with n variables and m clauses, we construct the lattice basis

$$B := \begin{pmatrix} \bar{\Phi} \\ 2\alpha I_n \end{pmatrix} , \tag{1}$$

where $\alpha > 0$ is some very large number and $\bar{\Phi} \in \mathbb{R}^{m \times n}$ is given by

$$\bar{\Phi}_{i,j} := \begin{cases} 2 & \text{if the } i\text{th clause contains } x_j, \\ -2 & \text{if the } i\text{th clause contains } \neg x_j, \\ 0 & \text{otherwise .} \end{cases} \tag{2}$$

I.e., the rows of $\bar{\Phi}$ correspond to clauses and the columns correspond to variables. Each entry encodes whether the relevant variable is included in the relevant clause unnegated, negated, or not at all, using 2, -2, and 0 respectively. (We assume without loss of generality that no clause contains repeated literals or a literal and its negation simultaneously.) The target $\boldsymbol{t} \in \mathbb{R}^{m+n}$ is given by

$$\boldsymbol{t} := (t_1, t_2, \ldots, t_m, \alpha, \alpha, \ldots, \alpha)^T , \tag{3}$$

where

$$t_i := 3 - 2\eta_i , \tag{4}$$

6

where η_i is the number of negated variables in the ith clause.

Notice that the copy of $2\alpha I_n$ at the bottom of B together with the sequence of α's in the last coordinates of $\boldsymbol{t}$ guarantee that any lattice vector $B\boldsymbol{z}$ with $\boldsymbol{z} \in \mathbb{Z}^n$ is at distance at least $\alpha n^{1/p}$ away from $\boldsymbol{t}$. Furthermore, if $\boldsymbol{z} \notin \{0,1\}^n$, then this distance increases to at least $\alpha(n-1+3^p)^{1/p}$. This is a standard gadget, which will allow us to ignore the case $\boldsymbol{z} \notin \{0,1\}^n$ (as long as α is large enough). I.e., we can view $\boldsymbol{z}$ as an assignment to the n variables of Φ.

Now, suppose $\boldsymbol{z}$ does not satisfy the ith clause. Then, notice that the ith coordinate of $B\boldsymbol{z}$ will be exactly $-2\eta_i$, so that $(B\boldsymbol{z} - \boldsymbol{t})_i = 0 - 3 = -3$. If, on the other hand, exactly one literal in the ith clause is satisfied, then the ith coordinate of $B\boldsymbol{z}$ will be $2 - 2\eta_i$, so that $(B\boldsymbol{z} - \boldsymbol{t})_i = 2 - 3 = -1$. Finally, if both literals are satisfied, then the ith coordinate will be $4 - 2\eta_i$, so that $(B\boldsymbol{z} - \boldsymbol{t})_i = 4 - 3 = 1$. In particular, if the ith clause is not satisfied, then $|(B\boldsymbol{z} - \boldsymbol{t})_i| = 3$. Otherwise, $|(B\boldsymbol{z} - \boldsymbol{t})_i| = 1$.

It follows that the distance to the target is exactly $\mathrm{dist}_p(\boldsymbol{t}, \mathcal{L})^p = \alpha^p n + S + 3^p(m - S) = \alpha^p n - (3^p - 1)S + 3^p m$, where S is the maximal number of satisfied clauses of Φ. So, the distance $\mathrm{dist}_p(\boldsymbol{t}, \mathcal{L})$ tells us exactly the maximum number of satisfiable clauses, which is what we needed.

Difficulties extending this to k-SAT. The above reduction relied on one very important fact: that $|4 - 3| = |2 - 3| < |0 - 3|$. In particular, a 2-SAT clause can be satisfied in two different ways; either one variable is satisfied or two variables are satisfied. We designed our CVP instance above so that the ith coordinate of $B\boldsymbol{z} - \boldsymbol{t}$ is $4 - 3$ if two literals in the ith clause are satisfied by $\boldsymbol{z} \in \{0,1\}^n$, $2 - 3$ if one literal is satisfied, and $0 - 3$ if the clause is unsatisfied. Since $|4 - 3| = |2 - 3|$, the "contribution" of this ith coordinate to the distance $\|B\boldsymbol{z} - \boldsymbol{t}\|_p^p$ is the same for any satisfied clause. Since $|0 - 3| > |4 - 3|$, the contribution to the ith coordinate is larger for unsatisfied clauses than satisfied clauses.

Suppose we tried the same construction for a k-SAT instance. I.e., suppose we take $\bar{\Phi} \in \mathbb{R}^{m \times n}$ to encode the literals in each clause as in Eq. (2) and construct our lattice basis B as in Eq. (1) and target $\boldsymbol{t}$ as in Eq. (3), perhaps with the number 3 in the definition of $\boldsymbol{t}$ replaced by an arbitrary $t^* \in \mathbb{R}$. Then, the ith coordinate of $B\boldsymbol{z} - \boldsymbol{t}$ would be $2S_i - t^*$, where S_i is the number of literals satisfied in the ith clause.

No matter how cleverly we choose $t^* \in \mathbb{R}$, some satisfied clauses will contribute more to the distance than others as long as $k \geq 3$. I.e., there will always be some "imbalance" in this contribution. As a result, we will not be able to distinguish between, e.g., an assignment that satisfies all clauses but has S_i far from $t^*/2$ for all i and an assignment that satisfies fewer clauses but has $S_i \approx t^*/2$ whenever i corresponds to a satisfying clause.

In short, for $k \geq 3$, we run into trouble because satisfying assignments to a clause may satisfy anywhere between 1 and k literals, but k distinct numbers obviously cannot all be equidistant from some number t^*. (See Section 6.2 for a simple way to get around this issue by adding to the rank of the lattice. Below, we show a more technical way to do this without adding to the rank of the lattice, which allows us to prove SETH-hardness.)

A solution via isolating parallelepipeds. To get around the issue described above for $k \geq 3$, we first observe that, while many distinct *numbers* cannot all be equidistant from some *number* t^*, it is trivial to find many distinct *vectors* in $\mathbb{R}^{d^*}$ that are equidistant from some *vector* $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$.

We therefore consider modifying the reduction from above by replacing the scalar ± 2 values in our matrix $\bar{\Phi}$ with vectors in $\mathbb{R}^{d^*}$ for some d^*. In particular, for some vectors $V = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k) \in \mathbb{R}^{d^* \times k}$,

we define $\bar{\Phi} \in \mathbb{R}^{d^* m \times n}$ as

$$\bar{\Phi}_{i,j} := \begin{cases} \boldsymbol{v}_s & \text{if } x_j \text{ is the } s\text{th literal in the } i\text{th clause,} \\ -\boldsymbol{v}_s & \text{if } \neg x_j \text{ is the } s\text{th literal in the } i\text{th clause,} \\ \boldsymbol{0}_d & \text{otherwise ,} \end{cases} \quad (5)$$

where we have abused notation and taken $\bar{\Phi}_{i,j}$ to be a column vector in d^* dimensions. By defining $\boldsymbol{t} \in \mathbb{R}^{d^* m + n}$ appropriately,[6] we will get that the "contribution of the ith clause to the distance" $\|B\boldsymbol{z} - \boldsymbol{t}\|_p^p$ is exactly $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p^p$ for some $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$, where $\boldsymbol{y} \in \{0, 1\}^k$ such that $y_s = 1$ if and only if $\boldsymbol{z}$ satisfies the sth literal of the relevant clause. (See Table 2 for a diagram showing the output of the reduction and Theorem 3.2 for the formal statement.) We stress that, while we have increased the *ambient dimension* by nearly a factor of d^*, the *rank* of the lattice is still n.

This motivates the introduction of our primary technical tool, which we call *isolating parallelepipeds*. For $1 \le p \le \infty$, a (p, k)-isolating parallelepiped is represented by a matrix $V \in \mathbb{R}^{d^* \times k}$ and a shift vector $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$ with the special property that one vertex of the parallelepiped $V\{0, 1\}^k - \boldsymbol{t}^*$ is "isolated." (Here, $V\{0, 1\}^k - \boldsymbol{t}^*$ is an affine transformation of the hypercube, i.e., a parallelepiped.) In particular, every vertex of the parallelepiped, $V\boldsymbol{y} - \boldsymbol{t}^*$ for $\boldsymbol{y} \in \{0, 1\}^k$ has unit length $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p = 1$ *except for the vertex* $-\boldsymbol{t}^*$, *which is longer, i.e.,* $\|\boldsymbol{t}^*\|_p > 1$. (See Figure 1.)

In terms of the reduction above, an isolating parallelepiped is exactly what we need. In particular, if we plug V and $\boldsymbol{t}^*$ into the above reduction, then all satisfied clauses (which correspond to non-zero $\boldsymbol{y}$ in the above description) will "contribute" 1 to the distance $\|B\boldsymbol{z} - \boldsymbol{t}\|_p^p$, while unsatisfied clauses (which correspond to $\boldsymbol{y} = \boldsymbol{0}$) will contribute $1 + \delta$ for some $\delta > 0$. Therefore, the total distance will be exactly $\|B\boldsymbol{z} - \boldsymbol{t}\|_p^p = \alpha^p n + m^+(\boldsymbol{z}) + (m - m^+(\boldsymbol{z}))(1 + \delta) = \alpha^p n - \delta m^+(\boldsymbol{z}) + (1 + \delta)m$, where $m^+(\boldsymbol{z})$ is the number of clauses satisfied by $\boldsymbol{z}$. So, the distance $\text{dist}_p(\boldsymbol{t}, \mathcal{L})$ exactly corresponds to the maximal number of satisfied clauses, as needed.

Constructing isolating parallelepipeds. Of course, in order for the above to be useful, we must show how to construct these (p, k)-isolating parallelepipeds. Indeed, it is not hard to find constructions for all $p \ge 1$ when $k = 2$, and even for all k in the special case when $p = 1$ (see Figure 1). Some other fairly nice examples can also be found for small k, as shown in Figure 2. For $p > 1$ and large k, these objects seem to be much harder to find. (In fact, in Section 4.2, we show that there is no (p, k)-isolating parallelepiped for any even integer $p \le k - 1$.) Our solution is therefore a bit technical.

At a high level, in Section 4, we consider a natural class of parallelepipeds $V \in \mathbb{R}^{2^k \times k}, \boldsymbol{t}^* \in \mathbb{R}^{2^k}$ parametrized by some weights $\alpha_0, \alpha_1, \dots, \alpha_k \ge 0$ and a scalar shift $t^* \in \mathbb{R}$. These parallelepipeds are constructed so that the length of the vertex $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p^p$ for $\boldsymbol{y} \in \{0, 1\}^k$ depends only on the Hamming weight of $\boldsymbol{y}$ and is linear in the α_i for fixed t^*. In other words, there is a matrix $M_k(p, t^*) \in \mathbb{R}^{(k+1) \times (k+1)}$ such that $M_k(p, t^*)(\alpha_0, \dots, \alpha_k)^T$ encodes the value of $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p^p$ for each possible Hamming weight of $\boldsymbol{y} \in \{0, 1\}^k$. (See Lemma 4.2.)

We show that, in order to find weights $\alpha_0, \dots, \alpha_k \ge 0$ such that V and $\boldsymbol{t}^*$ define a (p, k)-isolating parallelepiped, it suffices to find a t^* such that $M_k(p, t^*)$ is invertible. For each odd integer $p \ge 1$ and each $k \ge 2$, we show an algorithm that finds such a t^*. (See Section 4.1.)

---

[6]In particular, we replace the scalars t_i in Eq. (4) with vectors

$$\boldsymbol{t}_i := \boldsymbol{t}^* - \sum \boldsymbol{v}_s \in \mathbb{R}^{d^*} ,$$

where the sum is over s such that the sth literal in the ith clause is negated.
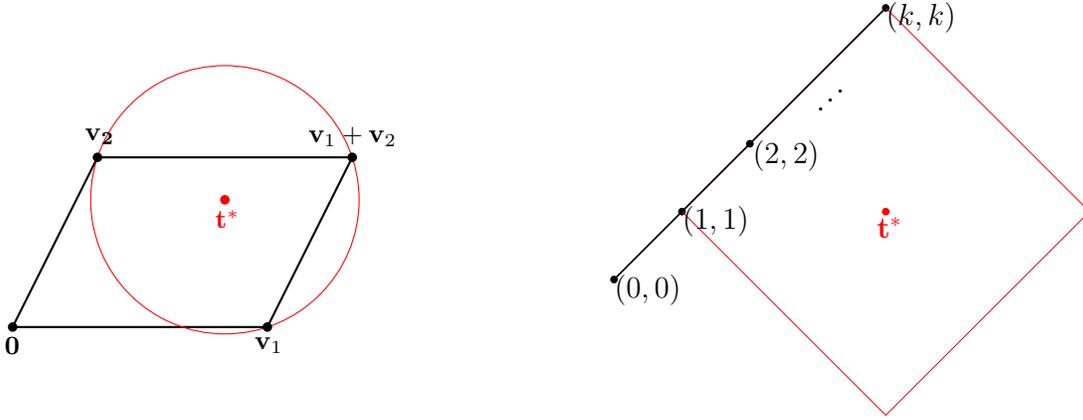
Figure 1: (p, k)-isolating parallelepipeds for $p = 2, k = 2$ (left) and $p = 1, k \geq 1$ (right). On the left, the vectors $\boldsymbol{v}_1$, $\boldsymbol{v}_2$, and $\boldsymbol{v}_1 + \boldsymbol{v}_2$ are all at the same distance from $\boldsymbol{t}^*$, while $\boldsymbol{0}$ is strictly farther away. On the right is the degenerate parallelepiped generated by k copies of the vector $(1, 1)$. The vectors (i, i) are all at the same ℓ_1 distance from $\boldsymbol{t}^*$ for $1 \leq i \leq m$, while $(0, 0)$ is strictly farther away. The (scaled) unit balls centered at $\boldsymbol{t}^*$ are shown in red, while the parallelepipeds are shown in black.

$$
V := \frac{1}{2 \cdot 12^{1/3}} \cdot \begin{pmatrix} 12^{1/3} & 12^{1/3} & 12^{1/3} \\ 1 & 1 & -1 \\ 1 & -1 & 1 \\ 1 & -1 & -1 \\ -1 & 1 & 1 \\ -1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix}, \qquad \boldsymbol{t}^* := \frac{1}{2 \cdot 12^{1/3}} \cdot \begin{pmatrix} 2 \cdot 12^{1/3} \\ 2 \\ 2 \\ 2 \\ 2 \\ 2 \\ 2 \end{pmatrix}.
$$

Figure 2: A $(3, 3)$-isolating parallelepiped in seven dimensions. One can verify that $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_3^3 = 1$ for all non-zero $\boldsymbol{y} \in \{0, 1\}^3$, and $\|\boldsymbol{t}^*\|_3^3 = 3/2$.

9

To extend this result to other $p \geq 1$, we consider the determinant of $M_k(p, t^*)$ for fixed k and t^*, viewed as a function of p. We observe that this function has a rather nice form—it is a Dirichlet polynomial. I.e., for fixed t^* and k, the determinant can be written as $\sum \exp(a_i p)$ for some $a_i \in \mathbb{R}$. Such a function has finitely many roots unless it is identically zero. So, we take the value of t^* from above such that, say, $M_k(1, t^*)$ is invertible. Since $M_k(1, t^*)$ does not have zero determinant, the Dirichlet polynomial corresponding to $\det(M_k(p, t^*))$ cannot be identically zero and therefore has finitely many roots. This is how we prove Theorem 1.2. (See Section 4.3.)

