

Short Paths on the Voronoi Graph and Closest Vector Problem with Preprocessing

Daniel Dadush*

Nicolas Bonifas†

Abstract

Improving on the Voronoi cell based techniques of [28, 24], we give a *Las Vegas* $\tilde{O}(2^n)$ expected time and space algorithm for CVPP (the preprocessing version of the Closest Vector Problem, CVP). This improves on the $\tilde{O}(4^n)$ deterministic runtime of the Micciancio Voulgaris algorithm [24] (henceforth MV) for CVPP¹ at the cost of a polynomial amount of randomness (which only affects runtime, not correctness).

As in MV, our algorithm proceeds by computing a short path on the Voronoi graph of the lattice, where lattice points are adjacent if their Voronoi cells share a common facet, from the origin to a closest lattice vector. Our main technical contribution is a randomized procedure that, given the Voronoi relevant vectors of a lattice – the lattice vectors inducing facets of the Voronoi cell – as preprocessing, and any “close enough” lattice point to the target, computes a path to a closest lattice vector of expected polynomial size.

This improves on the $\tilde{O}(2^n)$ path length given by the MV algorithm. Furthermore, as in MV, each edge of the path can be computed using a single iteration over the Voronoi relevant vectors.

As a byproduct of our work, we also give an optimal relationship between geometric and path distance on the Voronoi graph, which we believe to be of independent interest.

Keywords. Closest Vector Problem, Lattice Problems, Convex Geometry.

1 Introduction

An n dimensional lattice \mathcal{L} in \mathbb{R}^n is defined as all integer combinations of some basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ of \mathbb{R}^n . The most fundamental computational problems on lattices are the Shortest and Closest Vector Problems, which we denote by SVP and CVP respectively. Given a basis $B \in \mathbb{R}^{n \times n}$ of \mathcal{L} , the SVP is to compute $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ minimizing $\|\mathbf{y}\|_2$, and the CVP is, given an additional target $\mathbf{t} \in \mathbb{R}^n$, to compute a vector $\mathbf{y} \in \mathcal{L}$ minimizing $\|\mathbf{t} - \mathbf{y}\|_2$ ².

The study of the algorithms and complexity of lat-

tice problems has yielded many fundamental results in Computer Science and other fields over the last three decades. Lattice techniques were introduced to factor polynomials with rational coefficients [20] and to show the polynomial solvability of integer programs with a fixed number of integer variables [20, 21]. It has been used as a cryptanalytic tool for breaking the security of knapsack crypto schemes [19], and in coding theory for developing structured codes [8] and asymptotically optimal codes for power-constrained additive white Gaussian noise (AWGN) channels [10]. Most recently, the security of powerful cryptographic primitives such as fully homomorphic encryption [12, 13, 6] have been based on the worst case hardness of lattice problems.

The Closest Vector Problem with Preprocessing. In CVP applications, a common setup is the need to solve many CVP queries over the same lattice but with varying targets. This is the case in the context of coding over a Gaussian noise channel, a fundamental channel model in wireless communication theory. Lattice codes, where the codewords correspond to a subset of lattice points, asymptotically achieve the AWGN channel capacity (for fixed transmission power), and maximum likelihood decoding for a noisy codeword corresponds (almost) exactly to a CVP query on the coding lattice. In the context of lattice based public key encryption, in most cases the decryption routine can be interpreted as solving an approximate (decisional) CVP over a public lattice, where the encrypted bit is 0 if the point is close and 1 if it is far.

CVP algorithms in this setting (and in general), often naturally break into a preprocessing phase, where useful information about the lattice is computed (i.e. short lattice vectors, a short basis, important sublattices, etc.), and a query / search phase, where the computed advice is used to answer CVP queries quickly. Since the advice computed during preprocessing is used across all CVP queries, if the number of CVP queries is large the work done in the preprocessing phase can be effectively “amortized out”. This motivates the definition of the Closest Vector Problem with Preprocessing (CVPP), where we fix an n dimensional lattice \mathcal{L} and measure only the complexity of answering CVP queries

*Department of Computer Science, New York University, New York (USA). dadush@cs.nyu.edu

†LIX, École Polytechnique, Palaiseau and IBM, Gentilly (France). nicolas.bonifas@polytechnique.edu

¹The MV algorithm also solves CVP, as the preprocessing can be computed in the same time bound.

²The SVP and CVP can be defined over any norm, though we restrict our attention here to the Euclidean norm.

on \mathcal{L} after the preprocessing phase has been completed (crucially, the preprocessing is done before the CVP queries are known). To avoid trivial solutions to this problem, i.e. not allowing the preprocessing phase to compute a table containing all CVP solutions, we restrict the amount of space (as a function of the encoding size of the input lattice basis) needed to store the preprocessing advice.

