

Sieving for Closest Lattice Vectors (with Preprocessing)

Thijs Laarhoven^(✉)

IBM Research, Rüschlikon, Switzerland
mail@thijs.com

Abstract. Lattice-based cryptography has recently emerged as a prime candidate for efficient and secure post-quantum cryptography. The two main hard problems underlying its security are the shortest vector problem (SVP) and the closest vector problem (CVP). Various algorithms have been studied for solving these problems, and for SVP, lattice sieving currently dominates in terms of the asymptotic time complexity: one can heuristically solve SVP in time $2^{0.292d+o(d)}$ in high dimensions d [Becker–Ducas–Gama–Laarhoven, SODA’16]. Although several SVP algorithms can also be used to solve CVP, it is not clear whether this also holds for heuristic lattice sieving methods. The best time complexity for CVP is currently $2^{0.377d+o(d)}$ [Becker–Gama–Joux, ANTS’14].

In this paper we revisit sieving algorithms for solving SVP, and study how these algorithms can be modified to solve CVP and its variants as well. Our first method is aimed at solving one problem instance and minimizes the overall time complexity for a single CVP instance with a time complexity of $2^{0.292d+o(d)}$. Our second method minimizes the amortized time complexity for several instances on the same lattice, at the cost of a larger preprocessing cost. Using nearest neighbor searching with a balanced space-time tradeoff, with this method we can solve the closest vector problem with preprocessing (CVPP) with $2^{0.636d+o(d)}$ space and preprocessing, in $2^{0.136d+o(d)}$ time, while the query complexity can be further reduced to $2^{0.059d+o(d)}$ at the cost of $2^{d+o(d)}$ space and preprocessing, or even to $2^{\varepsilon d+o(d)}$ for arbitrary $\varepsilon > 0$, at the cost of preprocessing time and memory complexities of $(1/\varepsilon)^{O(d)}$.

For easier variants of CVP, such as approximate CVP and bounded distance decoding (BDD), we further show how the preprocessing method achieves even better complexities. For instance, we can solve approximate CVPP with large approximation factors κ with polynomial-sized advice in polynomial time if $\kappa = \Omega(\sqrt{d/\log d})$. This heuristically closes the gap between the decision-CVPP result of [Aharonov–Regev, FOCS’04] (with equivalent κ) and the search-CVPP result of [Dadush–Regev–Stephens–Davidowitz, CCC’14] (which required larger κ).

Keywords: Lattices · Sieving algorithms · Approximate nearest neighbors · Shortest vector problem (SVP) · Closest vector problem (CVP) · Bounded distance decoding (BDD)

1 Introduction

Hard lattice problems. Lattices are discrete subgroups of \mathbb{R}^d . More concretely, given a basis $B = \{\mathbf{b}_1, \dots, \mathbf{b}_d\} \subset \mathbb{R}^d$, the lattice $\mathcal{L} = \mathcal{L}(B)$ generated by B is defined as $\mathcal{L}(B) = \left\{ \sum_{i=1}^d \lambda_i \mathbf{b}_i : \lambda_i \in \mathbb{Z} \right\}$. Given a basis of a lattice \mathcal{L} , the Shortest Vector Problem (SVP) asks to find a shortest non-zero vector in \mathcal{L} under the Euclidean norm, i.e., a non-zero lattice vector \mathbf{s} of norm $\|\mathbf{s}\| = \lambda_1(\mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L} \setminus \{0\}} \|\mathbf{v}\|$. Given a basis of a lattice and a target vector $\mathbf{t} \in \mathbb{R}^d$, the Closest Vector Problem (CVP) asks to find a vector $\mathbf{s} \in \mathcal{L}$ closest to \mathbf{t} under the Euclidean distance, i.e. such that $\|\mathbf{s} - \mathbf{t}\| = \min_{\mathbf{v} \in \mathcal{L}} \|\mathbf{v} - \mathbf{t}\|$.

These two hard problems are fundamental in the study of lattice-based cryptography, as the security of these schemes is directly related to the hardness of SVP and CVP in high dimensions. Various other hard lattice problems, such as Learning With Errors (LWE) and the Shortest Integer Solution (SIS) problem are closely related to SVP and CVP, and many reductions between these and other hard lattice problems are known; see e.g. [LvdPdW12, Fig. 3.1] or [Ste16] for an overview. These reductions show that being able to solve CVP efficiently implies that almost all other lattice problems can also be solved efficiently in the same dimension, which makes the study of the hardness of CVP even more important for choosing parameters in lattice-based cryptography.

Algorithms for SVP and CVP. Although SVP and CVP are both central in the study of lattice-based cryptography, algorithms for SVP have received somewhat more attention, including a benchmarking website to compare different algorithms [SG15]. Various SVP methods have been studied which can solve CVP as well, such as enumeration (see e.g. [Kan83, FP85, GNR10, MW15]), discrete Gaussian sampling [ADRS15, ADS15], constructing the Voronoi cell of the lattice [AEVZ02, MV10a], and using a tower of sublattices [BGJ14]. On the other hand, for the asymptotically fastest method in high dimensions for SVP¹, lattice sieving, it is not known how to solve CVP with similar costs as SVP.

After a series of theoretical works on constructing efficient heuristic sieving algorithms [NV08, MV10b, WLTB11, ZPH13, Laa15a, LdW15, BGJ15, BL16, BDGL16] as well as practical papers studying how to speed up these algorithms even further [MS11, Sch11, Sch13, BNvdP14, FBB+14, IKMT14, MTB14, MODB14, MLB15, MB16, MLB16], the best time complexity for solving SVP currently stands at $2^{0.292d+o(d)}$ [BDGL16, MLB16]. Although for various other methods the complexities for solving SVP and CVP are similar [GNR10, MV10a, ADS15], one can only guess whether the same holds for lattice sieving methods.

¹ To obtain provable guarantees, sieving algorithms are commonly modified to facilitate a somewhat artificial proof technique, which drastically increases the time complexity beyond e.g. the discrete Gaussian sampler and the Voronoi cell algorithm [AKS01, NV08, PS09, MV10b]. On the other hand, if some natural heuristic assumptions are made to enable analyzing the algorithm's behavior, then sieving clearly outperforms these methods. We focus on heuristic sieving in this paper.

To date, the best heuristic time complexity for solving CVP in high dimensions stands at $2^{0.377d+o(d)}$, due to Becker–Gama–Joux [BGJ14].

1.1 Contributions

In this paper we revisit heuristic lattice sieving algorithms, as well as the recent trend to speed up these algorithms using nearest neighbor searching, and we investigate how these algorithms can be modified to solve CVP and its generalizations. We present two different approaches for solving CVP with sieving, each of which we argue has its own merits.

Adaptive sieving. In *adaptive sieving*, we adapt the entire sieving algorithm to the problem instance, including the target vector. As the resulting algorithm is tailored specifically to the given CVP instance, this leads to the best asymptotic complexity for solving a single CVP instance out of our two proposed methods: $2^{0.292d+o(d)}$ time and space. This method is very similar to solving SVP with lattice sieving, and leads to equivalent asymptotics on the space and time complexities as for SVP. The corresponding space-time tradeoff is illustrated in Fig. 1, and equals that of [BDGL16] for solving SVP.

Non-adaptive sieving. Our main contribution, *non-adaptive sieving*, takes a different approach, focusing on cases where several CVP instances are to be solved on the same lattice. The goal here is to minimize the costs of computations depending on the target vector, and spend more time on preprocessing the lattice, so that the amortized time complexity per instance is smaller when solving many CVP instances on the same lattice. This is very closely related to the Closest Vector Problem with Preprocessing (CVPP), where the difference is that we allow for exponential-size preprocessed space. Using nearest neighbor techniques with a balanced space-time tradeoff, we show how to solve CVPP with $2^{0.636d+o(d)}$ space and preprocessing, in $2^{0.136d+o(d)}$ time. A continuous tradeoff between the two complexities can be obtained, where in the limit we can solve CVPP with $(1/\varepsilon)^{O(d)}$ space and preprocessing, in $2^{\varepsilon d+o(d)}$ time. This tradeoff is depicted in Fig. 1.

A potential application of non-adaptive sieving is as a subroutine within enumeration methods. As described in e.g. [GNR10], at any given level in the enumeration tree, one is attempting to solve a CVP instance in a lower-dimensional sublattice of \mathcal{L} , where the target vector is determined by the path chosen from the root to the current node in the tree. That means that if we can preprocess this sublattice such that the amortized time complexity of solving CVPP is small, then this could speed up processing the bottom part of the enumeration tree. This in turn might help speed up the lattice basis reduction algorithm BKZ [Sch87, SE94, CN11], which commonly uses enumeration as its SVP subroutine, and is key in assessing the security of lattice-based schemes. As the preprocessing needs to be performed once, CVPP algorithms with impractically large preprocessing costs may not be useful, but we show that with sieving the preprocessing costs can be quite small.