Extension to constant-factor approximation. In order to extend our hardness results to approximate CVP_p for finite p, we can try simply using the same reduction with k-SAT replaced by approximate Max-k-SAT. Unfortunately, this does not quite work. Indeed, it is easy to see that the "identity matrix gadget" that we use to restrict our attention to lattice vectors whose coordinates are in $\{0, 1\}$ (Eq. (1)) cannot tolerate an approximation factor larger than $1 + O(1/n)$ (for finite p).

However, we observe that when $k = 2$, this identity matrix gadget is actually unnecessary. In particular, even without this gadget, it "never helps" to consider a lattice vector whose coordinates are not all in $\{0, 1\}$. It then follows immediately from the analysis above that Gap-2-SAT reduces to approximate CVP_p with a constant approximation factor strictly greater than one. We note that we do not know how to extend this result to larger $k > 2$ (except when $p = 1$, see Theorem 5.3). We show that the case $k = 2$ is sufficient for proving Gap-ETH-hardness (see Proposition 2.12), but we suspect that one can just "remove the identity matrix gadget" from all of our reductions for finite p. If this were true, it would show Gap-ETH-hardness of approximation for slightly larger constant approximation factors and imply even stronger hardness results under less common assumptions.

1.6 Open questions

The most important question that we leave open is the extension of our SETH-hardness result to arbitrary $p \geq 1$. In particular, while our result applies to $p = p(n) \neq 2$ that approaches 2 asymptotically, it does not apply to the specific case $p = 2$. An extension to $p = 2$ would settle the time complexity of CVP_2 up to a factor of $2^{o(n)}$ (assuming SETH). However, we know that our technique does not work in this case (in that $(2, k)$-parallelepipeds do not exist for $k \geq 3$), so substantial new ideas might be needed to resolve this issue.

Another direction would be to strengthen our hardness of approximation results in one of two possible directions. First, one could try to increase the approximation factor. (Prior techniques for amplifying the approximation factor increase the rank of the lattice quite a bit, so they do not yield very interesting quantitative hardness results.) Second, one could try to show a reduction from Gap-k-SAT to approximate CVP_p for $k \geq 3$. For $p \in \{1, \infty\}$, we already have such a reduction, and as we mentioned above, we suspect that we can simply "remove the identity matrix gadget" in our current reduction to achieve this for $1 < p < \infty$. But, we do not know how to prove that this works.

Finally, we note that our main reduction constructs lattices of *rank n*, but the ambient dimension d can be significantly larger. (Specifically, $d = n + O(m)$, where m is the number of clauses in the relevant SAT instance, and where the hidden constant depends on k and can be very large.) Lattice problems are typically parameterized in terms of the rank of the lattice (and for the ℓ_2 norm, one can assume without loss of generality that $d = n$), but it is still interesting to ask whether we can reduce the ambient dimension d.

Organization

In Section 2, we review some necessary background knowledge. In Section 3, we show how to use a (p, k)-isolating parallelepiped (for finite p) to reduce any n-variable instance of k-SAT to a CVP_p instance with rank n, and we show that this immediately gives SETH-hardness for $p = 1$. In Section 4, we show how to construct (p, k)-isolating parallelepipeds, first for odd integers $p \geq 1$ and then for "almost all" p. In Section 5, we show $2^{\Omega(n)}$-hardness of approximating CVP_p up to a constant factor. In Section 6, we prove a number of additional hardness results: $2^{\Omega(\sqrt{n})}$ ETH- and Max-2-SAT-hardness of CVPP_p (Section 6.1), ETH- and Max-2-SAT-hardness of CVP_p (Section 6.2), and SETH-hardness of CVP_∞ and SVP_∞ (Section 6.3).

2 Preliminaries

Throughout this paper, we work with lattice problems over $\mathbb{R}^d$ for convenience. Formally, we must pick a suitable representation of real numbers and consider both the size of the representation and the efficiency of arithmetic operations in the given representation. But, we omit such details throughout to ease readability.

2.1 Computational lattice problems

Let $\mathrm{dist}_p(\mathcal{L}, \boldsymbol{t}) := \min_{\boldsymbol{x} \in \mathcal{L}} \|\boldsymbol{x} - \boldsymbol{t}\|_p$ denote the ℓ_p distance of $\boldsymbol{t}$ to $\mathcal{L}$. In addition to SVP and CVP, we also consider a variant of CVP called the Closest Vector Problem with Preprocessing (CVPP), which allows arbitrary preprocessing of a lattice.

Definition 2.1. *For any $\gamma \geq 1$ and $1 \leq p \leq \infty$, the γ-approximate Shortest Vector Problem with respect to the ℓ_p norm (γ-SVP_p) is the promise problem defined as follows. Given a lattice $\mathcal{L}$ (specified by a basis $B \in \mathbb{R}^{d \times n}$) and a number $r > 0$, distinguish between a 'YES' instance where there exists a non-zero vector $\boldsymbol{v} \in \mathcal{L}$ such that $\|\boldsymbol{v}\|_p \leq r$, and a 'NO' instance where $\|\boldsymbol{v}\|_p > \gamma r$ for all non-zero $v \in \mathcal{L}$.*

Definition 2.2. *For any $\gamma \geq 1$ and $1 \leq p \leq \infty$, the γ-approximate Closest Vector Problem with respect to the ℓ_p norm (γ-CVP_p) is the promise problem defined as follows. Given a lattice $\mathcal{L}$ (specified by a basis $B \in \mathbb{R}^{d \times n}$), a target vector $\boldsymbol{t} \in \mathbb{R}^d$, and a number $r > 0$, distinguish between a 'YES' instance where $\mathrm{dist}_p(\mathcal{L}, \boldsymbol{t}) \leq r$, and a 'NO' instance where $\mathrm{dist}_p(\mathcal{L}, \boldsymbol{t}) > \gamma r$.*

When $\gamma = 1$, we refer to the problems simply as SVP_p and CVP_p, respectively.

Definition 2.3. *The Closest Vector Problem with Preprocessing with respect to the ℓ_p norm (CVPP_p) is the problem of finding a preprocessing function P and an algorithm Q which work as follows. Given a lattice $\mathcal{L}$ (specified by a basis $B \in \mathbb{R}^{d \times n}$), P outputs a new description of $\mathcal{L}$. Given $P(\mathcal{L})$, a target vector $\boldsymbol{t} \in \mathbb{R}^d$, and a number $r > 0$, Q decides whether $\mathrm{dist}_p(\mathcal{L}, \boldsymbol{t}) \leq r$.*

When we measure the running time of a CVPP algorithm, we only count the running time of Q, and not of the preprocessing algorithm P.

2.2 Satisfiability problems and the Max-Cut problem

A k-SAT formula Φ on n Boolean variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$ is a conjunction $\Phi = \wedge_{i=1}^{m} C_i$ of clauses $C_i = \vee_{s=1}^{k} \ell_{i,s}$, where the literals $\ell_{i,s}$ denote a variable x_j or its negation $\neg x_j$. The goal is to decide whether there exists an assignment $\boldsymbol{a} \in \{0,1\}^n$ to the variables of Φ such that all clauses have at least one "true" literal, i.e., so that all clauses are satisfied.

The *value* of a k-SAT formula Φ, denoted val(Φ), is the maximum fraction of clauses satisfied by an assignment to Φ.

Definition 2.4. *Given a k-SAT formula Φ and constants $0 \leq \delta < \varepsilon \leq 1$, the (δ, ε)-Gap-k-SAT problem is the promise problem defined as follows. The goal is to distinguish between a 'YES' instance in which val(Φ) $\geq \varepsilon$, and a 'NO' instance in which val(Φ) $< \delta$.*

Definition 2.5. *Given a k-SAT formula Φ with clauses $\mathcal{C} = \{C_1, \ldots, C_m\}$, a clause weight function $w : \mathcal{C} \to \mathbb{Z}^+$, and a weight threshold W, the Weighted Max-k-SAT problem is to decide whether there exists an assignment $\boldsymbol{a}$ to the variables of Φ such that $\sum_{C_i \text{ is sat. by } \boldsymbol{a}} w(C_i) \geq W$.*

Definition 2.6. *Given an undirected graph $G = (V, E)$, an edge weight function $w : E \to \mathbb{Z}^+$, and a weight threshold W, the Weighted Max-CUT problem is defined as follows. The goal is to decide whether V can be partitioned into sets V_1 and V_2 such that $\sum_{e_{ij} \in E : v_i \in V_1, v_j \in V_2} w(e_{ij}) \geq W$.*

There exists a folklore reduction from an instance of Weighted Max-Cut on a graph with n vertices to an instance of Weighted Max 2-SAT on a formula with n variables. See, e.g., [GHNR03].

2.3 Exponential Time Hypotheses

Impagliazzo and Paturi [IP99] introduced the following two hypotheses (ETH and SETH), which are now widely used to study the quantitative hardness of computational problems.

Definition 2.7. *The Exponential Time Hypothesis (ETH) is the hypothesis defined as follows. For every $k \geq 3$ there exists a constant $\varepsilon > 0$ such that no algorithm solves k-SAT formulas with n variables in $2^{\varepsilon n}$-time. In particular, there is no $2^{o(n)}$-time algorithm for 3-SAT.*

Definition 2.8. *The Strong Exponential Time Hypothesis (SETH) is the hypothesis defined as follows. For every constant $\varepsilon > 0$ there exists k such that no algorithm solves k-SAT formulas with n variables in $2^{(1-\varepsilon)n}$-time.*

An important tool in the study of the exact complexity of k-SAT is the Sparisification Lemma of Impagliazzo, Paturi, and Zane [IPZ01] which roughly says that any k-SAT formula can be replaced with $2^{\varepsilon n}$ formulas each with $O(n)$ clauses for some $\varepsilon > 0$.

Proposition 2.9 (Sparsification Lemma, [IPZ01])**.** *For every $k \in \mathbb{Z}^+$ and $\varepsilon > 0$ there exists a constant $c = c(k, \varepsilon)$ such that any k-SAT formula Φ with n variables can be expressed as $\Phi = \vee_{i=1}^{r} \Psi_i$ where $r \leq 2^{\varepsilon n}$ and each Ψ_i is a k-SAT formula with at most cn clauses. Moreover, this disjunction can be computed in $2^{\varepsilon n}$-time.*

In this paper we also consider the W-Max-SAT-SETH hypothesis, which corresponds to SETH but with Weighted Max-k-SAT in place of k-SAT. Our main result only relies on this weaker variant of SETH, and is therefore more robust.

Dinur [Din16] and Manurangsi and Raghavendra [MR17] recently introduced a "gap" version of ETH, which asserts that Gap-3-SAT takes $2^{\Omega(n)}$-time.

Definition 2.10. *The (randomized) Gap-Exponential Time Hypothesis ((randomized) Gap-ETH) is the hypothesis that there exist constants $\delta < 1$ and $\varepsilon > 0$ such that no (randomized) algorithm solves $(\delta, 1)$-Gap-3-SAT instances with n variables in $2^{\varepsilon n}$-time.*

As Dinur [Din16] notes, one can sparsify a Gap-SAT instance simply by sampling clauses. Therefore, we can assume (almost) without loss of generality that Gap-ETH applies only to formulas with $O(n)$ clauses. The caveat is that the sampling is randomized, so finding a $2^{o(n)}$-time algorithm for sparse Gap-3-SAT only implies a randomized $2^{o(n)}$-time algorithm for general Gap-3-SAT.

We give a variant of Dinur's sampling argument in Proposition 2.11. The idea is to show that both the total number of sampled clauses and the number of sampled clauses that are satisfied by any given assignment are highly concentrated around their expectation by using the Chernoff bound, and then to take a union bound over the bad events where these quantities deviate substantially from their expectation.

We will use the following multiplicative Chernoff bounds (see, e.g., [HP]). Let $X_1, \dots, X_n$ be independent identically distributed Bernoulli random variables with expectation p, so that the expectation of $\sum_{i=1}^{n} X_i$ is $\mu = pn$. Then:

$$\Pr[\sum_{i=1}^{n} X_i < (1-\alpha)\mu] < e^{-\mu\alpha^2/2} \ , \tag{6}$$

$$\Pr[\sum_{i=1}^{n} X_i > (1+\alpha)\mu] < e^{-\mu\alpha^2/4} \ . \tag{7}$$

Proposition 2.11 (Sparsification for Gap-SAT). *For any $0 < \delta < \delta' < 1$, there is a polynomial-time randomized reduction from a $(\delta, 1)$-Gap-k-SAT instance Φ with n variables and m clauses to a $(\delta', 1)$-Gap-k-SAT instance Φ' with n variables and $O(n)$ clauses.*

Proof. Let Φ' be the formula obtained by sampling each clause of Φ independently with probability $p := \min\{1, 10/(\delta\alpha^2) \cdot n/m\}$, where α is fixed so that $-(1-\delta'/\delta)/(1+\delta'/\delta) < \alpha < 1$. Clearly, if $\text{val}(\Phi) = 1$ then $\text{val}(\Phi') = 1$ as well. We analyze the case where $\text{val}(\Phi) < \delta$.

In expectation Φ' has pm clauses. Furthermore, because $\text{val}(\Phi') < \delta$, in expectation any fixed assignment will satisfy fewer than δpm clauses of Φ'. Therefore by Equation (6),

$$\Pr[\text{Number of clauses in } \Phi' \leq (1-\alpha)pm] \leq e^{-\alpha^2 pm/2} \leq e^{-2n}. \tag{8}$$

Furthermore, by Equation (7), we have that for each fixed assignment $\boldsymbol{a}$,

$$\Pr[\text{Number of clauses in } \Phi' \text{ sat. by } \boldsymbol{a} \geq (1+\alpha)\delta pm] \leq e^{-\alpha^2 \delta pm/4} \leq e^{-2n}. \tag{9}$$

By applying Equations (8) and (9), and taking a union bound we get that the probability that Φ' has at least $(1-\alpha)pm$ clauses and that no assignment to Φ' satisfies more than $(1+\alpha)\delta pm$ clauses is at least $1 - (e^{-2n} + 2^n e^{-2n}) \geq 1 - 2e^{-n}$. Therefore,

$$\text{val}(\Phi') \leq \frac{(1+\alpha)pm}{(1-\alpha)pm} \cdot \delta < \delta'$$

with high probability. $\qquad\square$

	x_1	x_2	$\cdots$	x_{n-1}	x_n	
$C_1\Big\{$	$\boldsymbol{v}_1$	$\boldsymbol{v}_2$	$\cdots$	$\boldsymbol{0}_{d^*}$	$-\boldsymbol{v}_3$	$\boldsymbol{t}^* - \boldsymbol{v}_3$
$\vdots$	$\vdots$	$\cdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$
$C_m\Big\{$	$\boldsymbol{0}_{d^*}$	$-\boldsymbol{v}_1$	$\cdots$	$\boldsymbol{v}_2$	$\boldsymbol{v}_3$	$\boldsymbol{t}^* - \boldsymbol{v}_1$
x_1	$2\alpha^{1/p}$	0	$\cdots$	0	0	$\alpha^{1/p}$
x_2	0	$2\alpha^{1/p}$	$\cdots$	0	0	$\alpha^{1/p}$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$
x_{n-1}	0	0	$\cdots$	$2\alpha^{1/p}$	0	$\alpha^{1/p}$
x_n	0	0	$\cdots$	0	$2\alpha^{1/p}$	$\alpha^{1/p}$
			B			$\boldsymbol{t}$

Table 2: A basis B and target vector $\boldsymbol{t}$ output by the reduction from Theorem 3.2 with some $(p,3)$-isolating parallelepiped given by $V = (\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3) \in \mathbb{R}^{d^* \times 3}$ and $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$. In this example, the first clause is $C_1 \equiv x_1 \vee x_2 \vee \neg x_n$ and the mth clause is $C_m \equiv \neg x_2 \vee x_{n-1} \vee x_n$. By the definition of an isolating parallelepiped (Definition 3.1), the contribution of the first d coordinates to the distance $\|B\boldsymbol{z} - \boldsymbol{t}\|_p^p$ will be 1 for any assignment $\boldsymbol{z} \in \{0,1\}^n$ satisfying C_1, while non-satisfying assignments contribute $(1 + \delta)$ for some $\delta > 0$. For example, if $z_1 = 1, z_2 = 0, z_n = 1$, the clause C_1 is satisfied, and the first d coordinates will contribute $\|\boldsymbol{v}_1 - \boldsymbol{v}_3 - (\boldsymbol{t}^* - \boldsymbol{v}_3)\|_p^p = \|\boldsymbol{v}_1 - \boldsymbol{t}^*\|_p^p = 1$. On the other hand, if $z_1 = 0, z_2 = 0, z_n = 1$, then C_1 is not satisfied, and $\| - \boldsymbol{v}_3 - (\boldsymbol{t}^* - \boldsymbol{v}_3)\|_p^p = \|\boldsymbol{t}^*\|_p^p = 1 + \delta$.