Complexity. While the ability to preprocess the lattice is very powerful, it was shown in [25] that CVPP is NP-hard when the size of the preprocessing advice is polynomial. Subsequently, approximation hardness for the gap version of CVPP (i.e. approximately deciding the distance of the target) was shown in [11, 27, 4], culminating in a hardness factor of $2^{\log^{1-\varepsilon} n}$ for any $\varepsilon > 0$ [17] under the assumption that NP is not in randomized quasi-polynomial time. On the positive side, polynomial time algorithms for the approximate search version of CVPP were studied (implicitly) in [5, 18], where the current best approximation factor $O(n/\sqrt{\log n})$ was recently achieved in [7]. For the gap decisional version of CVPP, the results are better, where the current best approximation factor is $O(\sqrt{n/\log n})$ [1].

Exact CVPP algorithms. Given the hardness results for polynomial sized preprocessing, we do not expect efficient algorithms for solving exact CVPP for general lattices. For applications in wireless coding however, one has control over the coding lattice, though constructing coding lattices with good error correcting properties (i.e. large minimum distance) for which decoding is “easy” remains an outstanding open problem. In this context, the study of fast algorithms for exact CVPP in general lattices can yield new tools in the context of lattice design, as well as new insights for solving CVP without preprocessing.

The extant algorithms for exact CVPP are in fact also algorithms for CVP, that is, the time to compute the preprocessing is bounded by query / search time. There are currently two classes of CVP algorithms which fit the preprocessing / search model (this excludes only the randomized sieving approaches [2, 3]).

The first class is based on lattice basis reduction [20], which use a “short” lattice basis as preprocessing to solve lattice problems, that is polynomial sized preprocessing. The fastest such algorithm is due to Kannan [16], with subsequent refinements in [15, 5, 14, 26], which computes a Hermite-Korkine-Zolotareff basis (HKZ) during the preprocessing phase in $\tilde{O}(n^{\frac{2n}{2\varepsilon}})^3$ time and $\text{poly}(n)$ space, and in the query phase uses a search tree to compute the coefficients of the closest vector under the HKZ basis in $\tilde{O}(n^{\frac{n}{2}})$ time

and $\text{poly}(n)$ space.

The second class, which are the most relevant to this work, use the Voronoi cell (see Section A.1.1 for precise definitions) of the lattice – the centrally symmetric polytope corresponding to the points closer to the origin than to other lattice points – as preprocessing, and were first introduced by Sommer, Feder and Shalvi [28]. In [28], they give an iterative procedure that uses the facet inducing lattice vectors of the Voronoi cell (known as the *Voronoi relevant vectors*) to move closer and closer to the target, and show that this procedure converges to a closest lattice vector in a finite number of steps. The number of Voronoi relevant vectors is $2(2^n - 1)$ in the worst-case (this holds for almost all lattices), and hence Voronoi cell based algorithms often require exponential size preprocessing. Subsequently, Micciancio and Voulgaris [24] (henceforth MV), showed how to compute the Voronoi relevant vectors during preprocessing and how to implement the search phase such that each phase uses $\tilde{O}(4^n)$ time and $\tilde{O}(2^n)$ space (yielding the first $2^{O(n)}$ time algorithm for exact CVP!).

While Voronoi cell based CVPP algorithms require exponential time and space on general lattices, it was recently shown in [23] that a variant of [28] can be implemented in polynomial time for lattices of Voronoi’s first kind – lattices which admit a set of $n+1$ generators whose Gram matrix is the Laplacian of a non-negatively weighted graph – using these generators as the preprocessing advice. Hence, it is sometimes possible to “scale down” the complexity of exact solvers for interesting classes of lattices.

Main Result. Our main result is a randomized $\tilde{O}(2^n)$ expected time and space algorithm for exact CVPP, improving the $\tilde{O}(4^n)$ (deterministic) running time of MV. Our preprocessing is the same as MV, that is we use the facet inducing lattice vectors of the Voronoi cell, known as the *Voronoi relevant vectors* (see Figure 1), as the preprocessing advice, which in the worst case consists of $2(2^n - 1)$ lattice vectors. Our main contribution, is a new search algorithm that requires only an expected polynomial number of iterations over the set of Voronoi relevant vectors to converge to a closest lattice vector, compared to $\tilde{O}(2^n)$ in MV.

One minor caveat to our iteration bound is that unlike that of MV, which only depends on n , ours also depends (at worst linearly) on the binary encoding length of the input lattice basis and target (though the $\tilde{O}(2^n)$ bound also holds for our procedure). Hence, while the bound is polynomial, it is only “weakly” so. In applications however, it is rather anomalous to encounter n dimensional lattice bases and targets whose individual coefficients require more than say $\text{poly}(n)$ bits to represent, and hence the iteration bound will be

³The \tilde{O} notation suppresses polylogarithmic factors.