Fig. 1. Heuristic complexities for solving the Closest Vector Problem (CVP), the Closest Vector Problem with Preprocessing (CVPP), Bounded Distance Decoding with Preprocessing (δ -BDDP), and the Approximate Closest Vector Problem with Preprocessing (κ -CVPP). The red curve shows CVP complexities of Becker–Gama–Joux [BGJ14]. The left blue curve denotes CVP complexities of adaptive sieving. The right blue curve shows exact CVPP complexities using non-adaptive sieving. Purple curves denote relaxations of CVPP corresponding to different parameters δ (BDD radius) and κ (approximation factor). Note that exact CVPP corresponds to δ -BDDP with $\delta = 1$ and to κ -CVPP with $\kappa = 1$. (Color figure online)

Outline. The remainder of the paper is organized as follows. In Sect. 2 we describe some preliminaries, such as sieving algorithms and a useful result on nearest neighbor searching. Section 3 describes adaptive sieving and its analysis for solving CVP without preprocessing. Section 4 describes the preprocessing approach to solving CVP, with complexity analyses for exact CVP and some of its relaxations.

2 Preliminaries

2.1 Lattice Sieving for Solving SVP

Heuristic lattice sieving algorithms for solving the shortest vector problem all use the following basic property of lattices: if $\mathbf{v}, \mathbf{w} \in \mathcal{L}$, then their sum/difference $\mathbf{v} \pm \mathbf{w} \in \mathcal{L}$ is a lattice vector as well. Therefore, if we have a long list L of lattice vectors stored in memory, we can consider combinations of these vectors to obtain new, shorter lattice vectors. To make sure the algorithm makes progress in finding shorter lattice vectors, L needs to contain a lot of lattice vectors; for vectors $\mathbf{v}, \mathbf{w} \in \mathcal{L}$ of similar norm, the vector $\mathbf{v} - \mathbf{w}$ is shorter than \mathbf{v}, \mathbf{w} iff the angle between \mathbf{v}, \mathbf{w} is smaller than $\pi/3$, which for random vectors \mathbf{v}, \mathbf{w} occurs with probability $(3/4)^{d/2+o(d)}$. The expected space complexity of heuristic

sieving algorithms follows directly from this observation: if we draw $(4/3)^{d/2+o(d)}$ random vectors from the unit sphere, we expect a large number of pairs of vectors to have angle less than $\pi/3$, leading to many short difference vectors. This is exactly the heuristic assumption used in analyzing these sieving algorithms: when normalized, vectors in L follow the same distribution as vectors sampled uniformly at random from the unit sphere.

Heuristic 1. *When normalized, the list vectors $\mathbf{w} \in L$ behave as i.i.d. uniformly distributed random vectors from the unit sphere $\mathcal{S}^{d-1} := \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| = 1\}$.*

Therefore, if we start by sampling a list L of $(4/3)^{d/2+o(d)}$ long lattice vectors, and iteratively consider combinations of vectors in L to find shorter vectors, we expect to keep making progress. Note that naively, combining pairs of vectors in a list of size $(4/3)^{d/2+o(d)} \approx 2^{0.208d+o(d)}$ takes time $(4/3)^{d+o(d)} \approx 2^{0.415d+o(d)}$.

The Nguyen-Vidick sieve. The heuristic sieve algorithm of Nguyen and Vidick [NV08] starts by sampling a list L of $(4/3)^{d/2+o(d)}$ long lattice vectors, and uses a *sieve* to map L , with maximum norm $R := \max_{\mathbf{v} \in L} \|\mathbf{v}\|$, to a new list L' , with maximum norm at most γR for $\gamma < 1$ close to 1. By repeatedly applying this sieve, after $\text{poly}(d)$ iterations we expect to find a long list of lattice vectors of norm at most $\gamma^{\text{poly}(d)} R = O(\lambda_1(\mathcal{L}))$. The final list is then expected to contain a shortest vector of the lattice. Algorithm 3 in Appendix A describes a sieve equivalent to Nguyen-Vidick’s original sieve, to map L to L' in $|L|^2$ time.

Micciancio and Voulgaris’ GaussSieve. Micciancio and Voulgaris used a slightly different approach in the GaussSieve [MV10b]. This algorithm reduces the memory usage by immediately *reducing* all pairs of lattice vectors that are sampled. The algorithm uses a single list L , which is always kept in a state where for all $\mathbf{w}_1, \mathbf{w}_2 \in L$, $\|\mathbf{w}_1 \pm \mathbf{w}_2\| \geq \|\mathbf{w}_1\|, \|\mathbf{w}_2\|$, and each time a new vector $\mathbf{v} \in \mathcal{L}$ is sampled, its norm is reduced with vectors in L . After the norm can no longer be reduced, the vectors in L are reduced with \mathbf{v} . Modified list vectors are added to a stack to be processed later (to maintain the pairwise reduction-property of L), and new vectors which are pairwise reduced with L are added to L . Immediately reducing all pairs of vectors means that the algorithm uses less time and memory in practice, but at the same time Nguyen and Vidick’s heuristic proof technique does not apply here. However, it is commonly believed that the same bounds $(4/3)^{d/2+o(d)}$ and $(4/3)^{d+o(d)}$ on the space and time complexities hold for the GaussSieve. Pseudocode of the GaussSieve is given in Algorithm 4 in Appendix A.

2.2 Nearest Neighbor Searching

Given a data set $L \subset \mathbb{R}^d$, the nearest neighbor problem asks to preprocess L such that, when given a query $\mathbf{t} \in \mathbb{R}^d$, one can quickly return a nearest neighbor $\mathbf{s} \in L$ with distance $\|\mathbf{s} - \mathbf{t}\| = \min_{\mathbf{w} \in L} \|\mathbf{w} - \mathbf{t}\|$. This problem is essentially identical to CVP, except that L is a finite set of unstructured points, rather than the infinite set of all points in a lattice \mathcal{L} .

Locality-Sensitive Hashing/Filtering (LSH/LSF). A celebrated technique for finding nearest neighbors in high dimensions is Locality-Sensitive Hashing (LSH) [IM98, WSSJ14], where the idea is to construct many random partitions of the space, and store the list L in hash tables with buckets corresponding to regions. Preprocessing then consists of constructing these hash tables, while a query \mathbf{t} is answered by doing a lookup in each of the hash tables, and searching for a nearest neighbor in these buckets. More details on LSH in combination with sieving can be found in e.g. [Laa15a, LdW15, BGJ15, BL16].

Similar to LSH, Locality-Sensitive Filtering (LSF) [BDGL16, Laa15b] divides the space into regions, with the added relaxation that these regions do not have to form a partition; regions may overlap, and part of the space may not be covered by any region. This leads to improved results compared to LSH when L has size exponential in d [BDGL16, Laa15b]. Below we restate one of the main results of [Laa15b] for our applications. The specific problem considered here is: given a data set $L \subset \mathcal{S}^{d-1}$ sampled uniformly at random, and a random query $\mathbf{t} \in \mathcal{S}^{d-1}$, return a vector $\mathbf{w} \in L$ such that the angle between \mathbf{w} and \mathbf{t} is at most θ . The following result further assumes that the list L contains $n = (1/\sin \theta)^{d+o(d)}$ vectors.

Lemma 1. [Laa15b, Corollary 1] *Let $\theta \in (0, \frac{1}{2}\pi)$, and let $u \in [\cos \theta, 1/\cos \theta]$. Let $L \subset \mathcal{S}^{d-1}$ be a list of $n = (1/\sin \theta)^{d+o(d)}$ vectors sampled uniformly at random from \mathcal{S}^{d-1} . Then, using spherical LSF with parameters $\alpha_q = u \cos \theta$ and $\alpha_u = \cos \theta$, one can preprocess L in time $n^{1+\rho_u+o(1)}$, using $n^{1+\rho_u+o(1)}$ space, and with high probability answer a random query $\mathbf{t} \in \mathcal{S}^{d-1}$ correctly in time $n^{\rho_q+o(1)}$, where:*

$$n^{\rho_q} = \left(\frac{\sin^2 \theta (u \cos \theta + 1)}{u \cos \theta - \cos 2\theta} \right)^{d/2}, \quad n^{\rho_u} = \left(\frac{\sin^2 \theta}{1 - \cot^2 \theta (u^2 - 2u \cos \theta + 1)} \right)^{d/2}. \tag{1}$$

Applying this result to sieving for solving SVP, where $n = \sin(\frac{\pi}{3})^{-d+o(d)}$ and we are looking for pairs of vectors at angle at most $\frac{\pi}{3}$ to perform reductions, this leads to a space and preprocessing complexity of $n^{0.292d+o(d)}$, and a query complexity of $2^{0.084d+o(d)}$. As the preprocessing in sieving is only performed once, and queries are performed $n \approx 2^{0.208d+o(d)}$ times, this leads to a reduction of the complexities of sieving (for SVP) from $2^{0.208d+o(d)}$ space and $2^{0.415d+o(d)}$ time, to $2^{0.292d+o(d)}$ space and time [BDGL16].