Additionally, we will use a reduction of Garey et al. [GJS76] from 3-SAT to Max-2-SAT which also works as a reduction from Gap-3-SAT to Gap-2-SAT. The reduction works by outputting ten 1- and 2-clauses for each 3-clause in the original formula. Any assignment which satisfies the original clause corresponds to an assignment which satisfies 7 of the output clauses, and any assignment which does not satisfy the original clause corresponds to an assignment which satisfies 6 of the output clauses.

Proposition 2.12 ([GJS76], Theorem 1.1]). *For every $0 \leq \delta < \varepsilon \leq 1$, there is a polynomial-time reduction from every instance of (δ, ε)-Gap-3-SAT with n variables and m clauses to an instance of $((6+\delta)/10, (6+\varepsilon)/10)$-Gap-2-SAT with $n + m$ variables and $10m$ clauses.*

3 SETH-hardness from isolating parallelepipeds

We start by giving a reduction from instances of weighted Max-k-SAT on formulas with n variables to instances of CVP_p with rank n for all p that uses a certain geometric object, which we define next. Let $\boldsymbol{1}_n$ and $\boldsymbol{0}_n$ denote the all 1s and all 0s vectors of length n respectively, and let I_n denote the $n \times n$ identity matrix.

Definition 3.1. *For any $1 \leq p \leq \infty$ and integer $k \geq 2$, we say that $V \in \mathbb{R}^{d^* \times k}$ and $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$ define a (p, k)-isolating parallelepiped if $\|\boldsymbol{t}\|_p > 1$ and $\|V\boldsymbol{x} - \boldsymbol{t}^*\|_p = 1$ for all $\boldsymbol{x} \in \{0, 1\}^k \setminus \{\boldsymbol{0}_k\}$.*

In order to give the reduction, we first introduce some notation related to SAT. Let Φ be a k-SAT formula on n variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$. Let $\text{ind}(\ell)$ denote the index of the variable underlying a literal ℓ. I.e., $\text{ind}(\ell) = j$ if $\ell = x_j$ or $\ell = \neg x_j$. Call a literal ℓ *positive* if $\ell = x_j$ and *negative* if $\ell = \neg x_j$ for some variable x_j. Given a clause $C_i = \vee_{s=1}^{k} \ell_{i,s}$, let $P_i := \{s \in [k] : \ell_{i,s} \text{ is positive}\}$ and let $N_i := \{s \in [k] : \ell_{i,s} \text{ is negative}\}$ denote the indices of positive and negative literals in C_i respectively. Given an assignment $\boldsymbol{a} \in \{0, 1\}^n$ to the variables of Φ, let $S_i(\boldsymbol{a})$ denote the indices of literals in C_i satisfied by $\boldsymbol{a}$. I.e., $S_i(\boldsymbol{a}) := \{s \in P_i : a_{\text{ind}(\ell_{i,s})} = 1\} \cup \{s \in N_i : a_{\text{ind}(\ell_{i,s})} = 0\}$. Finally, let $m^+(\boldsymbol{a})$ denote the number of clauses of Φ satisfied by the assignment $\boldsymbol{a}$, i.e., the number of clauses i for which $|S_i(\boldsymbol{a})| \geq 1$.

Theorem 3.2. *If there exists a computable (p, k)-isolating parallelepiped for some $p = p(n) \in [1, \infty)$ and integer $k \geq 2$, then there exists a polynomial-time reduction from any (weighted-)Max-k-SAT instance with n variables to a CVP_p instance of rank n.*

Proof. For simplicity, we give a reduction from unweighted Max-k-SAT, and afterwards sketch how to modify our reduction to handle the weighted case as well. Namely, we give a reduction from any Max-k-SAT instance (Φ, W) to an instance $(B, \boldsymbol{t}^*, r)$ of CVP_p. Here, the formula Φ is on n variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$. (Φ, W) is a 'YES' instance if there exists an assignment $\boldsymbol{a}$ such that $m^+(\boldsymbol{a}) \geq W$.

By assumption, there exist computable $d^* = d^*(p, k) \in \mathbb{Z}^+$, $V = [\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k] \in \mathbb{R}^{d^* \times k}$, and $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$ such that $\|\boldsymbol{t}^*\|_p = (1 + \delta)^{1/p}$ for some $\delta > 0$ and $\|V\boldsymbol{z} - \boldsymbol{t}^*\|_p = 1$ for all $\boldsymbol{z} \in \{0, 1\}^k \setminus \{\boldsymbol{0}_k\}$.

We define the output CVP_p instance as follows. Let $d := md^* + n$. The basis $B \in \mathbb{R}^{d \times n}$ and target vector $\boldsymbol{t} \in \mathbb{R}^d$ in the output instance have the form

$$B = \begin{pmatrix} B_1 \\ \vdots \\ B_m \\ 2\alpha^{1/p} \cdot I_n \end{pmatrix}, \qquad \boldsymbol{t} = \begin{pmatrix} \boldsymbol{t}_1 \\ \vdots \\ \boldsymbol{t}_m \\ \alpha^{1/p} \cdot \boldsymbol{1}_n \end{pmatrix},$$

with blocks $B_i \in \mathbb{R}^{d^* \times n}$ and $\boldsymbol{t}_i \in \mathbb{R}^{d^*}$ for $1 \leq i \leq m$ and $\alpha := m + (m - W)\delta$. Note that α is the maximum possible contribution of the clauses $C_1, \ldots, C_m$ to $\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p$ when (Φ, W) is a 'YES' instance. For every $1 \leq i \leq m$ and $1 \leq j \leq n$, set the jth column $(B_i)_j$ of block B_i (corresponding to the clause $C_i = \vee_{s=1}^{k} \ell_{i,s}$) as

$$(B_i)_j := \begin{cases} \boldsymbol{v}_s & \text{if } x_j \text{ is the } s\text{th literal of clause } i, \\ -\boldsymbol{v}_s & \text{if } \neg x_j \text{ is the } s\text{th literal of clause } i, \\ \boldsymbol{0}_{d^*} & \text{otherwise,} \end{cases}$$

and set $\boldsymbol{t}_i := \boldsymbol{t}^* - \sum_{s \in N_i} \boldsymbol{v}_s$. Set $r := (\alpha(n+1))^{1/p}$.

Clearly, the reduction runs in polynomial time. We next analyze for which $\boldsymbol{y} \in \mathbb{Z}^n$ it holds that $\|B\boldsymbol{y} - \boldsymbol{t}\|_p \leq r$. Given $\boldsymbol{y} \notin \{0, 1\}^n$,

$$\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p \geq \|2\alpha^{1/p}I_n\boldsymbol{y} - \alpha^{1/p}\boldsymbol{1}_n\|_p^p \geq \alpha(n+2) > r^p \,,$$

so we only need to analyze the case when $\boldsymbol{y} \in \{0,1\}^n$. Consider an assignment $\boldsymbol{y} \in \{0,1\}^n$ to the variables of Φ. Then,

$$
\begin{aligned}
\|B_i \boldsymbol{y} - \boldsymbol{t}_i\|_p &= \Big\| \sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} \cdot \boldsymbol{v}_s - \sum_{s \in N_i} y_{\mathrm{ind}(\ell_{i,s})} \cdot \boldsymbol{v}_s - \Big(\boldsymbol{t}^* - \sum_{s \in N_i} \boldsymbol{v}_s \Big) \Big\|_p \\
&= \Big\| \sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} \cdot \boldsymbol{v}_s + \sum_{s \in N_i} \big(1 - y_{\mathrm{ind}(\ell_{i,s})} \big) \cdot \boldsymbol{v}_s - \boldsymbol{t}^* \Big\|_p \\
&= \Big\| \sum_{s \in S_i(\boldsymbol{a})} \boldsymbol{v}_s - \boldsymbol{t}^* \Big\|_p .
\end{aligned}
$$

By assumption, the last quantity is equal to 1 if $|S_i(\boldsymbol{y})| \geq 1$, and is equal to $(1+\delta)^{1/p}$ otherwise. Because $|S_i(\boldsymbol{y})| \geq 1$ if and only if C_i is satisfied, it follows that

$$
\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p = \Big(\sum_{i=1}^m \|B_i \boldsymbol{y} - \boldsymbol{t}_i\|_p^p \Big) + \alpha n = m + (m - m^+(\boldsymbol{y}))\delta + \alpha n .
$$

Therefore, $\|B\boldsymbol{y} - \boldsymbol{t}\|_p \leq r$ if and only if $m^+(\boldsymbol{y}) \geq W$, and therefore there exists $\boldsymbol{y}$ such that $\|B\boldsymbol{y} - \boldsymbol{t}\|_p \leq r$ if and only if (Φ, W) is a 'YES' instance of Max-k-SAT, as needed.

To extend this to a reduction from *weighted* Max-k-SAT to CVP_p, simply multiply each block B_i and the corresponding target vector $\boldsymbol{t}_i$ by $w(C_i)^{1/p}$, where $w(C_i)$ denotes the weight of the clause C_i. Then, by adjusting α to depend on the weights $w(C_i)$ we obtain the desire reduction. $\qquad\square$

Because the rank n of the output CVP_p instance matches the number of variables in the input SAT formula, we immediately get the following corollary.

Corollary 3.3. *For any efficiently computable $p = p(n) \in [1, \infty)$ if there exists a computable (p, k)-isolating parallelepiped for infinitely many $k \in \mathbb{Z}^+$, then, for every constant $\varepsilon > 0$ there is no $2^{(1-\varepsilon)n}$-time algorithm for CVP_p assuming W-Max-SAT-SETH. In particular there is no $2^{(1-\varepsilon)n}$-time algorithm for CVP_p assuming SETH.*

It is easy to construct a (degenerate) family of isolating parallelepipeds for $p = 1$, and therefore we get hardness of CVP_1 as a simple corollary. (See Figure 1.)

Corollary 3.4. *For every constant $\varepsilon > 0$ there is no $2^{(1-\varepsilon)n}$-time algorithm for CVP_1 assuming W-Max-SAT-SETH, and in particular there is no $2^{(1-\varepsilon)n}$-time algorithm for CVP_1 assuming SETH.*

Proof. Let $k \in \mathbb{Z}^+$, let $V = [\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k]$ with $\boldsymbol{v}_1 = \cdots = \boldsymbol{v}_k := \frac{1}{k-1}(1,1)^T \in \mathbb{R}^2$, and let $\boldsymbol{t}^* := \frac{1}{k-1}(k,1)^T \in \mathbb{R}^2$. Then, $\|V\boldsymbol{x} - \boldsymbol{t}^*\|_1 = 1$ for every $\boldsymbol{x} \in \{0,1\}^k \setminus \{\boldsymbol{0}_k\}$, and $\|\boldsymbol{t}^*\|_1 = (k+1)/(k-1) > 1$. The result follows by Corollary 3.3. $\qquad\square$

4 Finding isolating parallelepipeds

We now show how to find a (p, k)-isolating parallelepiped given by $V \in \mathbb{R}^{d^* \times k}$ and $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$ as in Definition 3.1. We will first show a general strategy for trying to find such an object for any $p \geq 1$ and integer $k \geq 2$. In Section 4.1, we will show how to successfully implement this strategy in the case when p is an odd integer. In Section 4.2, we show that (p, k)-isolating parallelepipeds do not

exist for even integers $p \leq k - 1$. Finally, in Section 4.3 we show how to mostly get around this issue in order to find (p, k)-isolating parallelepipeds for "almost all" $p \geq 1$.

It will actually be convenient to find a slightly different object that "works with $\{\pm 1\}^k$ instead of $\{0, 1\}^k$." We observe below that this suffices.

Lemma 4.1. *There is an efficient algorithm that takes as input a matrix $V \in \mathbb{R}^{d^* \times k}$ and vector $\boldsymbol{t}^* \in \mathbb{R}^{d^*}$ such that $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p = 1$ for any $\boldsymbol{y} \in \{\pm 1\}^k \setminus \{-\boldsymbol{1}_k\}$ and $\| - V\boldsymbol{1}_k - \boldsymbol{t}^*\|_p > 1$, and outputs a matrix $V' \in \mathbb{R}^{d^* \times k}$ and vector $(\boldsymbol{t}^*)' \in \mathbb{R}^{d^*}$ that form a (p, k)-isolating parallelepiped.*

Proof. Define $V' := 2V$ and $(\boldsymbol{t}^*)' = V\boldsymbol{1}_k + \boldsymbol{t}^*$. Now consider the affine transformation $f \colon \mathbb{R}^k \to \mathbb{R}^k$ defined by $f(\boldsymbol{x}) := (2\boldsymbol{x} - \boldsymbol{1}_k)$, which maps $\{0, 1\}^k$ to $\{\pm 1\}^k$ and $\boldsymbol{0}_k$ to $-\boldsymbol{1}_k$. Then, for $\boldsymbol{x} \in \{0, 1\}^k$ and $\boldsymbol{y} = f(\boldsymbol{x}) = 2\boldsymbol{x} - \boldsymbol{1}_k \in \{\pm 1\}^k$, we have

$$\|V'\boldsymbol{x} - (\boldsymbol{t}^*)'\|_p = \left\|V' \frac{\boldsymbol{y} + \boldsymbol{1}_k}{2} - (\boldsymbol{t}^*)'\right\|_p = \left\|V' \frac{\boldsymbol{y}}{2} + V' \frac{\boldsymbol{1}_k}{2} - (\boldsymbol{t}^*)'\right\|_p = \|V\boldsymbol{y} - \boldsymbol{t}^*\|_p \, ,$$

as needed. $\qquad\qquad\square$

Intuitively, a "reasonable" matrix V should act symmetrically on bit strings. I.e., if $\boldsymbol{y}, \boldsymbol{y}' \in \{\pm 1\}^k$ have the same number of positive entries, then $V\boldsymbol{y}$ should be a permutation of $V\boldsymbol{y}'$. This implies that any row of V must be accompanied by all possible permutations of this row. If we further require that each row in V is $\alpha \cdot \boldsymbol{v}$ for some $\boldsymbol{v} \in \{\pm 1\}^k$ and $\alpha \in \mathbb{R}$, then we arrive at a very general construction that is still possible to analyze.