Figure 1: Voronoi relevant vectors

$\text{poly}(n)$ in almost all settings of relevance. Furthermore, it is unclear if this dependence of our algorithm is inherent, or whether it is just an artifact of the analysis.

While our algorithm is randomized, it is Las Vegas, and hence the randomness is in the runtime and not the correctness. Furthermore, the amount of randomness we require is polynomial: it corresponds to the randomness needed to generate a nearly-uniform sample from the Voronoi cell, which can be achieved using Monte Carlo Markov Chain (MCMC) methods over convex bodies [9, 22]. This requires a polynomial number of calls to a membership oracle. Each membership oracle test requires an enumeration over the $\tilde{O}(2^n)$ Voronoi-relevant vectors, resulting in a total complexity of $\tilde{O}(2^n)$.

Unfortunately, we do not know how to convert our CVPP improvement to one for CVP. The technical difficulty lies in the fact that computing the Voronoi relevant vectors, using the current approach, is reduced to solving $\tilde{O}(2^n)$ related lower dimensional CVPs on an $n - 1$ dimensional lattice (for which the Voronoi cell has already been computed). While the MV CVPP algorithm requires $\tilde{O}(4^n)$ for worst case targets (which we improve to $\tilde{O}(2^n)$), they are able to use the relations between the preprocessing CVPs to solve each of them in amortized $\tilde{O}(2^n)$ time per instance. Hence, with the current approach, reducing the running time of CVP to $\tilde{O}(2^n)$ would require reducing the amortized per instance complexity to polynomial, which seems very challenging.

Organization. In section 2, we explain how to solve CVPP by finding short paths over the Voronoi graph. In particular, we review the iterative slicer [28] and MV [24] algorithms for navigating the Voronoi graph, and describe our new *randomized straight line procedure* for this task. In section 3, we state the guarantees for the randomized straight line procedure and use it to give our expected $\tilde{O}(2^n)$ time CVPP algo-

rithm (Theorem 3.4), as well as an optimal relationship between geometric and path distance on the Voronoi graph (Theorem 3.3). The main geometric estimates underlying the analysis of the randomized straight path algorithm are proved in appendix B, and the omitted proofs from section 3 are provided in appendix C.

For lack of space, most definitions are deferred to the preliminaries in appendix A. In particular, see subsections A.1 for basic lattice definitions, and subsection A.1.1 for precise definitions and fundamental facts about the Voronoi cell and related concepts.

2 Navigating the Voronoi graph

In this section, we explain how one can solve CVP using an efficient navigation algorithm over the Voronoi graph of a lattice. We first describe the techniques used by [28, 24] for finding short paths on this graph, and then give our new (randomized) approach.

Paths on the Voronoi graph. Following the strategy of [28, 24], our search algorithm works on the Voronoi graph \mathcal{G} of an n dimensional lattice \mathcal{L} . Let \mathcal{V} denote the Voronoi cell and let VR denote the set of Voronoi relevant vectors of \mathcal{L} (see Section A for precise definitions). Defined geometrically, the Voronoi graph \mathcal{G} is the contact graph induced by the tiling of space by Voronoi cells, that is, two lattice vectors $\mathbf{x}, \mathbf{y} \in \mathcal{L}$ are adjacent if their associated Voronoi cells $\mathbf{x} + \mathcal{V}$ and $\mathbf{y} + \mathcal{V}$ touch in a shared facet (equivalently $\mathbf{x} - \mathbf{y} \in \text{VR}$). We denote the shortest path distance between $\mathbf{x}, \mathbf{y} \in \mathcal{L}$ on \mathcal{G} by $d_{\mathcal{G}}(\mathbf{x}, \mathbf{y})$.

To solve CVP on a target \mathbf{t} , the idea of Voronoi cell based methods is to compute a short path on the Voronoi graph \mathcal{G} from a “close enough” starting vertex $\mathbf{x} \in \mathcal{L}$ to \mathbf{t} (usually, a rounded version of \mathbf{t} under some basis), to the center $\mathbf{y} \in \mathcal{L}$ of a Voronoi cell containing \mathbf{t} , which we note is a closest lattice vector by definition. (see Figure 2).

Iterative slicer. The iterative slicer [28] was the first CVP algorithm to make use of an explicit description of the Voronoi cell, in the form of the VR vectors.

The path steps of the iterative slicer are computed by greedily choosing any Voronoi relevant vector that brings the current iterate $\mathbf{z} \in \mathcal{L}$ closer to the target \mathbf{t} . That is, if there exists a VR vector \mathbf{v} such that $\|\mathbf{z} + \mathbf{v} - \mathbf{t}\|_2 < \|\mathbf{z} - \mathbf{t}\|_2$, then we move to $\mathbf{z} + \mathbf{v}$. This procedure is iterated until there is no improving VR vector, at which point we have reached a closest lattice vector to \mathbf{t} . This procedure was shown to terminate in a finite number of steps, and currently, no good quantitative bound is known on its convergence time.