3 Adaptive Sieving for CVP

We present two methods for solving CVP using sieving, the first of which we call *adaptive sieving* – we adapt the entire sieving algorithm to the particular CVP instance, to obtain the best overall time complexity for solving one instance. When solving several CVP instances, the costs roughly scale linearly with the number of instances.

Algorithm 1. The adaptive Nguyen-Vidick sieve for finding closest vectors

Require: Lists $L_0, L_t \subset \mathcal{L}$ containing $(4/3)^{d/2+o(d)}$ vectors at distance $\leq R$ from $\mathbf{0}, \mathbf{t}$

Ensure: Lists $L'_0, L'_t \subset \mathcal{L}$ contain $(4/3)^{d/2+o(d)}$ vectors at distance $\leq \gamma R$ from $\mathbf{0}, \mathbf{t}$

- 1: Initialize empty lists L'_0, L'_t
 - 2: **for each** $(\mathbf{w}_1, \mathbf{w}_2) \in L_0 \times L_0$ **do**
 - 3: **if** $\|\mathbf{w}_1 - \mathbf{w}_2\| \leq \gamma R$ **then**
 - 4: Add $\mathbf{w}_1 - \mathbf{w}_2$ to the list L'_0
 - 5: **for each** $(\mathbf{w}_1, \mathbf{w}_2) \in L_t \times L_0$ **do**
 - 6: **if** $\|(\mathbf{w}_1 - \mathbf{w}_2) - \mathbf{t}\| \leq \gamma R$ **then**
 - 7: Add $\mathbf{w}_1 - \mathbf{w}_2$ to the list L'_t
 - 8: **return** (L'_0, L'_t)
-

Using one list. The main idea behind this method is to translate the SVP algorithm by the target vector \mathbf{t} ; instead of generating a long list of lattice vectors reasonably close to $\mathbf{0}$, we generate a list of lattice vectors close to \mathbf{t} , and combine lattice vectors to find lattice vectors even closer vectors to \mathbf{t} . The final list then hopefully contains a closest vector to \mathbf{t} .

One quickly sees that this does not work, as the fundamental property of lattices does not hold for the lattice coset $\mathbf{t} + \mathcal{L}$: if $\mathbf{w}_1, \mathbf{w}_2 \in \mathbf{t} + \mathcal{L}$, then $\mathbf{w}_1 \pm \mathbf{w}_2 \notin \mathbf{t} + \mathcal{L}$. In other words, two lattice vectors close to \mathbf{t} can only be combined to form lattice vectors close to $\mathbf{0}$ or $2\mathbf{t}$. So if we start with a list of vectors close to \mathbf{t} , and combine vectors in this list as in the Nguyen-Vidick sieve, then after one iteration we will end up with a list L' of lattice vectors close to $\mathbf{0}$.

Using two lists. To make the idea of translating the whole problem by \mathbf{t} work for the Nguyen-Vidick sieve, we make the following modification: we keep track of two lists $L = L_0$ and L_t of lattice vectors close to $\mathbf{0}$ and \mathbf{t} , and construct a sieve which maps two input lists L_0, L_t to two output lists L'_0, L'_t of lattice vectors slightly closer to $\mathbf{0}$ and \mathbf{t} . Similar to the original Nguyen-Vidick sieve, we then apply this sieve several times to two initial lists (L_0, L_t) with a large radius R , to end up with two lists L_0 and L_t of lattice vectors at distance at most approximately $\sqrt{4/3} \cdot \lambda_1(\mathcal{L})$ from $\mathbf{0}$ and \mathbf{t} ². The argumentation that this algorithm works is almost identical to that for solving SVP, where we now make the following slightly different heuristic assumption.

Heuristic 2. *When normalized, the list vectors L_0 and L_t in the modified Nguyen-Vidick sieve both behave as i.i.d. uniformly distributed random vectors from the unit sphere.*

The resulting algorithm, based on the Nguyen-Vidick sieve, is presented in Algorithm 1.

² Observe that by the Gaussian heuristic, there are $(4/3)^{d/2+o(d)}$ vectors in \mathcal{L} within any ball of radius $\sqrt{4/3} \cdot \lambda_1(\mathcal{L})$. So the list size of the NV-sieve will surely decrease below $(4/3)^{d/2}$ when $R < \sqrt{4/3} \cdot \lambda_1(\mathcal{L})$.

Main result. As the (heuristic) correctness of this algorithm follows directly from the correctness of the original NV-sieve, and nearest neighbor techniques can be applied to this algorithm in similar fashion as well, we immediately obtain the following result. Note that space-time tradeoffs for SVP, such as the one illustrated in [BDGL16, Fig. 1], similarly carry over to solving CVP, and the best tradeoff for SVP (and therefore CVP) is depicted in Fig. 1.

Theorem 1. *Assuming Heuristic 2 holds, the adaptive Nguyen-Vidick sieve with spherical LSF solves CVP in time T and space S , with*

$$S = (4/3)^{d/2+o(d)} \approx 2^{0.208d+o(d)}, \quad T = (3/2)^{d/2+o(d)} \approx 2^{0.292d+o(d)}. \quad (2)$$

An important open question is whether these techniques can also be applied to the faster GaussSieve algorithm to solve CVP. The GaussSieve seems to make even more use of the property that the sum/difference of two lattice vectors is also in the lattice, and operations in the GaussSieve in \mathcal{L} cannot as easily be *mimicked* for the coset $\mathbf{t} + \mathcal{L}$. Solving CVP with the GaussSieve with similar complexities is left as an open problem.

4 Non-adaptive Sieving for CVPP

Our second method for finding closest vectors with heuristic lattice sieving follows a slightly different approach. Instead of focusing only on the total time complexity for one problem instance, we split the algorithm into two phases:

- Phase 1: Preprocess the lattice \mathcal{L} , without knowledge of the target \mathbf{t} ;
- Phase 2: Process the query \mathbf{t} and output a closest lattice vector $\mathbf{s} \in \mathcal{L}$ to \mathbf{t} .

Intuitively it may be more important to keep the costs of Phase 2 small, as the preprocessed data can potentially be reused later for other instances on the same lattice. This approach is essentially equivalent to the Closest Vector Problem with Preprocessing (CVPP): preprocess \mathcal{L} such that when given a target vector \mathbf{t} later, one can quickly return a closest vector $\mathbf{s} \in \mathcal{L}$ to \mathbf{t} . For CVPP however the preprocessed space is usually restricted to be of polynomial size, and the time used for preprocessing the lattice is often not taken into account. Here we will keep track of the preprocessing costs as well, and we do not restrict the output from the preprocessing phase to be of size $\text{poly}(d)$.

Algorithm description. To minimize the costs of answering a query, and to do the preprocessing independently of the target vector, we first run a standard SVP sieve, resulting in a large list L of almost all short lattice vectors. Then, after we are given the target vector \mathbf{t} , we use L to reduce the target. Finally, once the resulting vector $\mathbf{t}' \in \mathbf{t} + \mathcal{L}$ can no longer be reduced with our list, we hope that this reduced vector \mathbf{t}' is the shortest vector in the coset $\mathbf{t} + \mathcal{L}$, so that $\mathbf{0}$ is the closest lattice vector to \mathbf{t}' and $\mathbf{s} = \mathbf{t} - \mathbf{t}'$ is the closest lattice vector to \mathbf{t} .

Algorithm 2. Non-adaptive sieving (Phase 2) for finding closest vectors

Require: A list $L \subset \mathcal{L}$ of $\alpha^{d/2+o(d)}$ vectors of norm at most $\alpha \cdot \lambda_1(\mathcal{L})$, and $\mathbf{t} \in \mathbb{R}^d$

Ensure: The output vector \mathbf{s} is the closest lattice vector to \mathbf{t} (w.h.p.)