For weights $\alpha_0, \ldots, \alpha_k \geq 0$, we define $V := V(\alpha_0, \ldots, \alpha_k) \in \mathbb{R}^{2^k \times k}$ as follows. The rows of V are indexed by the strings $\{\pm 1\}^k$, and row $\boldsymbol{v}$ is $\alpha_{k-|\boldsymbol{v}|}^{1/p} \boldsymbol{v}^T$, where

$$|\boldsymbol{v}| := \{\# \text{ of positive entries in } \boldsymbol{v}\}$$

is the number of positive entries in $\boldsymbol{v}$. For a shift $t^* \in \mathbb{R}$, we set $\boldsymbol{t}^* := \boldsymbol{t}^*(\alpha_0, \ldots, \alpha_k, t^*) \in \mathbb{R}^{2^k}$ such that the coordinate of $\boldsymbol{t}^*$ corresponding to $\boldsymbol{v}$ is $\alpha_{k-|\boldsymbol{v}|}^{1/p} t^*$. (Figure 2 is an example of this construction. In particular, it shows $V(\alpha_0, \alpha_1, \alpha_2, \alpha_3)$ with $\alpha_0 = 12$, $\alpha_1 = \alpha_2 = 1$ and $\alpha_3 = 0$ and $\boldsymbol{t}^*(\alpha_0, \alpha_1, \alpha_2, \alpha_3, t^*)$ with $t^* = 2$, where we have omitted the last row, whose weight is zero.)

In what follows, we will use the binomial coefficient $\binom{i}{j}$ extensively, and we adopt the convention that $\binom{i}{j} = 0$ if $j > i$ or $j < 0$ or $j \notin \mathbb{Z}$.

Lemma 4.2. *For any $\boldsymbol{y} \in \{\pm 1\}^k$, weights $\alpha_0, \ldots, \alpha_k \geq 0$, and shift $t^* \in \mathbb{R}$,*

$$\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p^p = \sum_{j=0}^k \alpha_{k-j} \sum_{\ell=0}^k \binom{|\boldsymbol{y}|}{\ell} \binom{k - |\boldsymbol{y}|}{j - \ell} \cdot \big|2|\boldsymbol{y}| + 2j - k - 4\ell - t^*\big|^p \, ,$$

where $V := V(\alpha_0, \ldots, \alpha_k) \in \mathbb{R}^{2^k \times k}$ and $\boldsymbol{t}^* := \boldsymbol{t}^*(\alpha_0, \ldots, \alpha_k, t^*)$ as above.

In other words, $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p^p$ depends only on $|\boldsymbol{y}|$, and if $\boldsymbol{w} \in (\mathbb{R}^{\geq 0})^{k+1}$ is the vector such that $w_j = \|V\boldsymbol{y}' - \boldsymbol{t}^*\|_p^p$ for all $\boldsymbol{y}' \in \{\pm 1\}$ with $|\boldsymbol{y}'| = j$, then

$$\boldsymbol{w} = M_k(p, t^*) \begin{pmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_k \end{pmatrix} ,$$

where $M_k(p, t^*) \in \mathbb{R}^{(k+1) \times (k+1)}$ is given by

$$M_k(p, t^*)_{i,j} := \sum_{\ell=0}^{k} \binom{i}{\ell} \binom{k-i}{j-\ell} \cdot |2i + 2j - k - 4\ell - t^*|^p . \tag{10}$$

Proof. We have

$$\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p^p = \sum_{j=0}^{k} \alpha_j \sum_{|\boldsymbol{v}|=k-j} |\langle \boldsymbol{v}, \boldsymbol{y} \rangle - t^*|^p .$$

Notice that $\langle \boldsymbol{v}, \boldsymbol{y} \rangle$ depends only on how many of the j negative entries of $\boldsymbol{v}$ align with the positive entries of $\boldsymbol{y}$. In particular,

$$\sum_{|\boldsymbol{v}|=k-j} |\langle \boldsymbol{v}, \boldsymbol{y} \rangle - t^*|^p = \sum_{\ell=0}^{k} \binom{|\boldsymbol{y}|}{\ell} \binom{k-|\boldsymbol{y}|}{j-\ell} \cdot |-\ell + (|\boldsymbol{y}| - \ell) + (j - \ell) - (k - |\boldsymbol{y}| - j + \ell) - t^*|^p$$

$$= \sum_{\ell=0}^{k} \binom{|\boldsymbol{y}|}{\ell} \binom{k-|\boldsymbol{y}|}{j-\ell} \cdot |2|\boldsymbol{y}| + 2j - k - 4\ell - t^*|^p ,$$

as needed. $\qquad \square$

Lemma 4.3. *For any $t^* \in \mathbb{R}$, the matrix $M_k(p, t^*)$ defined in Eq. (10) is stochastic. I.e., $M_k(p, t^*)\mathbf{1}_{k+1} = \lambda(t^*)\mathbf{1}_{k+1}$ for some $\lambda(t^*) \in \mathbb{R}$. Furthermore, $\lambda(t^*) > 0$.*

Proof. We rearrange the sum corresponding to the ith entry of $M_k(p, t^*)\mathbf{1}_{k+1}$, setting $r := (i+j)/2 - \ell$ to obtain

$$\sum_{j=0}^{k} \sum_{\ell=0}^{k} \binom{i}{\ell} \binom{k-i}{j-\ell} \cdot |2i + 2j - k - 4\ell - t^*|^p = \sum_{r} |4r - k - t^*|^p \sum_{j=0}^{k} \binom{i}{(i+j)/2 - r} \binom{k-i}{r + (j-i)/2}$$

$$= \sum_{r} |4r - k - t^*|^p \sum_{j=0}^{k} \binom{i}{r + (i-j)/2} \binom{k-i}{r + (j-i)/2} .$$

Finally, we recall Vandermonde's identity, which says that

$$\sum_{j=0}^{k} \binom{i}{r + (i-j)/2} \binom{k-i}{r + (j-i)/2} = \binom{k}{r} .$$

Therefore, the summation does not depend on i (and is clearly positive), as needed. $\qquad \square$

Lemma 4.3 tells us that for any $t^* \in \mathbb{R}$, $M_k(p, t^*)(\mathbf{1}_{k+1})/\lambda = \mathbf{1}_{k+1}$ for some $\lambda > 0$. We wish to show that, for some $t^* \in \mathbb{R}$, we can find $\alpha_0, \ldots, \alpha_k \geq 0$ such that $M_k(p, t^*)(\alpha_0, \alpha_1, \ldots, \alpha_k)^T = \mathbf{1}_{k+1} + \varepsilon \boldsymbol{e}_0$ for some $\varepsilon > 0$, where $\boldsymbol{e}_0 := (1, 0, \ldots, 0)^T$. In order to do this, it suffices to show that $M_k(p, t^*)$ is invertible. Then, we can take

$$(\alpha_0, \alpha_1, \ldots, \alpha_k)^T := \mathbf{1}_{k+1}/\lambda + \varepsilon M_k(p, t^*)^{-1} \boldsymbol{e}_0 .$$

If $\varepsilon := (\lambda \cdot \|M_k(p, t^*)^{-1} \boldsymbol{e}_0\|_\infty)^{-1} > 0$, then the α_i must be non-negative. We make this formal in the next proposition.

Proposition 4.4. *There is an efficient algorithm that takes as input any $p \geq 1$, an integer $k \geq 2$, and $t^* \in \mathbb{R}$ such that $\det(M_k(p, t^*)) \neq 0$, where $M_k(p, t^*)$ is defined as in Eq. (10) and outputs $V \in \mathbb{R}^{2^k \times k}$ and $\boldsymbol{t}^* \in \mathbb{R}^{2^k}$ that define a (p, k)-isolating parallelepiped.*

Proof. By Lemma 4.1, it suffices to construct a matrix that works for $\boldsymbol{y} \in \{\pm 1\}^k$. The algorithm behaves as follows on input $k \geq 2$ and $p \geq 1$ and $t^* \in \mathbb{R}$. By Lemma 4.3, $M_k(p, t^*)\boldsymbol{1}_{k+1} = \lambda \boldsymbol{1}_{k+1}$ for some $\lambda > 0$. Since we are promised that $\det(M_k(p, t^*)) \neq 0$, we see that $M_k(p, t^*)$ is invertible. The algorithm therefore sets

$$(\alpha_0, \ldots, \alpha_k)^T := \boldsymbol{1}_{k+1}/\lambda + \varepsilon M_k(p, t^*)^{-1}\boldsymbol{e}_0 \, , \tag{11}$$

where $\varepsilon := (\lambda \cdot \|M_k(p, t^*)^{-1}\boldsymbol{e}_0\|_\infty)^{-1} > 0$ is chosen to be small enough such that the α_i are all non-negative. Finally, it outputs $V := V(\alpha_0, \ldots, \alpha_k)$ and $\boldsymbol{t}^* := \boldsymbol{t}^*(\alpha_0, \ldots, \alpha_k, t^*)$ as defined above.

To prove correctness, we note that V and $\boldsymbol{t}^*$ have the desired property. Indeed, it follows from the definition of $M_k(p, t^*)$ in Lemma 4.2 that $\|V\boldsymbol{y} - \boldsymbol{t}^*\|_p^p$ is the jth coordinate of $\boldsymbol{w} := M_k(p, t^*)(\alpha_0, \ldots, \alpha_k)^T$, where $j := |\boldsymbol{y}|$. But, by Eq. (11), we see that the jth coordinate of $\boldsymbol{w}$ is $1 + \varepsilon$ if $j = 0$ and is 1 otherwise, as needed. $\qquad\square$

4.1   Finishing the proof for odd integer p

We now handle the case when $p \geq 1$ is an odd integer. Notice that, if $p \geq 1$ is an integer, then $\det(M_k(p, t^*))$ is some piecewise combination of polynomials of degree at most $(k+1)p$ in t^*. In particular, it is a polynomial in t^* if we restrict our attention to the interval $t^* \in [-k, -k+2]$. We wish to argue that this is not the zero polynomial when p is odd. To prove this, it suffices to show that the coefficient of $(t^*)^{(k+1)p}$ is non-zero, which we do below by studying a matrix whose determinant is this coefficient (when p is odd).

We first show an easy claim concerning matrices that can be written as sums of the identity plus a certain kind of rank-one matrix.

Claim 4.5. *For any matrix $A \in \mathbb{R}^{d \times d}$ with constant columns given by $A_{i,j} = a_j$ for some $a_0, \ldots, a_{d-1} \in \mathbb{R}$ and any $\lambda \in \mathbb{R}$,*

$$\det(A - \lambda I_d) = (-\lambda)^{d-1}\Big(\sum_j a_j - \lambda\Big) \, .$$

Proof. Notice that A is a rank-one stochastic matrix with one non-zero eigenvalue given by $\sum a_j$. Therefore, the characteristic polynomial of A is $\det(A - \lambda I_d) = (-\lambda)^{d-1}(\sum a_j - \lambda)$, as needed. $\qquad\square$

Lemma 4.6. *For an integer $k \geq 1$ and an odd integer $p \geq 1$, the function $t^* \mapsto \det(M_k(p, t^*))$, where $M_k(p, t^*)$ is defined as in Eq. (10), is a polynomial of degree at most $(k+1)p$ when restricted to the interval $t^* \in [-k, -k+2]$. Furthermore, the coefficient of $(t^*)^{(k+1)p}$ of this polynomial is exactly $2^k(2 - 2^k)$.*

In particular, $t^* \mapsto \det(M_k(p, t^*))$ is a non-zero polynomial of degree $(k+1)p$ on the interval $t^* \in [-k, -k+2]$ for $k \geq 2$.

Proof. For any $t^* \in [-k, -k+2]$, the matrix $M_k(p, t^*)$ is given by

$$M_k(p, t^*)_{i,j} = \sum_{\ell=0}^{k} \binom{i}{\ell}\binom{k-i}{j-\ell} \cdot |2i + 2j - k - 4\ell - t^*|^p = \sum_{\ell=0}^{k} \delta_{i+j-2\ell}\binom{i}{\ell}\binom{k-i}{j-\ell} \cdot (2i + 2j - k - 4\ell - t^*)^p \, ,$$

19

where $\delta_r = -1$ if $r = 0$ and 1 otherwise. (Here, we have used the fact that $\binom{i}{\ell}\binom{k-i}{j-\ell}$ is only non-zero when $\ell \leq \min(i,j)$. Therefore, $2i + 2j \geq 4\ell$, so that $2i + 2j - k - 4\ell - t^* \geq 2i + 2j - 4\ell - 2 \geq 0$ unless $2i - 2j - 4\ell = 0$.)

The coefficient of $(t^*)^{(k+1)p}$ in the polynomial $t^* \mapsto \det(M_k(p, t^*))$ is therefore given by $\det(M')$, where M' is defined as

$$
\begin{aligned}
(M')_{i,j} &:= -\sum_{\ell=0}^{k} \delta_{2\ell-i-j} \binom{i}{\ell}\binom{k-i}{j-\ell} \\
&= 2\binom{i}{(i+j)/2}\binom{k-i}{(j-i)/2} - \sum_{\ell=0}^{k}\binom{i}{\ell}\binom{k-i}{j-\ell} \\
&= 2\binom{i}{(i-j)/2}\binom{k-i}{(j-i)/2} - \binom{k}{j} ,
\end{aligned}
$$

where we have again applied Vandermonde's identity. Notice that the first term is non-zero if and only if $i = j$, in which case it is equal to 2. In other words, $M' = A + 2I_{k+1}$, where $A_{i,j} := -\binom{k}{j}$. The result then follows from Claim 4.5. $\qquad\square$

Corollary 4.7. *There is an efficient algorithm that takes as input an integer $k \geq 2$ and odd integer $p \geq 1$ and outputs $t^* \in \mathbb{Q}$ such that $\det(M_k(p, t^*)) \neq 0$, with $M_k(p, t^*)$ defined as in Eq. (10).*

Proof. The algorithm works as follows. It chooses $(k+1)p + 1$ distinct points $t_0, \ldots, t_{(k+1)p} \in [-k, -k+2]$ arbitrarily. (E.g., it chooses $t_i = -k + 2i/((k+1)p)$.) For each t_i, it computes $\det(M_k(p, t_i))$. It outputs the first t_i such that the determinant is non-zero.

We claim that $\det(M_k(p, t_i)) \neq 0$ for at least one index i. Indeed, by Lemma 4.6, $t^* \mapsto \det(M_k(p, t^*))$ is a non-zero polynomial of degree $(k+1)p$. The result then follows from the fact that such a polynomial can have at most $(k+1)p$ roots. $\qquad\square$

Theorem 1.1 for finite p now follows immediately from Theorems 3.2 together with Proposition 4.4, and Corollary 4.7.