- 1: Initialize $\mathbf{t}' \leftarrow \mathbf{t}$
- 2: **for each** $\mathbf{w} \in L$ **do**
- 3: **if** $\|\mathbf{t}' - \mathbf{w}\| \leq \|\mathbf{t}'\|$ **then**
- 4: Replace $\mathbf{t}' \leftarrow \mathbf{t}' - \mathbf{w}$ and restart the **for**-loop
- 5: **return** $\mathbf{s} = \mathbf{t} - \mathbf{t}'$

The first phase of this algorithm consists in running a sieve and storing the resulting list in memory (potentially in a nearest neighbor data structure for faster lookups). For this phase either the Nguyen-Vidick sieve or the GaussSieve can be used. The second phase is the same for either method, and is described in Algorithm 2 for the general case of an input list essentially consisting of the $\alpha^{d+o(d)}$ shortest vectors in the lattice. Note that a standard SVP sieve would produce a list of size $(4/3)^{d/2+o(d)}$ corresponding to $\alpha = \sqrt{4/3}$.

List size. We first study how large L must be to guarantee that the algorithm succeeds. One might wonder why we do not fix $\alpha = \sqrt{4/3}$ immediately in Algorithm 2. To see why this choice of α does not suffice, suppose we have a vector $\mathbf{t}' \in \mathbf{t} + \mathcal{L}$ which is no longer reducible with L . This implies that \mathbf{t}' has norm approximately $\sqrt{4/3} \cdot \lambda_1(\mathcal{L})$, similar to what happens in the GaussSieve. Now, unfortunately the fact that \mathbf{t}' cannot be reduced with L anymore, does *not* imply that the closest lattice point to \mathbf{t}' is $\mathbf{0}$. In fact, it is more likely that there exists an $\mathbf{s} \in \mathbf{t} + \mathcal{L}$ of norm slightly more than $\sqrt{4/3} \cdot \lambda_1(\mathcal{L})$ which is closer to \mathbf{t}' , but which is not used for reductions.

By the Gaussian heuristic, we expect the distance from \mathbf{t} and \mathbf{t}' to the lattice to be $\lambda_1(\mathcal{L})$. So to guarantee that $\mathbf{0}$ is the closest lattice vector to the reduced vector \mathbf{t}' , we need \mathbf{t}' to have norm at most $\lambda_1(\mathcal{L})$. To analyze and prove correctness of this algorithm, we will therefore prove that, under the assumption that the input is a list of the $\alpha^{d+o(d)}$ shortest lattice vectors of norm at most $\alpha \cdot \lambda_1(\mathcal{L})$ for a particular choice of α , w.h.p. the algorithm reduces \mathbf{t} to a vector $\mathbf{t}' \in \mathbf{t} + \mathcal{L}$ of norm at most $\lambda_1(\mathcal{L})$.

To study how to set α , we start with the following elementary lemma regarding the probability of reduction between two uniformly random vectors with given norms.

Lemma 2. *Let $v, w > 0$ and let $\mathbf{v} = v \cdot \mathbf{e}_v$ and $\mathbf{w} = w \cdot \mathbf{e}_w$. Then:*

$$\mathbb{P}_{\mathbf{e}_v, \mathbf{e}_w \sim \mathcal{S}^{d-1}} \left(\|\mathbf{v} - \mathbf{w}\|^2 < \|\mathbf{v}\|^2 \right) \sim \left[1 - \left(\frac{w}{2v} \right)^2 \right]^{d/2+o(d)}. \tag{3}$$

Proof. Expanding $\|\mathbf{v} - \mathbf{w}\|^2 = v^2 + w^2 - 2vw \langle \mathbf{e}_v, \mathbf{e}_w \rangle$ and $\|\mathbf{v}\|^2 = v^2$, the condition $\|\mathbf{v} - \mathbf{w}\|^2 < \|\mathbf{v}\|^2$ equals $\frac{w}{2v} < \langle \mathbf{e}_v, \mathbf{e}_w \rangle$. The result follows from [BDGL16, Lemma 2.1].

Under Heuristic 1, we then obtain a relation between the choice of α for the input list and the expected norm of the reduced vector \mathbf{t}' as follows.

Lemma 3. *Let $L \subset \alpha \cdot \mathcal{S}^{d-1}$ be a list of $\alpha^{d+o(d)}$ uniformly random vectors of norm $\alpha > 1$, and let $\mathbf{v} \in \beta \cdot \mathcal{S}^{d-1}$ be sampled uniformly at random. Then, for high dimensions d , there exists a $\mathbf{w} \in L$ such that $\|\mathbf{v} - \mathbf{w}\| < \|\mathbf{v}\|$ if and only if*

$$\alpha^4 - 4\beta^2\alpha^2 + 4\beta^2 < 0. \tag{4}$$


Proof. By Lemma 2 we can reduce \mathbf{v} with $\mathbf{w} \in L$ with probability similar to $p = [1 - \frac{\alpha^2}{4\beta^2}]^{d/2+o(d)}$. Since we have $n = \alpha^{d+o(d)}$ such vectors \mathbf{w} , the probability that none of them can reduce \mathbf{v} is $(1-p)^n$, which is $o(1)$ if $n \gg 1/p$ and $1-o(1)$ if $n \ll 1/p$. Expanding $n \cdot p$, we obtain the given Eq. (4), where $\alpha^4 - 4\beta^2\alpha^2 + 4\beta^2 < 0$ implies $n \gg 1/p$.

Note that in our applications, we do not just have a list of $\alpha^{d+o(d)}$ lattice vectors of norm $\alpha \cdot \lambda_1(\mathcal{L})$; for any $\alpha_0 \in [1, \alpha]$ we expect L to contain $\alpha_0^{d+o(d)}$ lattice vectors of norm at most $\alpha_0 \cdot \lambda_1(\mathcal{L})$. To obtain a reduced vector \mathbf{t}' of norm $\beta \cdot \lambda_1(\mathcal{L})$, we therefore obtain the condition that for *some* value $\alpha_0 \in [1, \alpha]$, it must hold that $\alpha_0^4 - 4\beta^2\alpha_0^2 + 4\beta^2 < 0$.

From (4) it follows that $p(\alpha^2) = \alpha^4 - 4\beta^2\alpha^2 + 4\beta^2$ has two roots $r_1 < 2 < r_2$ for α^2 , which lie close to 2 for $\beta \approx 1$. The condition that $p(\alpha_0^2) < 0$ for some $\alpha_0 \leq \alpha$ is equivalent to $\alpha > r_1$, which for $\beta = 1 + o(1)$ implies that $\alpha^2 \geq 2 + o(1)$. This means that asymptotically we must set $\alpha = \sqrt{2}$, and use $n = 2^{d/2+o(d)}$ input vectors, to guarantee that w.h.p. the algorithm succeeds. A sketch of the situation is also given in Fig. 2a.

Modifying the first phase. As we will need a larger list of size $2^{d/2+o(d)}$ to make sure we can solve CVP exactly, we need to adjust Phase 1 of the algorithm as well. Recall that with standard sieving, we reduce vectors iff their angle is at most $\theta = \frac{\pi}{3}$, resulting in a list of size $(\sin \theta)^{-d+o(d)}$. As we now need the output list of the first phase to consist of $2^{d/2+o(d)} = (\sin \theta')^{-d+o(d)}$ vectors for $\theta' = \frac{\pi}{4}$, we make the following adjustment: only reduce \mathbf{v} and \mathbf{w} if their common angle is less than $\frac{\pi}{4}$. For unit length vectors, this condition is equivalent to reducing \mathbf{v} with \mathbf{w} iff $\|\mathbf{v} - \mathbf{w}\|^2 \leq (2 - \sqrt{2}) \cdot \|\mathbf{v}\|^2$. This further accelerates nearest neighbor techniques due to the smaller angle θ . Pseudocode for the modified first phase is given in Appendix B.

Main result. With the algorithm in place, let us now analyze its complexity for solving CVP. The first phase of the algorithm generates a list of size $2^{d/2+o(d)}$ by combining pairs of vectors, and naively this can be done in time $T_1 = 2^{d+o(d)}$ and space $S = 2^{d/2+o(d)}$, with query complexity $T_2 = 2^{d/2+o(d)}$. Using nearest neighbor searching (Lemma 1), the query and preprocessing complexities can be further reduced, leading to the following result.



(a) For solving **exact CVP**, we must reduce the vector t to a vector $t' \in t + \mathcal{L}$ of norm at most $\lambda_1(\mathcal{L})$. The nearest lattice point to t' lies in a ball of radius approximately $\lambda_1(\mathcal{L})$ around t' (blue), and almost all the mass of this ball is contained in the (black) ball around 0 of radius $\sqrt{2} \cdot \lambda_1(\mathcal{L})$. So if $s \in \mathcal{L} \setminus \{0\}$ had lain closer to t' than 0 , we would have reduced t' with s , since $s \in L$.

(b) For **variants of CVP**, a choice α for the list size implies a norm $\beta \cdot \lambda_1(\mathcal{L})$ of t' . The nearest lattice vector s to t' lies within $\delta \cdot \lambda_1(\mathcal{L})$ of t' ($\delta = 1$ for approx-CVP), so with high probability s has norm approximately $(\sqrt{\beta^2 + \delta^2}) \cdot \lambda_1(\mathcal{L})$. For δ -BDD, if $\sqrt{\beta^2 + \delta^2} \leq \alpha$ then we expect the nearest point s to be in the list L . For κ -CVP, if $\beta \leq \kappa$, then the lattice vector $t - t'$ has norm at most $\kappa \cdot \lambda_1(\mathcal{L})$.

Fig. 2. Comparison of the list size complexity analysis for CVP (left) and BDD/approximate CVP (right). The point t represents the target vector, and after a series of reductions with Algorithm 2, we obtain $t' \in t + \mathcal{L}$. Blue balls around t' depict regions in which we expect the closest lattice point to t' to lie, where the blue shaded area indicates a negligible fraction of this ball [BDGL16, Lemma 2]. (Color figure online)

Theorem 2. Let $u \in (\frac{1}{2}\sqrt{2}, \sqrt{2})$. Using non-adaptive sieving, we can solve CVP with preprocessing time T_1 , space complexity S , and query time complexity T_2 as follows:

$$S = T_1 = \left(\frac{1}{u(\sqrt{2} - u)} \right)^{d/2+o(d)}, \quad T_2 = \left(\frac{\sqrt{2} + u}{2u} \right)^{d/2+o(d)}. \quad (5)$$

Proof. These complexities follow from Lemma 1 with $\theta = \frac{\pi}{4}$, noting that the first phase can be performed in time and space $T_1 = S = n^{1+\rho_u}$, and the second phase in time $T_2 = n^{\rho_u}$.

To illustrate the time and space complexities of Theorem 2, we highlight three special cases u as follows. The full tradeoff curve for $u \in (\frac{1}{2}\sqrt{2}, \sqrt{2})$ is depicted in Fig. 1.

- Setting $u = \frac{1}{2}\sqrt{2}$, we obtain $S = T_1 = 2^{d/2+o(d)}$ and $T_2 \approx 2^{0.2925d+o(d)}$.
- Setting $u = 1$, we obtain $S = T_1 \approx 2^{0.6358d+o(d)}$ and $T_2 \approx 2^{0.1358d+o(d)}$.
- Setting $u = \frac{1}{2}(\sqrt{2} + 1)$, we get $S = T_1 = 2^{d+o(d)}$ and $T_2 \approx 2^{0.0594d+o(d)}$.

The first result shows that the query complexity of non-adaptive sieving is never worse than for adaptive sieving; only the space and preprocessing complexities are worse. The second and third results show that CVP can be solved in significantly less time, even with preprocessing and space complexities bounded by $2^{d+o(d)}$.

Minimizing the query complexity. As $u \rightarrow \sqrt{2}$, the query complexity keeps decreasing while the memory and preprocessing costs increase. For arbitrary $\varepsilon > 0$, we can set $u = u_\varepsilon \approx \sqrt{2}$ as a function of ε , resulting in asymptotic complexities $S = T_1 = (1/\varepsilon)^{O(d)}$ and $T_2 = 2^{\varepsilon d+o(d)}$. This shows that it is possible to obtain a slightly subexponential query complexity, at the cost of superexponential space, by taking $\varepsilon = o(1)$ as a function of d .

Corollary 1. *For arbitrary $\varepsilon > 0$, using non-adaptive sieving we can solve CVPP with preprocessing time and space complexities $(1/\varepsilon)^{O(d)}$, in time $2^{\varepsilon d+o(d)}$. In particular, we can solve CVPP in $2^{o(d)}$ time, using $2^{\omega(d)}$ space and preprocessing.*

Being able to solve CVPP in subexponential time with superexponential preprocessing and memory is neither trivial nor quite surprising. A naive approach to the problem, with this much memory, could for instance be to index the entire fundamental domain of \mathcal{L} in a hash table. One could partition this domain into small regions, solve CVP for the centers of each of these regions, and store all the solutions in memory. Then, given a query, one looks up which region \mathbf{t} is in, and returns the answer corresponding to that vector. With a sufficiently fine-grained partitioning of the fundamental domain, the answers given by the look-ups are accurate, and this algorithm probably also runs in subexponential time.

Although it may not be surprising that it is possible to solve CVPP in subexponential time with (super)exponential space, it is not clear what the complexities of other methods would be. Our method presents a clear tradeoff between the complexities, where the constants in the preprocessing exponent are quite small; for instance, we can solve CVPP in time $2^{0.06d+o(d)}$ with less than $2^{d+o(d)}$ memory, which is the same amount of memory/preprocessing of the best provable SVP and CVP algorithms [ADRS15, ADS15]. Indexing the fundamental domain may well require much more memory than this.

4.1 Bounded Distance Decoding with Preprocessing

We finally take a look at specific instances of CVP which are easier than the general problem, such as when the target \mathbf{t} lies unusually close to the lattice.

This problem naturally appears in practice, when a private key consists of a *good basis* of a lattice with short basis vectors, and the public key is a *bad basis* of the same lattice. An encryption of a message could then consist of the message being mapped to a lattice point $\mathbf{v} \in \mathcal{L}$, and a small error vector \mathbf{e} being added to \mathbf{v} ($\mathbf{t} = \mathbf{v} + \mathbf{e}$) to hide \mathbf{v} . If the noise \mathbf{e} is small enough, then with a good basis one can decode \mathbf{t} to the closest lattice vector \mathbf{v} , while someone with the bad basis cannot decode correctly. As decoding for arbitrary \mathbf{t} (solving CVP) is known to be hard even with knowledge of a good basis [Mic01, FM02, Reg04, AKKV05], \mathbf{e} needs to be very short, and \mathbf{t} must lie unusually close to the lattice.

So instead of assuming target vectors $\mathbf{t} \in \mathbb{R}^d$ are sampled at random, suppose that \mathbf{t} lies at distance at most $\delta \cdot \lambda_1(\mathcal{L})$ from \mathcal{L} , for $\delta \in (0, 1)$. For adaptive sieving, recall that the list size $(4/3)^{d/2+o(d)}$ is the minimum initial list size one can hope to use to obtain a list of short lattice vectors; with fewer vectors, one would not be able to solve SVP.³ For non-adaptive sieving however, it may be possible to reduce the list size below $2^{d/2+o(d)}$.

List size. Let us again assume that the preprocessed list L contains almost all $\alpha^{d+o(d)}$ lattice vectors of norm at most $\alpha \cdot \lambda_1(\mathcal{L})$. The choice of α implies a maximum norm $\beta_\alpha \cdot \lambda_1(\mathcal{L})$ of the reduced vector \mathbf{t}' , as described in Lemma 3. The nearest lattice vector $\mathbf{s} \in \mathcal{L}$ to \mathbf{t}' lies within radius $\delta \cdot \lambda_1(\mathcal{L})$ of \mathbf{t}' , and w.h.p. $\mathbf{s} - \mathbf{t}'$ is approximately orthogonal to \mathbf{t}' ; see Fig. 2b, where the shaded area is asymptotically negligible. Therefore w.h.p. \mathbf{s} has norm at most $(\sqrt{\beta_\alpha^2 + \delta^2}) \cdot \lambda_1(\mathcal{L})$. Now if $\sqrt{\beta_\alpha^2 + \delta^2} \leq \alpha$, then we expect the nearest vector to be contained in L , so that ultimately $\mathbf{0}$ is nearest to \mathbf{t}' . Substituting $\alpha^4 - 4\beta^2\alpha^2 + 4\beta^2 = 0$ and $\beta^2 + \delta^2 \leq \alpha^2$, and solving for α , this leads to the following condition on α .

$$\alpha^2 \geq \frac{2}{3}(1 + \delta^2) + \frac{2}{3}\sqrt{(1 + \delta^2)^2 - 3\delta^2}. \tag{6}$$

Taking $\delta = 1$, corresponding to exact CVP, leads to the condition $\alpha \geq \sqrt{2}$ as expected, while in the limiting case of $\delta \rightarrow 0$ we obtain the condition $\alpha \geq \sqrt{4/3}$. This matches experimental observations using the GaussSieve, where after finding the shortest vector, newly sampled vectors often cause *collisions* (i.e. being reduced to the $\mathbf{0}$ -vector). In other words, Algorithm 2 often reduces target vectors \mathbf{t} which essentially lie *on* the lattice ($\delta \rightarrow 0$) to the $\mathbf{0}$ -vector when the list has size $(4/3)^{d/2+o(d)}$. This explains why collisions in the GaussSieve are common when the list size grows to size $(4/3)^{d/2+o(d)}$.