4.2   Limitations of the approach

In the previous section, we showed that for every odd $p \geq 1$ and every integer $k \geq 2$, there exists a (p, k)-isolating parallelepiped. This allowed us to conclude that CVP_p is SETH-hard for odd values of p. Now, we show that this approach necessarily fails for even $p \geq 2$. Namely, we show that for every even p, there is no (p, k)-isolating parallelepiped for any $k > p$.[7] For simplicity, we show this for $p = 2$, but a straightforward generalization works for all even p.

Lemma 4.8. *For any integer $k \geq 3$ and vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k, \boldsymbol{t}^* \in \mathbb{R}^{d^*}$, we have*

$$
\sum_{S \subseteq [k]} (-1)^{|S|} \left\| \boldsymbol{t}^* - \sum_{i \in S} \boldsymbol{v}_i \right\|_2^2 = 0 .
$$

---

[7]When $k \leq p$, it is possible to construct (p, k)-isolating parallelepiped for even p. See, e.g., Figure 1.

Proof. We have

$$\sum_{S\subseteq[k]}(-1)^{|S|}\left\|\boldsymbol{t}^*-\sum_{i\in S}\boldsymbol{v}_i\right\|_2^2 = \sum_{S\subseteq[k]}(-1)^{|S|}\left(\|\boldsymbol{t}^*\|_2^2-2\left\langle\boldsymbol{t}^*,\sum_{i\in S}\boldsymbol{v}_i\right\rangle+\left\|\sum_{i\in S}\boldsymbol{v}_i\right\|_2^2\right)$$

$$= \|\boldsymbol{t}^*\|_2^2\cdot\sum_{S\subseteq[k]}(-1)^{|S|}-2\sum_{i\in[k]}\langle\boldsymbol{t}^*,\boldsymbol{v}_i\rangle\cdot\sum_{S\ni i}(-1)^{|S|}$$

$$+\sum_{i\in[k]}\|\boldsymbol{v_i}\|_2^2\cdot\sum_{S\ni i}(-1)^{|S|}+2\sum_{i<j}\langle\boldsymbol{v}_i,\boldsymbol{v}_j\rangle\cdot\sum_{S\ni i,j}(-1)^{|S|}$$

$$= \|\boldsymbol{t}^*\|_2^2\cdot 0-2\sum_{i\in[k]}\langle\boldsymbol{t}^*,\boldsymbol{v}_i\rangle\cdot 0$$

$$+\sum_{i\in[k]}\|\boldsymbol{v_i}\|_2^2\cdot 0+2\sum_{i<j}\langle\boldsymbol{v}_i,\boldsymbol{v}_j\rangle\cdot 0$$

$$= 0\;,$$

where the penultimate equality uses the fact that

$$\sum_{S\subseteq[n]}(-1)^{|S|}=(1-1)^n=0$$

for $n\geq 1$. $\qquad\square$

Corollary 4.9. *There is no $(2,k)$-isolating parallelepiped for any integer $k\geq 3$.*

Proof. Assume $V=[\boldsymbol{v}_1,\dots,\boldsymbol{v}_k]\in\mathbb{R}^{d^*\times k}$ and $\boldsymbol{t}^*\in\mathbb{R}^d$ form a $(2,k)$-isolating parallelepiped. For any $S\neq\emptyset$, $\|\boldsymbol{t}^*-\sum_{i\in S}\boldsymbol{v}_i\|_2^2=1$ by the definition of an isolating parallelepiped. Thus, applying Lemma 4.8, we have

$$\|\boldsymbol{t}^*\|_2^2=\sum_{\emptyset\neq S\subseteq[k]}(-1)^{|S|+1}\left\|\boldsymbol{t}^*-\sum_{i\in S}\boldsymbol{v}_i\right\|_2^2=1\;,$$

which contradicts the assumption that V and $\boldsymbol{t}^*$ form an isolating parallelepiped. $\qquad\square$

4.3   Extending our result to almost all p

We now wish to extend Theorem 1.1 to arbitrary $p\geq 1$. Unfortunately, we know that we cannot do this for *all* p, since we showed in Section 4.2 that no such construction is possible when p is an even integer. However, we show a construction that works for "almost all values of p." In particular, for any fixed k, the construction works for all but finitely many choices of p. We also observe that this implies that, for every fixed p_0,k, there is an $\varepsilon>0$ such that the construction works for every $p\in(p_0-\varepsilon)$ or $p\in(p_0+\varepsilon)$. In particular, for any non-zero $\delta=\delta(n)=o(1)$, the construction works for $p=p_0+\delta(n)$ for sufficiently large integers n.

In Section 4.1, we observed that the function $t^*\mapsto\det(M_k(p,t^*))$ is a piecewise polynomial when p is an integer. This is what allowed us to analyze this case relatively easily (in both Section 4.1 and in Section 4.2). For non-integer p, the function $t^*\mapsto\det(M_k(p,t^*))$ is much less nice. So, instead of holding p fixed and varying t^*, we will be interested in studying the function $f_{k,t^*}(p):=\det(M_k(p,t^*))$ for fixed t^* and k. We first observe that this function has a fairly nice structure.

Lemma 4.10. *For any $t^* \in \mathbb{R}$, integer $k \geq 1$, and $p \geq 1$, let*

$$f_{k,t^*}(p) := \det(M_k(p, t^*)) , \tag{12}$$

where $M_k(p, t^*)$ is as defined in Eq. (10). Then, for fixed k, t^*, $f_{k,t^*}(p)$ is a Dirichlet polynomial. I.e., there are some real numbers $b_0, \ldots, b_r, c_0, \ldots, c_r \in \mathbb{R}$ (depending on t^* and k) such that

$$f_{k,t^*}(p) = \sum_{i=0}^{r} b_i \exp(c_i p) \tag{13}$$

for some finite r.

Proof. To see that $f_{k,t^*}(p)$ is a Dirichlet polynomial for fixed t^*, k, it suffices to note that (1) each entry of $M_k(p, t^*)$ is a Dirichlet polynomial; (2) the determinant of a matrix can be written as a polynomial in the coordinates; and (3) a polynomial of Dirichlet polynomials is itself a Dirichlet polynomial. $\square$

Corollary 4.11. *There is an efficient algorithm that takes as input $k \geq 2$ and any $p \geq 1$ and either fails or outputs $V \in \mathbb{R}^{2^k \times k}$ and $\boldsymbol{t}^* \in \mathbb{R}^{2^k}$ that define a (p, k)-isolating parallelepiped. Furthermore, for any fixed $k \geq 2$, the algorithm only fails for finitely many choices of $p \geq 1$.*

Proof. By Corollary 4.7, for any $k \geq 2$, we can find a $t^* \in \mathbb{Q}$ such that, say, $f_{k,t^*}(1) \neq 0$, where $f_{k,t^*}(p)$ is defined as in Eq. (12). Clearly, $f_{k,t^*}(p)$ is non-zero as a function of p for these values of t^*, k. Furthermore, by Lemma 4.10, $f_{k,t^*}(p)$ is a Dirichlet polynomial. The result follows by the fact that any non-zero Dirichlet polynomial has finitely many roots (see, e.g., Theorem 3.1 in [Jam06]). $\square$

Theorem 4.12. *There is an efficient algorithm that takes as input an integer $k \geq 2$ and any $p \geq 1$ and either fails or outputs $V \in \mathbb{R}^{2^k \times k}$ and $\boldsymbol{t}^* \in \mathbb{R}^{2^k}$ that define a (p, k)-isolating parallelepiped. Furthermore, for any fixed $k \geq 2$, the algorithm only fails on finitely many values of $p \geq 1$.[8]*

Proof. The result follows immediately from Proposition 4.4 and Corollary 4.11. $\square$

Item 2 of Theorem 1.2 now follows immediately from Theorem 3.2 and Theorem 4.12.

We now provide what amounts to a different interpretation of the above.

Lemma 4.13. *There is an efficient algorithm that takes as input any $p_0 \geq 1$ and an integer $k \geq 2$ and outputs a value t^* such that $f_{k,t^*}(p_0 + \delta)$ and $f_{k,t^*}(p_0 - \delta)$ are non-zero for sufficiently small $\delta > 0$, where $f_{k,t^*}(p)$ is as defined in Eq. (12).*

Proof. The algorithm simply calls the procedure from Corollary 4.7 with, say, $p = 1$ and outputs the result. I.e., it suffices to choose any t^* such that $f_{k,t^*}(1) \neq 0$. As in the proof of Corollary 4.11, we observe that the function $f_{k,t^*}(p)$ is zero on a finite set of values X. The result then follows by taking $\delta := \min_{x \in X \setminus \{p_0\}} |x - p_0|/2$ if $X \setminus \{p_0\}$ is non-empty, and $\delta := c$ for any $c > 0$ otherwise. $\square$

Finally, we derive the main theorem of this section.

---

[8] As we observed in Section 4.2, the set of failure points necessarily includes all even integers $p \leq k - 1$.

Theorem 4.14. *For any efficiently computable $\delta(n) \neq 0$ that converges to zero as $n \to \infty$ and $p_0 \geq 1$, there is an efficient algorithm that takes as input an integer $k \geq 2$ and sufficiently large positive integer n and outputs a matrix $V \in \mathbb{R}^{2^k \times k}$ and vector $\boldsymbol{t}^* \in \mathbb{R}^{2^k}$ that define a $(p_0 + \delta(n), k)$-isolating parallelepiped.*

Proof. The result follows immediately from Proposition 4.4 and Lemma 4.13. In particular, the algorithm runs the procedure from Lemma 4.13, receiving as output some $t^* \in \mathbb{R}$ such that $f_{k,t^*}(p_0 \pm \varepsilon)$ is non-zero for sufficiently small $\varepsilon > 0$. In particular, if n is sufficiently large, then $f_{k,t^*}(p_0 + \delta(n))$ will be non-zero. The result then follows from Proposition 4.4. $\qquad\square$

Item 3 of Theorem 1.2 now follows from Theorem 3.2 and Theorem 4.14.

5 Gap-ETH-based Hardness of Approximation

In this section we prove Gap-ETH-based hardness of approximation for CVP_p for every $p \in [1, \infty)$. We also show a stronger hardness of approximation result for CVP_1. (In Section 6.3, we additionally show a stronger result for $p = \infty$.) The main idea behind our proof is to use the reduction from Max-2-SAT to CVP described in Section 1.5 without the "identity matrix gadget" that we used to force any closest vector to be a 0-1 combination of basis vectors. We will show that this is permissible because the resulting CVP instance always has a 0-1 combination of basis vectors that is at least as close to the target as any other lattice vector.

Theorem 5.1. *For every $0 \leq \delta < \varepsilon \leq 1$ and every $p \in [1, \infty)$, there is a polynomial-time reduction from any (δ, ε)-Gap-2-SAT instance with n variables to an instance of γ-CVP_p of rank n, where*

$$\gamma := \left(\frac{\delta + (1 - \delta) \cdot 3^p}{\varepsilon + (1 - \varepsilon) \cdot 3^p} \right)^{1/p} > 1 .$$

In particular, when $\varepsilon = 1$, the corresponding approximation factor is $(\delta + (1 - \delta) \cdot 3^p)^{1/p}$.

Proof. Given a (δ, ε)-Gap-2-SAT instance with n variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$, we construct a CVP_p instance $(B, \boldsymbol{t}, r)$ for some fixed $p \in [1, \infty)$ as follows. Let $B \in \mathbb{Z}^{m \times n}$ be the basis defined by

$$b_{i,j} := \begin{cases} 2 & \text{if } C_i \text{ contains } x_j, \\ -2 & \text{if } C_i \text{ contains } \neg x_j, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\boldsymbol{t} \in \mathbb{R}^m$ be the target vector defined by $t_i := 3 - 2|N_i|$, and let $r := (\varepsilon m + (1 - \varepsilon)m \cdot 3^p)^{1/p}$.

We claim that there is always a 0-1 combination of basis vectors which is a closest vector to $\boldsymbol{t}$. Assuming this claim, we analyze $\|B\boldsymbol{y} - \boldsymbol{t}\|_p$ only for $\boldsymbol{y} \in \{0, 1\}^n$ without loss of generality, while deferring the proof of the claim until the end.

Let $\boldsymbol{y} \in \{0, 1\}^n$ be an assignment to the variables of Φ. For every $1 \leq i \leq m$,

$$\begin{aligned} |(B\boldsymbol{y} - \boldsymbol{t})_i| &= \left| 2 \Big(\sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} - \sum_{s \in N_i} y_{\mathrm{ind}(\ell_{i,s})} \Big) - (3 - 2|N_i|) \right| \\ &= \left| 2 \Big(\sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} + \sum_{s \in N_i} (1 - y_{\mathrm{ind}(\ell_{i,s})}) \Big) - 3 \right| \qquad (14) \\ &= \left| 2 \cdot |S_i(\boldsymbol{y})| - 3 \right| . \end{aligned}$$

The last expression is equal to 1 if $\boldsymbol{y}$ satisfies C_i (i.e. if $|S_i(\boldsymbol{y})| \geq 1$) and 3 if not. We therefore have

$$\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p = \sum_{i=1}^m |(B\boldsymbol{y} - \boldsymbol{t})_i|^p = m^+(\boldsymbol{y}) + (m - m^+(\boldsymbol{y})) \cdot 3^p.$$

Therefore, if $\text{val}(\Phi) \geq \varepsilon$ then there exists $\boldsymbol{y} \in \{0,1\}^n$ such that $\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p \leq \varepsilon m + (1-\varepsilon)m \cdot 3^p = r^p$, and if $\text{val}(\Phi) < \delta$, then for every $\boldsymbol{y} \in \{0,1\}^n$, it holds that $\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p > \delta m + (1-\delta)m \cdot 3^p = \gamma^p r^p$. It follows that the reduction achieves the claimed approximation factor of γ.

It remains to prove the claim that there is always a 0-1 combination of basis vectors which is a closest vector to $\boldsymbol{t}$. We show this by demonstrating that for every $\boldsymbol{y} \in \mathbb{Z}^n$ there exists $\chi(\boldsymbol{y}) \in \{0,1\}^n$ such that $\|B \cdot \chi(\boldsymbol{y}) - \boldsymbol{t}\|_p \leq \|B\boldsymbol{y} - \boldsymbol{t}\|_p$. Given $\boldsymbol{y} \in \mathbb{Z}^n$, let $\chi(\boldsymbol{y}) \in \{0,1\}^n$ denote the vector whose ith coordinate is set to 1 if $y_i \geq 1$ and is set to 0 otherwise.