Main result. To solve BDD with a target \mathbf{t} at distance $\delta \cdot \lambda_1(\mathcal{L})$ from the lattice, we need the preprocessing to produce a list of almost all $\alpha^{d+o(d)}$ vectors of norm at most $\alpha \cdot \lambda_1(\mathcal{L})$, with α satisfying (6). Similar to the analysis for CVP, we can produce such a list by only doing reductions between two vectors if their

³ The recent paper [BLS16] discusses how to use less memory in sieving, by using triple- or tuple-wise reductions, instead of the standard pairwise reductions. These techniques may also be applied to adaptive sieving to solve CVP with less memory, at the cost of an increase in the time complexity.

angle is less than θ , where now $\theta = \arcsin(1/\alpha)$. Combining this with Lemma 2, we obtain the following result.

Theorem 3. *Let α satisfy (6) and let $u \in (\sqrt{\frac{\alpha^2-1}{\alpha^2}}, \sqrt{\frac{\alpha^2}{\alpha^2-1}})$. Using non-adaptive sieving, we can heuristically solve BDD for targets \mathbf{t} at distance $\delta \cdot \lambda_1(\mathcal{L})$ from the lattice, with preprocessing time T_1 , space complexity S , and query time complexity T_2 as follows:*

$$S = \left(\frac{1}{1 - (\alpha^2 - 1)(u^2 - \frac{2u}{\alpha}\sqrt{\alpha^2 - 1} + 1)} \right)^{d/2+o(d)}, \tag{7}$$

$$T_1 = \max \left\{ S, (3/2)^{d/2+o(d)} \right\}, \quad T_2 = \left(\frac{\alpha + u\sqrt{\alpha^2 - 1}}{2\alpha - \alpha^3 + \alpha^2 u\sqrt{\alpha^2 - 1}} \right)^{d/2+o(d)}. \tag{8}$$

Proof. These complexities directly follow from applying Lemma 1 with $\theta = \arcsin(1/\alpha)$, and again observing that Phase 1 can be performed in time $T_1 = n^{1+\rho_u}$ and space $S = n^{1+\rho_u}$, while Phase 2 takes time $T_2 = n^{\rho_q}$. Note that we cannot combine vectors whose angles are larger than $\frac{\pi}{3}$ in Phase 1, which leads to a lower bound on the preprocessing time complexity T_1 based on the costs of solving SVP.

Theorem 3 is a generalization of Theorem 2, as the latter can be derived from the former by substituting $\delta = 1$ above. To illustrate the results, Fig. 1 considers two special cases:

- For $\delta = \frac{1}{2}$, we find $\alpha \approx 1.1976$, leading to $S \approx 2^{0.2602d+o(d)}$ and $T_2 = 2^{0.1908d+o(d)}$ when minimizing the space complexity.
- For $\delta \rightarrow 0$, we have $\alpha \rightarrow \sqrt{4/3} \approx 1.1547$. The minimum space complexity is therefore $S = (4/3)^{d/2+o(d)}$, with query complexity $T_2 = 2^{0.1610d+o(d)}$.

In the limit of $u \rightarrow \sqrt{\frac{\alpha^2}{\alpha^2-1}}$ we need superexponential space/preprocessing $S, T_1 \rightarrow 2^{\omega(d)}$ and a subexponential query time $T_2 \rightarrow 2^{o(d)}$ for all $\delta > 0$.

4.2 Approximate Closest Vector Problem with Preprocessing

Given a lattice \mathcal{L} and a target vector $\mathbf{t} \in \mathbb{R}^d$, approximate CVP with approximation factor κ asks to find a vector $\mathbf{s} \in \mathcal{L}$ such that $\|\mathbf{s} - \mathbf{t}\|$ is at most a factor κ larger than the real distance from \mathbf{t} to \mathcal{L} . For random instances \mathbf{t} , by the Gaussian heuristic this means that a lattice vector counts as a solution iff it lies at distance at most $\kappa \cdot \lambda_1(\mathcal{L})$ from \mathbf{t} .

List size. Instead of reducing \mathbf{t} to a vector \mathbf{t}' of norm at most $\lambda_1(\mathcal{L})$ as is needed for solving exact CVP, we now want to make sure that the reduced vector \mathbf{t}' has norm at most $\kappa \cdot \lambda_1(\mathcal{L})$. If this is the case, then the vector $\mathbf{t} - \mathbf{t}'$

is a lattice vector lying at distance at most $\kappa \cdot \lambda_1(\mathcal{L})$, which w.h.p. qualifies as a solution. This means that instead of substituting $\beta = 1$ in Lemma 3, we now substitute $\beta = \kappa$. This leads to the condition that $\alpha_0^4 - 4\kappa^2\alpha_0^2 + 4\beta^2 < 0$ for some $\alpha_0 \leq \alpha$. By a similar analysis α^2 must therefore be larger than the smallest root $r_1 = 2\kappa(\kappa - \sqrt{\kappa^2 - 1})$ of this quadratic polynomial in α^2 . This immediately leads to the following condition on α :

$$\alpha^2 \geq 2\kappa \left(\kappa - \sqrt{\kappa^2 - 1} \right). \tag{9}$$

A sanity check shows that $\kappa = 1$, corresponding to exact CVP, indeed results in $\alpha \geq \sqrt{2}$, while in the limit of $\kappa \rightarrow \infty$ a value $\alpha \approx 1$ suffices to obtain a vector \mathbf{t}' of norm at most $\kappa \cdot \lambda_1(\mathcal{L})$. In other words, to solve approximate CVP with very large (constant) approximation factors, a preprocessed list of size $(1 + \varepsilon)^{d+o(d)}$ suffices.

Main result. Similar to the analysis of CVPP, we now take $\theta = \arcsin(1/\alpha)$ as the angle with which to reduce vectors in Phase 1, so that the output of Phase 1 is a list of almost all $\alpha^{d+o(d)}$ shortest lattice vectors of norm at most $\alpha \cdot \lambda_1(\mathcal{L})$. Using a smaller angle θ for reductions again means that nearest neighbor searching can speed up the reductions in both Phase 1 and Phase 2 even further. The exact complexities follow from Lemma 1.

Theorem 4. *Using non-adaptive sieving with spherical LSF, we can heuristically solve κ -CVP with similar complexities as in Theorem 3, where now α must satisfy (9).*

Note that only the dependence of α on κ is different, compared to the dependence of α on δ for bounded distance decoding. The complexities for κ -CVP arguably decrease *faster* than for δ -BDD: for instance, for $\kappa \approx 1.0882$ we obtain the same complexities as for BDD with $\delta = \frac{1}{2}$, while $\kappa = \sqrt{4/3} \approx 1.1547$ leads to the same complexities as for BDD with $\delta \rightarrow 0$. Two further examples are illustrated in Fig. 1:

- For $\kappa = 2$, we have $\alpha \approx 1.1976$, which for $u \approx 0.5503$ leads to $S = T_1 = 2^{0.2602d+o(d)}$ and $T_2 = 2^{0.1908d+o(d)}$, and for $u = 1$ leads to $S = T_1 = 2^{0.3573d+o(d)}$ and $T_2 = 2^{0.0971d+o(d)}$.
- For $\kappa \rightarrow \infty$, we have $\alpha \rightarrow 1$, i.e. the required preprocessed list size approaches $2^{o(d)}$ as κ grows. For sufficiently large κ , we can solve κ -CVP with a preprocessed list of size $2^{\varepsilon d+o(d)}$ in at most $2^{\varepsilon d+o(d)}$ time. The preprocessing time is given by $2^{0.2925d+o(d)}$.

The latter result shows that for any superconstant approximation factor $\kappa = \omega(1)$, we can solve the corresponding approximate closest vector problem with preprocessing in subexponential time, with an exponential preprocessing time complexity $2^{0.292d+o(d)}$ for solving SVP and generating a list of short lattice vectors, and a subexponential space complexity required for Phase 2.

In other words, even without superexponential preprocessing/memory we can solve CVPP with large approximation factors in subexponential time.

To compare this result with previous work, note that the lower bound on α from (9) tends to $1 + 1/(8\kappa^2) + O(\kappa^{-4})$ as κ grows. The query space and time complexities are further both proportional to $\alpha^{\Theta(d)}$. To obtain a polynomial query complexity and polynomial storage after the preprocessing phase, we can solve for κ , leading to the following result.