Fix $\boldsymbol{y} \in \mathbb{Z}^n$ and $1 \leq i \leq n$, and refer to Equation (14). If $\chi(\boldsymbol{y})$ satisfies C_i then $|(B\boldsymbol{y} - \boldsymbol{t})_i| \geq |(B \cdot \chi(\boldsymbol{y}) - \boldsymbol{t})_i| = 1$. On the other hand, if $\chi(\boldsymbol{y})$ does not satisify C_i, then $y_{\text{ind}(\ell_{i,s})} \leq 0$ for all $s \in P_i$ and $1 - y_{\text{ind}(\ell_{i,s})} \leq 0$ for all $s \in N_i$ so that $|(B\boldsymbol{y} - \boldsymbol{t})_i| \geq 3$ and therefore $|(B\boldsymbol{y} - \boldsymbol{t})_i| \geq |(B \cdot \chi(\boldsymbol{y}) - \boldsymbol{t})_i| = 3$. Combining these cases it follows that for all $\boldsymbol{y} \in \mathbb{Z}^n$ and $1 \leq i \leq n$, $|(B \cdot \chi(\boldsymbol{y}) - \boldsymbol{t})_i| \leq |(B\boldsymbol{y} - \boldsymbol{t})_i|$, and therefore $\|B \cdot \chi(\boldsymbol{y}) - \boldsymbol{t}\|_p \leq \|B\boldsymbol{y} - \boldsymbol{t}\|_p$, proving the claim. $\square$

Corollary 5.2. *For every $p \in [1, \infty)$ and $0 < \delta < \delta' < 1$, there is no $2^{o(n)}$-time algorithm for γ-CVP_p with*

$$\gamma := \left(\frac{6 + \delta' + (4 - \delta') \cdot 3^p}{7 + 3^{p+1}} \right)^{1/p} > 1$$

unless randomized Gap-ETH (with respect to $(\delta, 1)$-Gap-3-SAT) fails.

Proof. Concatenate the reductions described in Proposition 2.11, Proposition 2.12, and Theorem 5.1.[9]
$\square$

5.1 A stronger result for ℓ_1

In Theorem 5.1 we showed that there was no need to use the identity matrix gadget in our reduction from Gap-2-SAT to CVP. An interesting question is whether the identity matrix gadget is also unnecessary in our reduction in Theorem 3.2 from Gap-k-SAT to CVP_p using (p, k)-isolating parallelepipeds. If so, then we get stronger "Gap-SETH-hardness" for CVP_p for all p for which there exist (p, k)-isolating parallelepipeds for infinitely many $k \in \mathbb{Z}^+$.

Although we do not know how to show this in general, we are able to get stronger hardness of approximation for CVP_1 in this way by using the family of $(1, k)$-isolating parallelepipeds described in Corollary 3.4. The analysis is similar to the analysis in Theorem 5.1. In particular, a 0-1 combination of basis vectors will always be at least as close to the target vector as any other lattice vector will.

Theorem 5.3. *For every $0 \leq \delta < \varepsilon \leq 1$, there is a polynomial-time reduction from any (δ, ε)-Gap-k-SAT instance with n variables to an instance of γ-CVP_1 of rank n, where*

$$\gamma := \left(\frac{\delta \cdot (k-1) + (1-\delta) \cdot (k+1)}{\varepsilon \cdot (k-1) + (1-\varepsilon) \cdot (k+1)} \right) > 1.$$

---

[9]Note that the Gap-2-SAT instance output by Proposition 2.12 contains 1-clauses as well as 2-clauses. We therefore stress that the reduction described in Theorem 5.1 works for this case as well.

Proof. Given a (δ, ε)-Gap-k-SAT instance with n variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$, we construct a CVP$_1$ instance $(B, \boldsymbol{t}, r)$ as follows. The basis $B \in \mathbb{Z}^{(2m) \times n}$ and target vector $\boldsymbol{t} \in \mathbb{Z}^{2m}$ in the output instance have the form

$$B = \begin{pmatrix} B_1 \\ \vdots \\ B_m \end{pmatrix}, \qquad \boldsymbol{t} = \begin{pmatrix} \boldsymbol{t}_1 \\ \vdots \\ \boldsymbol{t}_m \end{pmatrix},$$

with blocks $B_i \in \mathbb{Z}^{2 \times n}$ and $\boldsymbol{t}_i \in \mathbb{Z}^2$ for $1 \leq i \leq m$. For every $1 \leq i \leq m$ and $1 \leq j \leq n$, set the jth column $(B_i)_j$ of block B_i (corresponding to the clause $C_i = \vee_{s=1}^k \ell_{i,s}$) as

$$(B_i)_j := \begin{cases} (1,1)^T & \text{if } C_i \text{ contains } x_j, \\ -(1,1)^T & \text{if } C_i \text{ contains } \neg x_j, \\ (0,0)^T & \text{otherwise,} \end{cases}$$

and set $\boldsymbol{t}_i := (k,1)^T - |N_i| \cdot (1,1)^T$. Set $r := \varepsilon m(k-1) + (1-\varepsilon) \cdot m(k+1)$.

We claim that there is always a 0-1 combination of basis vectors that is closest to $\boldsymbol{t}$. Assuming the claim, we analyze $\|B\boldsymbol{y} - \boldsymbol{t}\|$ only for $\boldsymbol{y} \in \{0,1\}^n$ without loss of generality, while deferring the proof of the claim until the end.

Let $\boldsymbol{y} \in \{0,1\}^n$ be an assignment to the variables of Φ. For every $1 \leq i \leq m$,

$$\|B_i\boldsymbol{y} - \boldsymbol{t}_i\|_1 = \left\| \sum_{s \in P_i} y_{\text{ind}(\ell_{i,s})} \cdot (1,1)^T - \sum_{s \in N_i} y_{\text{ind}(\ell_{i,s})} \cdot (1,1)^T - \left((k,1)^T - |N_i| \cdot (1,1)^T \right) \right\|_1$$

$$= \left\| \sum_{s \in P_i} y_{\text{ind}(\ell_{i,s})} \cdot (1,1)^T + \sum_{s \in N_i} \left(1 - y_{\text{ind}(\ell_{i,s})} \right) \cdot (1,1)^T - (k,1)^T \right\|_1 \qquad (15)$$

$$= \left\| |S_i(\boldsymbol{y})| \cdot (1,1)^T - (k,1)^T \right\|_1.$$

The last expression is equal to $k-1$ if $\boldsymbol{y}$ satisfies C_i (i.e. if $|S_i(\boldsymbol{y})| \geq 1$) and $k+1$ if not. We therefore have

$$\|B\boldsymbol{y} - \boldsymbol{t}\|_1 = \sum_{i=1}^m \|B_i\boldsymbol{y} - \boldsymbol{t}_i\|_1 = m^+(\boldsymbol{y}) \cdot (k-1) + (m - m^+(\boldsymbol{y})) \cdot (k+1).$$

Therefore, if $\text{val}(\Phi) \geq \varepsilon$ then there exists $\boldsymbol{y} \in \{0,1\}^n$ such that $\|B\boldsymbol{y} - \boldsymbol{t}\|_1 \leq \varepsilon m(k-1) + (1 - \varepsilon) \cdot m(k+1) = r$, and if $\text{val}(\Phi) < \delta$ then for every $\boldsymbol{y} \in \{0,1\}^n$, it holds that $\|B\boldsymbol{y} - \boldsymbol{t}\|_1 > \delta m(k-1) + (1-\delta)m(k+1) = \gamma r$. It follows that the reduction achieves the claimed approximation factor of γ.

It remains to prove the claim that there is always a 0-1 combination of basis vectors that is a closest vector to $\boldsymbol{t}$. We show this by demonstrating that for every $\boldsymbol{y} \in \mathbb{Z}^n$ there exists $\chi(\boldsymbol{y}) \in \{0,1\}^n$ such that $\|B\chi(\boldsymbol{y}) - \boldsymbol{t}\|_1 \leq \|B\boldsymbol{y} - \boldsymbol{t}\|_1$. Given $\boldsymbol{y} \in \mathbb{Z}^n$, let $\chi(\boldsymbol{y}) \in \{0,1\}^n$ denote the vector whose coordinate is set to 1 if $y_i \geq 1$ and is set to 0 otherwise.

Note that $\|c(1,1)^T - (k,1)^T\|_1 = k-1$ for all $c \in [k]$, and $\|c(1,1)^T - (k,1)^T\|_1 \geq k+1$ for all $c \in \mathbb{Z} \setminus [k]$. Fix $\boldsymbol{y} \in \mathbb{Z}^n$ and $1 \leq i \leq n$, and refer to Equation (15). If $\chi(\boldsymbol{y})$ satisfies C_i then $\|B_i \cdot \boldsymbol{y} - \boldsymbol{t}_i\|_1 \geq \|B_i \cdot \chi(\boldsymbol{y}) - \boldsymbol{t}_i\|_1 = k-1$. On the other hand, if $\chi(\boldsymbol{y})$ does not satisfy C_i, then $y_{\text{ind}(\ell_{i,s})} \leq 0$ for all $s \in P_i$ and $1 - y_{\text{ind}(\ell_{i,s})} \leq 0$ for all $s \in N_i$ so that $\|B_i \cdot \boldsymbol{y} - \boldsymbol{t}_i\|_1 \geq$

25

$\|B_i \cdot \chi(\boldsymbol{y}) - \boldsymbol{t}_i\|_1 = k + 1$. Combining these cases it follows that for all $\boldsymbol{y} \in \mathbb{Z}^n$ and $1 \leq i \leq n$, $\|B_i \cdot \chi(\boldsymbol{y}) - \boldsymbol{t}_i\|_1 \leq \|B_i \cdot \boldsymbol{y} - \boldsymbol{t}_i\|_1$, and therefore $\|B \cdot \chi(\boldsymbol{y}) - \boldsymbol{t}\|_1 \leq \|B\boldsymbol{y} - \boldsymbol{t}\|_1$, proving the claim. $\qquad\square$

6 Additional hardness results

In this section we prove a number of additional results about the quantitative hardness of CVP and related problems. In Section 6.1, we give a reduction from Max-2-SAT to CVPP_p for all $p \in [1, \infty)$, proving Theorem 1.4. In Section 6.2, we give a reduction from Max-k-SAT (and in particular Max-2-SAT) to CVP_p for all $p \in [1, \infty)$, proving Theorem 1.5. Finally, in Section 6.3 we give a reduction from k-SAT to CVP_∞ and SVP_∞, proving the special case of Theorem 1.1 for $p = \infty$.

Our reductions all use the same high-level idea as the reduction given in Theorem 3.2, but each uses new ideas as well. Throughout this section we adopt the notation from Section 3.

6.1 Hardness of CVPP_p

In this section, we prove ETH-hardness of CVPP. To do this, for every n, we construct a single lattice $\mathcal{L}_n \subset \mathbb{R}^d$ of rank $O(n^2)$, such that for every n-variable instance of Max-2-SAT, there exists an efficiently computable $\boldsymbol{t} \in \mathbb{R}^d$ that is close to the lattice if and only if Φ is satisfiable. Clearly, any efficient algorithm for CVPP on this lattice would imply a similarly efficient algorithm for Max-2-SAT (and also 3-SAT, as described below).

Our basis B_n for $\mathcal{L}_n$ will encode all possible $O(n^2)$ clauses of a Max-2-SAT instance on n variables, together with a gadget that will allow us to "switch on or off" each clause by only changing the coordinates of the *target vector* $\boldsymbol{t}$. (This gadget costs us a quadratic blow-up in the lattice rank.) Then, given an instance (Φ, W) of Max-2-SAT, we define the target vector $\boldsymbol{t}$ such that it "switches on" all clauses from Φ and "switches off" all the remaining clauses.

Lemma 6.1. *For every $p \in [1, \infty)$, there is a pair of polynomial-time algorithms (P, Q) (in analogy to the definition of* CVPP*) that behave as follows.*

1. *On input an integer $n \geq 1$, P outputs a basis $B_n \in \mathbb{R}^{d \times N}$ of a rank N lattice $\mathcal{L}_n \subset \mathbb{R}^d$, where $d = d(n) = O(n^2)$ and $N = N(n) = O(n^2)$.*

2. *On input a Max-2-SAT instance with n variables, Q outputs a target vector $\boldsymbol{t} \in \mathbb{R}^d$ and a distance bound $r \geq 0$ such that $\mathrm{dist}_p(\boldsymbol{t}, \mathcal{L}_n) \leq r$ if and only if the input is a 'YES' instance.*

Proof. Let $M = 4\binom{n}{2} = O(n^2)$ be the total possible number of 2-clauses on n variables, and let $C_1, \ldots, C_M$ denote those clauses.

The algorithm P constructs the basis $B_n \in \mathbb{R}^{d \times N}$, where $d := n + 2M, N := n + M$, as

$$B := \begin{pmatrix} (\boldsymbol{b}_1^T, \boldsymbol{c}_1^T) \\ \vdots \\ (\boldsymbol{b}_M^T, \boldsymbol{c}_M^T) \\ 2\alpha^{1/p} I_N \end{pmatrix},$$

with rows $(\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)$ of B satisfying $\boldsymbol{b}_i \in \mathbb{R}^n$ and $\boldsymbol{c}_i \in \mathbb{R}^M$ for $1 \leq i \leq M$, and where $\alpha := 2^p M$. For every $1 \leq i \leq M$ and $1 \leq j \leq n$, set the jth coordinate of $\boldsymbol{b}_i$ (corresponding to the clause $C_i = \ell_{i,1} \lor \ell_{i,2}$) as

$$(\boldsymbol{b}_i)_j := \begin{cases} 1 & \text{if } x_j \text{ appears in the } i\text{th clause,} \\ -1 & \text{if } \neg x_j \text{ appears in the } i\text{th clause,} \\ 0 & \text{otherwise.} \end{cases}$$

For every $1 \leq i \leq M$, set $\boldsymbol{c}_i^T := (\boldsymbol{0}_{(i-1)}^T, 1, \boldsymbol{0}_{(M-i)}^T)$, i.e., set $(\boldsymbol{c}_1, \ldots, \boldsymbol{c}_M) = I_M$.

Given an instance (Φ, W) of Max-2-SAT with m clauses, the algorithm Q outputs

$$r := \left((M - m + W) \cdot 1/2^p + (m - W) \cdot (3/2)^p + \alpha(n + M - m)\right)^{1/p}$$

and $\boldsymbol{t} \in \mathbb{R}^d$ defined as

$$\boldsymbol{t} := \begin{pmatrix} u_1 \\ \vdots \\ u_M \\ \alpha^{1/p} \cdot \boldsymbol{1}_n \\ v_1 \\ \vdots \\ v_M \end{pmatrix},$$

where for $1 \leq i \leq M$, $u_i = 3/2 - |N_i|$, and $v_i = 0$ if the clause C_i appears in the formula Φ and $v_i = \alpha^{1/p}$ otherwise.

Clearly both algorithms are efficient. We now analyze for which $\boldsymbol{y} \in \mathbb{Z}^N$ we have $\|B\boldsymbol{y} - \boldsymbol{t}\|_p \leq r$. Note that the vector $\boldsymbol{v} = (v_1, \ldots, v_M)^T$ has exactly m coordinates equal to zero, and $M - m$ coordinates equal to $\alpha^{1/p}$. Given $\boldsymbol{y} \notin \{0, 1\}^N$, we have

$$\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p \geq \|2\alpha^{1/p} I_N \boldsymbol{y} - (\alpha^{1/p}\boldsymbol{1}_n^T, \boldsymbol{v}^T)^T\|_p^p \geq \alpha(n + M - m + 1) \geq 2^p M + \alpha(M + n - m) > r^p.$$

Furthermore, if $\boldsymbol{y} \in \{0, 1\}^N$ has a non-zero coordinate y_{n+M+i} (for $1 \leq i \leq M$) at a position corresponding to a $t_{n+M+i} = 0$ in $\boldsymbol{t}$ (i.e., $C_i \in \Phi$), then again $\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p \geq \alpha(n + M - m + 1) > r^p$. So, we can restrict our attention to $\boldsymbol{y} \in \{0, 1\}^N$ with $y_{n+M+i} = 0$ whenever $C_i \in \Phi$.

Consider an assignment $\boldsymbol{a} \in \{0, 1\}^n$ to the n variables of Φ. Take $(\boldsymbol{y}')^T = (y_1', \ldots y_M')^T \in \{0, 1\}^M$, and set $\boldsymbol{y} := (\boldsymbol{a}^T, (\boldsymbol{y}')^T)^T$. Then, for $1 \leq i \leq M$,

$$\begin{aligned} \left|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y}\rangle - t_i\right| &= \left|\sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} - \sum_{s \in N_i} y_{\mathrm{ind}(\ell_{i,s})} + \langle \boldsymbol{c}_i, \boldsymbol{y}'\rangle - (3/2 - |N_i|)\right| \\ &= \left|\sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} - \sum_{s \in N_i} (1 - y_{\mathrm{ind}(\ell_{i,s})}) + y_i' - 3/2\right| \\ &= \left||S_i(\boldsymbol{a})| + y_i' - 3/2\right|. \end{aligned}$$

If $C_i \notin \Phi$, then there exists $y_i' \in \{0, 1\}$, such that $|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y}\rangle - t_i| = 1/2$. Moreover, the choice of y_i' does not affect $|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y}\rangle - t_{i'}|$ for $i' \neq i$. If $C_i \in \Phi$ and $|S_i(\boldsymbol{a})| > 0$ for $y_i' = 0$, then $|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y}\rangle - t_i| = 1/2$. On the other hand, if $C_i \in \Phi$ and $|S_i(\boldsymbol{a})| = 0$, then $y_i' = 0$ implies $|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y}\rangle - t_i| \geq 3/2$.

Because $|S_i(\boldsymbol{a})| \geq 1$ if and only if C_i is satisfied, we see that the following holds if and only if the number of satisfied clauses $m^+(\boldsymbol{a})$ is at least W:

There exists a $\boldsymbol{y}'$ such that, setting $\boldsymbol{y} := (\boldsymbol{a}, \boldsymbol{y}')$, we have

$$\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p = \Big(\sum_{i=1}^{M} |\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y} \rangle - t_i|^p \Big) + \alpha(n + M - m)$$

$$= (M - m + m^+(\boldsymbol{a})) \cdot (1/2)^p + (m - m^+(\boldsymbol{a})) \cdot (3/2)^p + \alpha(n + M - m)$$

$$\leq (M - m + W) \cdot (1/2)^p + (m - W) \cdot (3/2)^p + \alpha(n + M - m)$$

$$= r^p \ .$$

Therefore, there exists $\boldsymbol{y}$ with $\|B\boldsymbol{y} - \boldsymbol{t}\|_p \leq r$ if and only if (Φ, W) is a 'YES' instance of Max-2-SAT. $\qquad\square$

Proof of Theorem 1.4. The main statement follows from Lemma 6.1. The "in particular" part follows from Lemma 6.1 and the existence of a reduction from 3-SAT with n variables clauses to (many) Max-2-SAT instances with $O(n)$ variables. Indeed, such a reduction follows by applying the Sparsification Lemma (Proposition 2.9), and then reducing each sparse 3-SAT instance to a Max-2-SAT instance with only a linear blow-up in the number of variables (see, e.g., the reduction in [GJS76, Theorem 1.1]). $\qquad\square$

6.2 ETH- and Max-2-SAT-hardness for all $p \in [1, \infty)$

Theorem 6.2. *For every $p = p(n) \in [1, \infty)$ and $k \geq 2$ there is a polynomial-time reduction from any (Weighted-)Max-k-SAT instance with n variables and m clauses to a CVP_p instance of rank $n + (k - 2)m$.*

Proof. For simplicity we give a reduction from unweighted Max-k-SAT. The modification sketched for reducing from Weighted Max-k-SAT in Theorem 3.2 works here as well. Namely, we give a reduction from a Max-k-SAT instance (Φ, W) to an instance $(B, \boldsymbol{t}, r)$ of CVP_p. Here the formula Φ is on n variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$. (Φ, W) is a 'YES' instance if there exists an assignment $\boldsymbol{a}$ such that $m^+ \geq W$. We assume without loss of generality that each variable appears at most once per clause. We define the output CVP_p instance as follows. Let $d := n + (k - 1)m$ and let $N := n + (k - 2)m$. The basis $B \in \mathbb{R}^{d \times N}$ in the output instance has the form

$$B = \begin{pmatrix} (\boldsymbol{b}_1^T, \boldsymbol{c}_1^T) \\ \vdots \\ (\boldsymbol{b}_m^T, \boldsymbol{c}_m^T) \\ 2\alpha^{1/p} \cdot I_N \end{pmatrix}, \qquad \boldsymbol{t} = \begin{pmatrix} t_1 \\ \vdots \\ t_m \\ \alpha^{1/p} \cdot \mathbf{1}_N \end{pmatrix}$$

with rows $(\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)$ of B satisfying $\boldsymbol{b}_i \in \mathbb{R}^n$ and $\boldsymbol{c}_i \in \mathbb{R}^{m(k-2)}$ for $1 \leq i \leq m$, and where $\alpha := W \cdot (\frac{1}{2})^p + (m - W) \cdot (\frac{3}{2})^p$. For every $1 \leq i \leq m$ and $1 \leq j \leq n$, set the jth coordinate of $\boldsymbol{b}_i$ (corresponding to the clause $C_i = \vee_{s=1}^{k} \ell_{i,s}$) as

$$(\boldsymbol{b}_i)_j := \begin{cases} 1 & \text{if } x_j \text{ is in the } i\text{th clause,} \\ -1 & \text{if } \neg x_j \text{ is in the } i\text{th clause,} \\ 0 & \text{otherwise.} \end{cases}$$

For every $1 \leq i \leq m$, set $\boldsymbol{c}_i^T := (\mathbf{0}_{(i-1)(k-2)}^T, -\mathbf{1}_{k-2}^T, \mathbf{0}_{(m-i)(k-2)}^T)$. I.e., each $\boldsymbol{c}_i$ has a block of -1s of length $k-2$, and $\boldsymbol{c}_i, \boldsymbol{c}_{i'}$ are coordinate-wise disjoint for $i \neq i'$. For every $1 \leq i \leq m$ set $t_i := 3/2 - |N_i|$. Finally, set $r := (\alpha(N + 1))^{1/p}$.

We next analyze for which $\boldsymbol{y} \in \mathbb{Z}^N$ it holds that $\|B\boldsymbol{y} - \boldsymbol{t}\|_p \leq r$. Given $\boldsymbol{y} \notin \{0,1\}^N$, $\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p \geq \|2\alpha^{1/p}I_N\boldsymbol{y} - \alpha^{1/p}\mathbf{1}_N\|_p^p \geq \alpha(N+2) > r^p$, so we only need to analyze the case when $\boldsymbol{y} \in \{0,1\}^n$.

Consider an assignment $\boldsymbol{a} \in \{0,1\}^n$ to the variables of Φ, take $(\boldsymbol{y}')^T = ((\boldsymbol{y}'_1)^T, \ldots, (\boldsymbol{y}'_m)^T) \in \{0,1\}^{(k-2)m}$ with $\boldsymbol{y}'_i \in \{0,1\}^{k-2}$ for $1 \leq i \leq m$, and set $\boldsymbol{y}^T := (\boldsymbol{a}^T, (\boldsymbol{y}')^T)$. Then, for $1 \leq i \leq m$,

$$
\begin{aligned}
\left| \langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y} \rangle - t_i \right| &= \Big| \sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} - \sum_{s \in N_i} y_{\mathrm{ind}(\ell_{i,s})} + \langle \boldsymbol{c}_i, \boldsymbol{y}'_i \rangle - (3/2 - |N_i|) \Big| \\
&= \Big| \sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} + \sum_{s \in N_i} (1 - y_{\mathrm{ind}(\ell_{i,s})}) - \|\boldsymbol{y}'_i\|_1 - 3/2 \Big| \\
&= \Big| |S_i(\boldsymbol{a})| - \|\boldsymbol{y}'_i\|_1 - 3/2 \Big|.
\end{aligned}
$$

It follows that if $|S_i(\boldsymbol{a})| = 0$ then $|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y} \rangle - t_i| \geq \frac{3}{2}$. On the other hand, if $|S_i(\boldsymbol{a})| \geq 1$, then there exists $\boldsymbol{y}'_i \in \{0,1\}^{k-2}$ such that $|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y} \rangle - t_i| = \frac{1}{2}$. Indeed, picking any $\boldsymbol{y}'_i$ with Hamming weight $|S_i(\boldsymbol{a})| - 2$ or $|S_i(\boldsymbol{a})| - 1$ achieves this. Moreover, the choice of $\boldsymbol{y}'_i$ does not affect $|\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y} \rangle - t_{i'}|$ for $i' \neq i$.

Because $|S_i(\boldsymbol{a})| \geq 1$ if and only if C_i is satisfied, we see that the following holds if and only if the number of satisfied clauses $m^+(\boldsymbol{a})$ is at least W:

There exists a $\boldsymbol{y}'$ such that, setting $\boldsymbol{y} := (\boldsymbol{a}, \boldsymbol{y}')$, we have

$$
\begin{aligned}
\|B\boldsymbol{y} - \boldsymbol{t}\|_p^p &= \Big(\sum_{i=1}^m |\langle (\boldsymbol{b}_i^T, \boldsymbol{c}_i^T)^T, \boldsymbol{y} \rangle - t_i|^p \Big) + \alpha N \\
&= m^+(\boldsymbol{a}) \cdot (1/2)^p + (m - m^+(\boldsymbol{a})) \cdot (3/2)^p + \alpha N \\
&\leq W \cdot (1/2)^p + (m - W) \cdot (3/2)^p + \alpha N \\
&= r^p.
\end{aligned}
$$

Therefore, there exists $\boldsymbol{y}$ such that $\|B\boldsymbol{y} - \boldsymbol{t}\|_p \leq r$ if and only if (Φ, W) is a 'YES' instance of Max-k-SAT. $\qquad\square$

We remark that a straightforward modification of the preceding reduction gives a reduction from an instance of Max-k-SAT with $k \geq 3$ on n variables and m clauses to a CVP_p instance of rank $n + (\lfloor \log_2(k-2) \rfloor + 1)m$ (as opposed to rank $n + (k-2)m$). The idea is to encode the value $k-2$ (corresponding to the row-specific blocks $-\mathbf{1}_{k-2}$ used in the reduction) in binary rather than unary.

Corollary 6.3. *For every $p \in [1, \infty)$ there is no $2^{o(n)}$-time algorithm for CVP_p assuming ETH.*

Proof. The claim follows by combining the Sparsification Lemma (Proposition 2.9) with the reduction in Theorem 6.2.[10] $\qquad\square$

When $k = 2$, the rank of the CVP_p instance output by the reduction in Theorem 6.2 is n. Therefore, we get the following corollary.

Corollary 6.4. *If there exists a $2^{(\omega/3-\varepsilon)n}$-time (resp. $2^{(1-\varepsilon)n}$-time and polynomial space) algorithm for CVP_p for any $p \in [1, \infty)$ and for any constant $\varepsilon > 0$, then there exists a $2^{(\omega/3-\varepsilon)n}$-time (resp. $2^{(1-\varepsilon)n}$-time and polynomial space) algorithm for (Weighted-)Max-2-SAT and (Weighted-)Max-Cut.*

---

[10]As a technical point, we must reduce from k-SAT rather than Max-k-SAT to show hardness under ETH because the Sparsification Lemma works for k-SAT.

6.3 The hardness of SVP_∞ and CVP_∞

Theorem 6.5. *There is a polynomial-time reduction from a k-SAT instance with n variables to a CVP_∞ instance of rank n.*

Proof. We give a reduction from a k-SAT formula Φ with n variables $x_1, \dots, x_n$ and m clauses $C_1, \dots, C_m$ to an instance $(B, \boldsymbol{t}, r)$ of CVP_∞. We assume without loss of generality that each variable appears at most once per clause. We define the output CVP_∞ instance as follows. Let $d := m + n$. The basis $B \in \mathbb{R}^{d \times n}$ in the output instance has the form

$$B = \begin{pmatrix} \boldsymbol{b}_1^T \\ \vdots \\ \boldsymbol{b}_m^T \\ (k+1) \cdot I_n \end{pmatrix}, \qquad \boldsymbol{t} = \begin{pmatrix} t_1 \\ \vdots \\ t_m \\ (k+1)/2 \cdot \boldsymbol{1}_n \end{pmatrix},$$

with rows $\boldsymbol{b}_i$ of B satisfying $\boldsymbol{b}_i \in \mathbb{R}^n$ for $1 \le i \le m$. For every $1 \le i \le m$, set $\boldsymbol{b}_i$ as in the proof of Theorem 6.2 and set $t_i := (k+1)/2 - |N_i|$. Set $r := (k-1)/2$.