Corollary 2. *With non-adaptive sieving we can heuristically solve approximate CVPP with approximation factor κ in polynomial time with polynomial-sized advice iff $\kappa = \Omega(\sqrt{d/\log d})$.*

Proof. The query time and space complexities are given by $\alpha^{\Theta(d)}$, where $\alpha = 1 + \Theta(1/\kappa^2)$. To obtain polynomial complexities in d , we must have $\alpha^{\Theta(d)} = d^{O(1)}$, or equivalently:

$$1 + \Theta\left(\frac{1}{\kappa^2}\right) = \alpha = d^{O(1/d)} = \exp O\left(\frac{\log d}{d}\right) = 1 + O\left(\frac{\log d}{d}\right). \quad (10)$$

Solving for κ leads to the given relation between κ and d .

Apart from the heuristic assumptions we made, this is equivalent to a result of Aharonov and Regev [AR04], who previously showed that the decision version of CVPP with approximation factor $\kappa = \Omega(\sqrt{d/\log d})$ can provably be solved in polynomial time. This further improves upon results of [LLS90, DRS14], who are able to solve search-CVPP with polynomial time and space complexities for $\kappa = O(d^{3/2})$ and $\kappa = \Omega(d/\sqrt{\log d})$ respectively. Assuming the heuristic assumptions are valid, Corollary 2 closes the gap between these previous results for decision-CVPP and search-CVPP with a rather simple algorithm: (1) preprocess the lattice by storing all $d^{O(1)}$ shortest vectors of the lattice in a list; and (2) apply Algorithm 2 to this list and the target vector to find an approximate closest vector. Note that nearest neighbor techniques only affect leading constants; even without nearest neighbor searching this would heuristically result in a polynomial time and space algorithm for κ -CVPP with $\kappa = \Omega(\sqrt{d/\log d})$.

Acknowledgments. The author is indebted to Léo Ducas, whose initial ideas and suggestions on this topic motivated work on this paper. The author further thanks Vadim Lyubashevsky and Oded Regev for their comments on the relevance of a subexponential time CVPP algorithm requiring (super)exponential space. The author is supported by the SNSF ERC Transfer Grant CRETP2-166734 FELICITY.

A Pseudocode of SVP Algorithms

Algorithms 3 and 4 present pseudo-code for the (sieve part of the) original Nguyen-Vidick sieve and the GaussSieve, respectively, as described in Sect. 2. For the Nguyen-Vidick sieve, the presented algorithm is a more intuitive but equivalent version of the original sieve; see [Laa15a, Appendix B] for details on this equivalence.

Algorithm 3. The quadratic Nguyen-Vidick sieve for finding shortest vectors

Require: An input list $L \subset \mathcal{L}$ of $(4/3)^{d/2+o(d)}$ vectors of norm at most R

Ensure: The output list $L' \subset \mathcal{L}$ has $(4/3)^{d/2+o(d)}$ vectors of norm at most $\gamma \cdot R$

- 1: Initialize an empty list L'
 - 2: **for each** $w_1, w_2 \in L$ **do**
 - 3: **if** $\|w_1 - w_2\| \leq \gamma R$ **then**
 - 4: Add $w_1 - w_2$ to the list L'
 - 5: **return** L'
-

Algorithm 4. The GaussSieve algorithm for finding shortest vectors

Require: A basis B of a lattice $\mathcal{L}(B)$

Ensure: The algorithm returns a shortest lattice vector

- 1: Initialize an empty list L and an empty stack S
 - 2: **repeat**
 - 3: Get a vector v from the stack (or sample a new one if $S = \emptyset$)
 - 4: **for each** $w \in L$ **do**
 - 5: **if** $\|v - w\| \leq \|v\|$ **then**
 - 6: Replace $v \leftarrow v - w$
 - 7: **if** $\|w - v\| \leq \|w\|$ **then**
 - 8: Replace $w \leftarrow w - v$
 - 9: Move w from the list L to the stack S (unless $w = \mathbf{0}$)
 - 10: **if** v has changed **then**
 - 11: Add v to the stack S (unless $v = \mathbf{0}$)
 - 12: **else**
 - 13: Add v to the list L (unless $v = \mathbf{0}$)
 - 14: **until** v is a shortest vector
 - 15: **return** v
-

B Pseudocode of Phase 1 for Non-adaptive Sieving

To generate a list of the $\alpha^{d+o(d)}$ shortest lattice vectors with the GaussSieve, rather than the $(4/3)^{d/2+o(d)}$ lattice vectors one would get with standard sieving, we relax the reductions: reducing if $\|v - w\| < \|v\|$ corresponds to an angle $\pi/3$ between v and w , leading to a list size $(1/\sin(\frac{\pi}{3}))^{d+o(d)} = (4/3)^{d/2+o(d)}$. To obtain a list of size $\alpha^{d+o(d)}$, we reduce vectors if their angle is less than $\theta = \arcsin(1/\alpha)$, which for vectors v, w of similar norm corresponds to the following condition:

$$\|v - w\| < \sqrt{2(1 - \cos \theta)} \cdot \|v\| = \sqrt{2 - \frac{2}{\alpha} \sqrt{\alpha^2 - 1}} \cdot \|v\|. \quad (11)$$

This leads to the modified GaussSieve described in Algorithm 5.

Algorithm 5. The non-adaptive GaussSieve (Phase 1) for finding closest vectors

Require: A basis B of a lattice $\mathcal{L}(B)$, a parameter $\alpha > 1$

Ensure: The output list L contains $\alpha^{d+\alpha(d)}$ vectors of norm at most $\alpha \cdot \lambda_1(\mathcal{L})$

```

1: Initialize an empty list  $L$  and an empty stack  $S$ 
2: Let  $\alpha_0 = \max\{\alpha, \sqrt{4/3}\}$ 
3: repeat
4:   Get a vector  $v$  from the stack (or sample a new one if  $S = \emptyset$ )
5:   for each  $w \in L$  do
6:     if  $\|v - w\|^2 \leq (2 - \frac{2}{\alpha_0} \sqrt{\alpha_0^2 - 1}) \cdot \|v\|^2$  then
7:       Replace  $v \leftarrow v - w$ 
8:     if  $\|w - v\|^2 \leq (2 - \frac{2}{\alpha_0} \sqrt{\alpha_0^2 - 1}) \cdot \|w\|^2$  then
9:       Replace  $w \leftarrow w - v$ 
10:    Move  $w$  from the list  $L$  to the stack  $S$  (unless  $w = \mathbf{0}$ )
11:  if  $v$  has changed then
12:    Add  $v$  to the stack  $S$  (unless  $v = \mathbf{0}$ )
13:  else
14:    Add  $v$  to the list  $L$  (unless  $v = \mathbf{0}$ )
15: until  $v$  is a shortest vector
16: return  $L$ 

```

References

- [ADRS15] Aggarwal, D., Dadush, D., Regev, O., Stephens-Davidowitz, N.: Solving the shortest vector problem in 2^n time via discrete Gaussian sampling. In: STOC, pp. 733–742 (2015)
- [ADS15] Aggarwal, D., Dadush, D., Stephens-Davidowitz, N.: Solving the closest vector problem in 2^n time - the discrete Gaussian strikes again! In: FOCS (2015)
- [AEVZ02] Agrell, E., Eriksson, T., Vardy, A., Zeger, K.: Closest point search in lattices. *IEEE Trans. Inf. Theory* **48**(8), 2201–2214 (2002)
- [AKKV05] Alekhovich, M., Khot, S., Kindler, G., Vishnoi, N.: Hardness of approximating the closest vector problem with pre-processing. In: FOCS, pp. 216–225 (2005)
- [AKS01] Ajtai, M., Kumar, R., Sivakumar, D.: A sieve algorithm for the shortest lattice vector problem. In: STOC, pp. 601–610 (2001)
- [AR04] Aharonov, D., Regev, O.: Lattice problems in $\text{NP} \cap \text{coNP}$. In: FOCS, pp. 362–371 (2004)
- [BDGL16] Becker, A., Ducas, L., Gama, N., Laarhoven, T.: New directions in nearest neighbor searching with applications to lattice sieving. In: SODA, pp. 10–24 (2016)
- [BGJ14] Becker, A., Gama, N., Joux, A.: A Sieve algorithm based on overlattices. In: ANTS, pp. 49–70 (2014)
- [BGJ15] Becker, A., Gama, N., Joux, A.: Speeding-up lattice sieving without increasing the memory, using sub-quadratic nearest neighbor search. *Cryptology ePrint Archive, Report 2015/522*, pp. 1–14 (2015)
- [BL16] Becker, A., Laarhoven, T.: Efficient (ideal) lattice sieving using cross-polytope LSH. In: Pointcheval, D., Nitaj, A., Rachidi, T. (eds.) *AFRICACRYPT 2016*. LNCS, vol. 9646, pp. 3–23. Springer, Cham (2016). doi:[10.1007/978-3-319-31517-1_1](https://doi.org/10.1007/978-3-319-31517-1_1)