We next analyze for which $\boldsymbol{y} \in \mathbb{Z}^n$ it holds that $\|B\boldsymbol{y} - \boldsymbol{t}\|_\infty \le r$. Given $\boldsymbol{y} \notin \{0,1\}^n$, $\|B\boldsymbol{y} - \boldsymbol{t}\|_\infty \ge \|(k-1) \cdot I_n \boldsymbol{y} - (k-1)/2 \cdot \boldsymbol{1}_n\|_\infty \ge 3(k-1)/2 > r$, so we only need to analyze the case when $\boldsymbol{y} \in \{0,1\}^n$. Consider an assignment $\boldsymbol{y} \in \{0,1\}^n$ to the variables of Φ. Then

$$\left| \langle \boldsymbol{b}_i, \boldsymbol{y} \rangle - t_i \right| = \left| \sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} - \sum_{s \in N_i} y_{\mathrm{ind}(\ell_{i,s})} - ((k+1)/2 - |N_i|) \right|$$

$$= \left| \sum_{s \in P_i} y_{\mathrm{ind}(\ell_{i,s})} - \sum_{s \in N_i} (1 - y_{\mathrm{ind}(\ell_{i,s})}) - (k+1)/2 \right|$$

$$= \left| |S_i(\boldsymbol{a})| - (k+1)/2 \right|.$$

It follows that if $|S_i(\boldsymbol{a})| = 0$ then $|\langle \boldsymbol{b}_i, \boldsymbol{y} \rangle - t_i| = (k+1)/2$, and otherwise $|\langle \boldsymbol{b}_i, \boldsymbol{y} \rangle - t_i| \le (k-1)/2$. Because $|S_i(\boldsymbol{a})| \ge 1$ if and only if C_i is satisfied, we then have that $\|B\boldsymbol{y} - \boldsymbol{t}\|_\infty = \max\{|\langle \boldsymbol{b}_1, \boldsymbol{y} \rangle - t_1|, \dots, |\langle \boldsymbol{b}_m, \boldsymbol{y} \rangle - t_m|, (k-1)/2\} = (k-1)/2 = r$ if $\boldsymbol{y}$ satisfies Φ, and $\|B\boldsymbol{y} - \boldsymbol{t}\|_\infty = (k+1)/2 > r$ otherwise. Therefore there exists $\boldsymbol{y}$ such that $\|B\boldsymbol{y} - \boldsymbol{t}\|_\infty \le r$ if and only if Φ is satisfiable. $\qquad \square$

Lemma 6.6. *There is a polynomial-time reduction from a k-SAT instance with n variables to an SVP_∞ instance of rank $n + 1$.*

Proof. We give a reduction from a k-SAT formula Φ with n variables $x_1, \dots, x_n$ and m clauses $C_1, \dots, C_m$ to an instance (B', r) of SVP_∞. Define the basis $B' \in \mathbb{R}^{(m+n+1) \times (n+1)}$ as

$$B' := \begin{pmatrix} B & -\boldsymbol{t} \\ \boldsymbol{0}_n^T & -(k-1)/2 \end{pmatrix},$$

where B and $\boldsymbol{t}$ are as defined in the proof of Theorem 6.5, and set $r := (k-1)/2$. We consider for which $\boldsymbol{y} \in \mathbb{Z}^{n+1} \setminus \{\boldsymbol{0}_{n+1}\}$ it holds that $\|B\boldsymbol{y}\|_\infty \le r$. It is not hard to check that if $|y_i| \ge 2$ for some $1 \le i \le n+1$, or if the signs of y_i and y_{n+1} differ for some $1 \le i \le n$, then $\|B\boldsymbol{y}\|_\infty > r$. Therefore we need only consider $\boldsymbol{y}$ of the form $\boldsymbol{y} = \pm(\boldsymbol{a}^T, 1)^T$ where $\boldsymbol{a} \in \{0,1\}^n$. But for such a $\boldsymbol{y}$ we have that $\|B'\boldsymbol{y}\|_\infty = \|B\boldsymbol{a} - \boldsymbol{t}\|_\infty$, and $\|B\boldsymbol{a} - \boldsymbol{t}\|_\infty \le (k-1)/2$ if and only if $\boldsymbol{a}$ is a satisfying assignment to Φ by the analysis in the proof of Theorem 6.5. $\qquad \square$

Corollary 6.7. *For any constant $\varepsilon > 0$ there is no $2^{(1-\varepsilon)n}$-time algorithm for* SVP_∞ *or* CVP_∞ *assuming SETH, and there is no $2^{o(n)}$-time algorithm for* SVP_∞ *or* CVP_∞ *assuming ETH.*

Proof. Combine Theorem 6.5 and Lemma 6.6. □

Note that the preceding reduction in fact achieves an approximation factor of $\gamma = \gamma(k) := 1 + 2/(k-1)$. This implies that for every constant $\varepsilon > 0$, there is a $\gamma_\varepsilon > 1$ such that no $2^{(1-\varepsilon)n}$-time algorithm that approximates SVP_∞ or CVP_∞ to within a factor of γ_ε unless SETH fails.

Finally, we remark that the reduction given in Theorem 6.2 is *parsimonious* when used as a reduction from 2-SAT. I.e., there is a one-to-one correspondence between satisfying assignments in the input instance and close vectors in the output instance. The reductions given in Theorem 6.5 and Lemma 6.6 are also parsimonious.[11]

Because #2-SAT is #P-hard, our reductions therefore show that the counting version of CVP_p (called the Vector Counting Problem) is #P-hard for all $1 \le p \le \infty$, and that the counting version of SVP_∞ is #P-hard. This improves (and arguably simplifies the proof of) a result of Charles [Cha07], which showed that the counting version of CVP_2 is #P-hard.

Acknowledgments

References

[ABSS93] Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. In *FOCS*, 1993.

[ABW15] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In *FOCS*, 2015.

[ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange — A new hope. In *USENIX Security Symposium*, 2016.

[ADRS15] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the Shortest Vector Problem in 2^n time via discrete Gaussian sampling. In *STOC*, 2015.

[ADS15] Divesh Aggarwal, Daniel Dadush, and Noah Stephens-Davidowitz. Solving the Closest Vector Problem in 2^n time— The discrete Gaussian strikes again! In *FOCS*, 2015.

[AJ08] V. Arvind and Pushkar S. Joglekar. Some sieving algorithms for lattice problems. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, 2008.

[Ajt98] Miklós Ajtai. The Shortest Vector Problem in L2 is NP-hard for randomized reductions. In *STOC*, 1998.

---

[11] Actually, in the case of SVP_∞, each satisfying assignment to the SAT formula corresponds to a pair $\pm \boldsymbol{v}$ of shortest non-zero vectors, so that there are exactly twice as many such vectors as there are satisfying assignments.

[Ajt04]   Miklós Ajtai. Generating hard instances of lattice problems. In *Complexity of compu-tations and proofs*, volume 13 of *Quad. Mat.*, pages 1–32. Dept. Math., Seconda Univ. Napoli, Caserta, 2004. Preliminary version in STOC'96.

[AKKV11]   Mikhail Alekhnovich, Subhash Khot, Guy Kindler, and Nisheeth K. Vishnoi. Hardness of approximating the Closest Vector Problem with Pre-processing. *Computational Complexity*, 20, 2011.

[AKS01]   Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the Shortest Lattice Vector Problem. In *STOC*, 2001.

[AKS02]   Miklós Ajtai, Ravi Kumar, and D. Sivakumar. Sampling short lattice vectors and the Closest Lattice Vector Problem. In *CCC*, 2002.

[APS15]   Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of Learning with Errors. *J. Mathematical Cryptology*, 9(3):169–203, 2015.

[AR05]   Dorit Aharonov and Oded Regev. Lattice problems in NP ∩ coNP. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS 2004.

[BCD+16]   Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, quantum-secure key exchange from LWE. In *CCS*, 2016.

[BD15]   Nicolas Bonifas and Daniel Dadush. Short paths on the Voronoi graph and the Closest Vector Problem with Preprocessing. In *SODA*, 2015.

[BDGL16]   Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In *SODA*, 2016.

[BI15]   Arturs Backurs and Piotr Indyk. Edit Distance cannot be computed in strongly subquadratic time (unless SETH is false). In *STOC*, 2015.

[BN09]   Johannes Blömer and Stefanie Naewe. Sampling methods for shortest vectors, closest vectors and successive minima. *Theoret. Comput. Sci.*, 410(18):1648–1665, 2009.

[CDL+12]   Marek Cygan, Holger Dell, Daniel Lokshtanov, Dániel Marx, Jesper Nederlof, Yoshio Okamoto, Ramamohan Paturi, Saket Saurabh, and Magnus Wahlström. On problems as hard as CNF-SAT. In *CCC*, 2012.

[CFK+15]   Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.

[Cha07]   Denis Xavier Charles. Counting lattice vectors. *J. Comput. Syst. Sci.*, 73(6):962–972, 2007.

[CLR+14]   Shiri Chechik, Daniel H. Larkin, Liam Roditty, Grant Schoenebeck, Robert Endre Tarjan, and Virginia Vassilevska Williams. Better approximation algorithms for the graph diameter. In *SODA*, 2014.

[CN98]   Jin-Yi Cai and Ajay Nerurkar. Approximating the SVP to within a factor $(1 + 1/\dim^\varepsilon)$ is NP-hard under randomized conditions. In *CCC*, 1998.

[Dad12]   Daniel Dadush. A $O(1/\varepsilon^2)^n$-time sieving algorithm for approximate Integer Programming. In *LATIN*, 2012.

[Din16]   Irit Dinur. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:128, 2016.

[DKRS03]   Irit Dinur, Guy Kindler, Ran Raz, and Shmuel Safra. Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003.

[DPV11]   Daniel Dadush, Chris Peikert, and Santosh Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *FOCS*, 2011.

[DRS14]   Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. On the Closest Vector Problem with a distance guarantee. In *CCC*, pages 98–109, 2014.

[FM04]   Uriel Feige and Daniele Micciancio. The inapproximability of lattice and coding problems with preprocessing. *Journal of Computer and System Sciences*, 69(1):45–67, 2004. Preliminary version in CCC 2002.

[GG00]   Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000. Preliminary version in STOC 1998.

[GHNR03]   Jens Gramm, Edward A. Hirsch, Rolf Niedermeier, and Peter Rossmanith. Worst-case upper bounds for MAX-2-SAT with an application to MAX-CUT. *Discrete Applied Mathematics*, 130(2):139–155, 2003.

[GJS76]   Michael R. Garey, David S. Johnson, and Larry Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976.

[GMSS99]   Oded Goldreich, Daniele Micciancio, Shmuel Safra, and Jean-Pierre Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Information Processing Letters*, 71(2):55 – 61, 1999.

[GN08]   Nicolas Gama and Phong Q. Nguyen. Finding short lattice vectors within Mordell's inequality. In *STOC*, 2008.

[GPV08]   Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.

[HP]   S. Har-Peled. Concentration of Random Variables – Chernoff's Inequality. Available at http://sarielhp.org/teach/13/b_574_rand_alg/lec/07_chernoff.pdf.

[HR12]   Ishay Haviv and Oded Regev. Tensor-based hardness of the Shortest Vector Problem to within almost polynomial factors. *Theory of Computing*, 8(23):513–531, 2012. Preliminary version in STOC'07.

[IP99]   Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. In *CCC*, pages 237–240, 1999.

[IPZ01]   Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.

[Jam06]   G. J. O. Jameson. Counting zeros of generalised polynomials: Descartes' rule of signs and Laguerre's extensions. *The Mathematical Gazette*, 90(518):223–234, 2006.

[JS98]   Antoine Joux and Jacques Stern. Lattice reduction: A toolbox for the cryptanalyst. *Journal of Cryptology*, 11(3):161–185, 1998.

[Kan87]   Ravi Kannan. Minkowski's convex body theorem and Integer Programming. *Math. Oper. Res.*, 12(3):415–440, 1987.

[Kho05]   Subhash Khot. Hardness of approximating the Shortest Vector Problem in lattices. *Journal of the ACM*, 52(5):789–808, September 2005. Preliminary version in FOCS'04.

[KPV14]   Subhash Khot, Preyas Popat, and Nisheeth K. Vishnoi. $2^{\log^{1-\varepsilon} n}$ hardness for Closest Vector Problem with Preprocessing. *SIAM Journal on Computing*, 43(3):1184–1205, 2014.

[Laa15]   Thijs Laarhoven. Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In *CRYPTO*, 2015.

[Len83]   Hendrik W. Lenstra, Jr. Integer Programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983.

[LG14]   François Le Gall. Powers of tensors and fast matrix multiplication. In *ISAAC*, 2014.

[LLL82]   Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.

[LWXZ11]   Mingjie Liu, Xiaoyun Wang, Guangwu Xu, and Xuexin Zheng. Shortest lattice vectors in the presence of gaps. *IACR Cryptology ePrint Archive*, 2011:139, 2011.

[Mic01a]   Daniele Micciancio. The hardness of the Closest Vector Problem with Preprocessing. *IEEE Transactions on Information Theory*, 47(3):1212–1215, 2001.

[Mic01b]   Daniele Micciancio. The Shortest Vector Problem is NP-hard to approximate to within some constant. *SIAM Journal on Computing*, 30(6):2008–2035, March 2001. Preliminary version in FOCS 1998.

[MR17]   Pasin Manurangsi and Prasad Raghavendra. A birthday repetition theorem and complexity of approximating dense csps. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, pages 78:1–78:15, 2017.

[MV10]   Daniele Micciancio and Panagiotis Voulgaris. Faster exponential time algorithms for the Shortest Vector Problem. In *SODA*, pages 1468–1480, 2010.

[MV13]   Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. *SIAM Journal on Computing*, 42(3):1364–1391, 2013.

[MW15]   Daniele Micciancio and Michael Walter. Fast lattice point enumeration with minimal overhead. In *SODA*, 2015.

[MW16]   Daniele Micciancio and Michael Walter. Practical, predictable lattice basis reduction. In *Eurocrypt*, 2016.

[NIS16]   NIST post-quantum standardization call for proposals. http://csrc.nist.gov/groups/ST/post-quantum-crypto/cfp-announce-dec2016.html, 2016. Accessed: 2017-04-02.

[NS01]   Phong Q. Nguyen and Jacques Stern. The two faces of lattices in cryptology. In *Cryptography and Lattices*, pages 146–180. Springer, 2001.

[NV08]   Phong Q. Nguyen and Thomas Vidick. Sieve algorithms for the shortest vector problem are practical. *J. Math. Cryptol.*, 2(2):181–207, 2008.

[Odl90]   Andrew M. Odlyzko. The rise and fall of knapsack cryptosystems. *Cryptology and Computational Number Theory*, 42:75–88, 1990.

[Pei08]   Chris Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, May 2008. Preliminary version in CCC 2007.

[Pei16]   Chris Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4):283–424, 2016.

[PS09]   Xavier Pujol and Damien Stehlé. Solving the Shortest Lattice Vector Problem in time $2^{2.465n}$. *IACR Cryptology ePrint Archive*, 2009:605, 2009.

[PW10]   Mihai Pǎtraşcu and Ryan Williams. On the possibility of faster SAT algorithms. In *SODA*, 2010.

[Reg04]   Oded Regev. Improved inapproximability of lattice and coding problems with preprocessing. *IEEE Transactions on Information Theory*, 50(9):2031–2037, 2004. Preliminary version in CCC'03.

[Reg09]   Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):Art. 34, 40, 2009.

[RR06]   Oded Regev and Ricky Rosen. Lattice problems and norm embeddings. In *STOC*, 2006.

[Sch87]   Claus P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 53:201 – 224, 1987.

[SFS09]   Naftali Sommer, Meir Feder, and Ofir Shalvi. Finding the closest lattice point by iterative slicing. *SIAM J. Discrete Math.*, 23(2):715–731, 2009.

[Sha84]   Adi Shamir. A polynomial-time algorithm for breaking the basic Merkle-Hellman cryptosystem. *IEEE Trans. Inform. Theory*, 30(5):699–704, 1984.

[Vai15]   Vinod Vaikuntanathan. Private communication, 2015.

[vEB81]   Peter van Emde Boas. Another NP-complete problem and the complexity of computing
            short vectors in a lattice. Technical report, University of Amsterdam, Department of
            Mathematics, Netherlands, 1981. Technical Report 8104.

[Wil05]   Ryan Williams. A new algorithm for optimal 2-Constraint Satisfaction and its implica-
            tions. *Theoretical Computer Science*, 348(2-3):357–365, 2005.

[Wil12]   Virginia Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd.
            In *STOC*, 2012.

[Wil15]   Virginia Vassilevska Williams. Hardness of easy problems: Basing hardness on popular
            conjectures such as the Strong Exponential Time Hypothesis (invited talk). In *IPEC*,
            pages 17–29, 2015.

[Wil16]   Ryan Williams. Strong ETH breaks with Merlin and Arthur: Short non-interactive
            proofs of batch evaluation. In *CCC*, 2016.

[WLTB11]   Xiaoyun Wang, Mingjie Liu, Chengliang Tian, and Jingguo Bi. Improved Nguyen-Vidick
            heuristic sieve algorithm for shortest vector problem. In *ASIACCS*, 2011.

[Woe08]   Gerhard J. Woeginger. Open problems around exact algorithms. *Discrete Applied
            Mathematics*, 156(3):397–405, 2008.