- [BLS16] Bai, S., Laarhoven, T., Stehlé, D.: Tuple lattice sieving. In: ANTS (2016)
- [BNvdP14] Bos, J.W., Naehrig, M., van de Pol, J.: Sieving for shortest vectors in ideal lattices: a practical perspective. Cryptology ePrint Archive, Report 2014/880, pp. 1–23 (2014)
- [CN11] Chen, Y., Nguyễn, P.Q.: BKZ 2.0: better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 1–20. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25385-0_1](https://doi.org/10.1007/978-3-642-25385-0_1)
- [DRS14] Dadush, D., Regev, O., Stephens-Davidowitz, N.: On the closest vector problem with a distance guarantee. In: CCC, pp. 98–109 (2014)
- [FBB+14] Fitzpatrick, R., Bischof, C., Buchmann, J., Dagdelen, Ö., Göpfert, F., Mariano, A., Yang, B.-Y.: Tuning GaussSieve for speed. In: Aranha, D.F., Menezes, A. (eds.) LATINCRYPT 2014. LNCS, vol. 8895, pp. 288–305. Springer, Cham (2015). doi:[10.1007/978-3-319-16295-9_16](https://doi.org/10.1007/978-3-319-16295-9_16)
- [FM02] Feige, U., Micciancio, D.: The inapproximability of lattice and coding problems with preprocessing. In: CCC, pp. 32–40 (2002)
- [FP85] Fincke, U., Pohst, M.: Improved methods for calculating vectors of short length in a lattice. Math. Comput. **44**(170), 463–471 (1985)
- [GNR10] Gama, N., Nguyễn, P.Q., Regev, O.: Lattice enumeration using extreme pruning. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 257–278. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-13190-5_13](https://doi.org/10.1007/978-3-642-13190-5_13)
- [IKMT14] Ishiguro, T., Kiyomoto, S., Miyake, Y., Takagi, T.: Parallel Gauss Sieve algorithm: solving the SVP challenge over a 128-dimensional ideal lattice. In: PKC, pp. 411–428 (2014)
- [IM98] Indyk, P., Motwani, R.: Approximate nearest neighbors: towards removing the curse of dimensionality. In: STOC, pp. 604–613 (1998)
- [Kan83] Kannan, R.: Improved algorithms for integer programming and related lattice problems. In: STOC, pp. 193–206 (1983)
- [Laa15a] Laarhoven, T.: Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 3–22. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-47989-6_1](https://doi.org/10.1007/978-3-662-47989-6_1)
- [Laa15b] Laarhoven, T.: Tradeoffs for nearest neighbors on the sphere (2015)
- [LdW15] Laarhoven, T., Weger, B.: Faster sieving for shortest lattice vectors using spherical locality-sensitive hashing. In: Lauter, K., Rodríguez-Henríquez, F. (eds.) LATINCRYPT 2015. LNCS, vol. 9230, pp. 101–118. Springer, Cham (2015). doi:[10.1007/978-3-319-22174-8_6](https://doi.org/10.1007/978-3-319-22174-8_6)
- [LLS90] Lagarias, J.C., Lenstra, H.W., Schnorr, C.-P.: Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. Combinatorica **10**(4), 333–348 (1990)
- [LvdPdW12] Laarhoven, T., van de Pol, J., de Weger, B.: Solving hard lattice problems and the security of lattice-based cryptosystems. Cryptology ePrint Archive, Report 2012/533, pp. 1–43 (2012)
- [MB16] Mariano, A., Bischof, C.: Enhancing the scalability and memory usage of HashSieve on multi-core CPUs. In: PDP (2016)
- [Mic01] Micciancio, D.: The hardness of the closest vector problem with preprocessing. IEEE Trans. Inf. Theory **47**(3), 1212–1215 (2001)
- [MLB15] Mariano, A., Laarhoven, T., Bischof, C.: Parallel (probable) lock-free HashSieve: a practical sieving algorithm for the SVP. In: ICPP, pp. 590–599 (2015)
- [MLB16] Mariano, A., Laarhoven, T., Bischof, C.: A parallel variant of LDSieve for the SVP on lattices (2016)

- [MODB14] Mariano, A., Dagdelen, Ö., Bischof, C.: A comprehensive empirical comparison of parallel ListSieve and GaussSieve. In: Lopes, L., Žilinskas, J., Costan, A., Cascella, R.G., Kecskemeti, G., Jeannot, E., Cannataro, M., Ricci, L., Benkner, S., Petit, S., Scarano, V., Gracia, J., Hunold, S., Scott, S.L., Lankes, S., Lengauer, C., Carretero, J., Breitbart, J., Alexander, M. (eds.) Euro-Par 2014. LNCS, vol. 8805, pp. 48–59. Springer, Cham (2014). doi:[10.1007/978-3-319-14325-5_5](https://doi.org/10.1007/978-3-319-14325-5_5)
- [MS11] Milde, B., Schneider, M.: A parallel implementation of GaussSieve for the shortest vector problem in lattices. In: PACT, pp. 452–458 (2011)
- [MTB14] Mariano, A., Timnat, S., Bischof, C.: Lock-free GaussSieve for linear speedups in parallel high performance SVP calculation. In: SBAC-PAD, pp. 278–285 (2014)
- [MV10a] Micciancio, D., Voulgaris, P.: A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In: STOC, pp. 351–358 (2010)
- [MV10b] Micciancio, D., Voulgaris, P.: Faster exponential time algorithms for the shortest vector problem. In: SODA, pp. 1468–1480 (2010)
- [MW15] Micciancio, D., Walter, M.: Fast lattice point enumeration with minimal overhead. In: SODA, pp. 276–294 (2015)
- [NV08] Nguyen, P.Q., Vidick, T.: Sieve algorithms for the shortest vector problem are practical. *J. Math. Cryptology* **2**(2), 181–207 (2008)
- [PS09] Pujol, X., Stehlé, D.: Solving the shortest lattice vector problem in time $2^{2 \cdot 465n}$. *Cryptology ePrint Archive*, Report 2009/605, pp. 1–7 (2009)
- [Reg04] Regev, O.: Improved inapproximability of lattice and coding problems with preprocessing. *IEEE Trans. Inf. Theory* **50**(9), 2031–2037 (2004)
- [Sch87] Schnorr, C.-P.: A hierarchy of polynomial time lattice basis reduction algorithms. *Theoret. Comput. Sci.* **53**(2–3), 201–224 (1987)
- [Sch11] Schneider, M.: Analysis of Gauss-Sieve for solving the shortest vector problem in lattices. In: Katoh, N., Kumar, A. (eds.) WALCOM 2011. LNCS, vol. 6552, pp. 89–97. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-19094-0_11](https://doi.org/10.1007/978-3-642-19094-0_11)
- [Sch13] Schneider, M.: Sieving for short vectors in ideal lattices. In: AFRICACRYPT, pp. 375–391 (2013)
- [SE94] Schnorr, C.-P., Euchner, M.: Lattice basis reduction: improved practical algorithms and solving subset sum problems. *Math. Program.* **66**(2–3), 181–199 (1994)
- [SG15] Schneider, M., Gama, N.: SVP challenge (2015)
- [Ste16] Stephens-Davidowitz, N.: Dimension-preserving reductions between lattice problems (2016). <http://noahsd.com/latticeproblems.pdf>
- [WLTB11] Wang, X., Liu, M., Tian, C., Bi, J.: Improved Nguyen-Vidick heuristic sieve algorithm for shortest vector problem. In: ASIACCS, pp. 1–9 (2011)
- [WSSJ14] Wang, J., Shen, H.T., Song, J., Ji, J.: Hashing for similarity search: a survey. [arXiv:1408.2927](https://arxiv.org/abs/1408.2927) [cs.DS], pp. 1–29 (2014)
- [ZPH13] Zhang, F., Pan, Y., Hu, G.: A three-level sieve algorithm for the shortest vector problem. In: Lange, T., Lauter, K., Lisoněk, P. (eds.) SAC 2013. LNCS, vol. 8282, pp. 29–47. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-43414-7_2](https://doi.org/10.1007/978-3-662-43414-7_2)