The Voronoi norm. We now make precise which notion of closeness to the target we use (as well as MV) for the starting lattice vector \mathbf{x} to the target \mathbf{t} . Notice




Figure 2: CVP solution is the center of target-containing Voronoi cell

that for the path finding approach to make sense from the perspective of CVP, we need to start the process from a point $\mathbf{x} \in \mathcal{L}$ that we know is a priori close in graph distance to a closest lattice vector \mathbf{y} to \mathbf{t} . Given the complexity of \mathcal{G} and the fact that we do not know \mathbf{y} , we will need a robust proxy for graph distance that we can estimate knowing only \mathbf{x} and \mathbf{t} . From this perspective, it was shown in [24] that the Voronoi norm

$$\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}} = \inf \{s \geq 0 : \mathbf{t} - \mathbf{x} \in s\mathcal{V}\} = \sup_{\mathbf{v} \in \text{VR}} 2 \frac{\langle \mathbf{v}, \mathbf{t} - \mathbf{x} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle}$$

of $\mathbf{t} - \mathbf{x}$ (i.e. the smallest scaling of \mathcal{V} containing $\mathbf{t} - \mathbf{x}$) can be used to bound the shortest path distance between \mathbf{x} and \mathbf{y} . Here the quantity $\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}}$ is robust in the sense that $\|\mathbf{y} - \mathbf{x}\|_{\mathcal{V}} \leq \|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}} + \|\mathbf{y} - \mathbf{t}\|_{\mathcal{V}} \leq \|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}} + 1$ by the triangle inequality. Hence from the perspective of the Voronoi norm, \mathbf{t} is simply a “noisy” version of \mathbf{y} . Furthermore, given that each Voronoi relevant vector has Voronoi norm 2, one can construct a lattice vector \mathbf{x} such that $\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}} \leq n$, by simply expressing $\mathbf{t} = \sum_{i=1}^n a_i \mathbf{v}_i$, for some linearly independent $\mathbf{v}_1, \dots, \mathbf{v}_n \in \text{VR}$, and letting $\mathbf{x} = \sum_{i=1}^n \lceil a_i \rceil \mathbf{v}_i$.

The MV Path. We now present the MV path finding approach, and give the relationship they obtain between $\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}}$ and the path distance to a closest lattice vector \mathbf{y} to \mathbf{t} .

The base principle of MV [24] is similar to that of the iterative slicer, but it uses a different strategy to select the next VR vector to follow, resulting in a provably single exponential path length.

In MV, a path step consists of tracing the straight line from the current path vertex $\mathbf{z} \in \mathcal{L}$ to the target \mathbf{t} , and moving to $\mathbf{z} + \mathbf{v}$ where $\mathbf{v} \in \text{VR}$ induces a facet (generically unique) of $\mathbf{z} + \mathcal{V}$ crossed by the line

segment $[\mathbf{z}, \mathbf{t}]$. It is not hard to check that each step can be computed using $O(n|\text{VR}|) = \tilde{O}(2^n)$ arithmetic operations, and hence the complexity of computing the path is $O(n|\text{VR}| \times \text{path length})$.

The main bound they give on the path length, is that if the start vertex $\mathbf{x} \in 2\mathcal{V} + \mathbf{t}$ (i.e. Voronoi distance less than 2), then the path length is bounded by 2^n . To prove the bound, they show that the path always stays inside $\mathbf{t} + 2\mathcal{V}$, that the ℓ_2 distance to the target monotonically decreases along the path (and hence it is acyclic), and that the number of lattice vectors in the interior of $\mathbf{t} + 2\mathcal{V}$ is at most 2^n .

To build the full path, they run this procedure on the Voronoi graph for decreasing exponential scalings of \mathcal{L}^4 , and build a path (on a supergraph of \mathcal{G}) of length $O(2^n \log_2 \|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}})$. One can also straightforwardly adapt the MV procedure to stay on \mathcal{G} , by essentially breaking up the line segment $[\mathbf{x}, \mathbf{t}]$ in pieces of length at most 2, yielding a path length of $O(2^n \|\mathbf{x} - \mathbf{t}\|_{\mathcal{V}})$. Since we can always achieve a starting distance of $\|\mathbf{x} - \mathbf{t}\|_{\mathcal{V}} \leq n$ by straightforward basis rounding, note that the distance term is lower order compared to the proportionality factor 2^n .

Randomized Straight Line. Given the 2^n proportionality factor between geometric and path distance achieved by the MV algorithm, the main focus of our work will be to reduce the proportionality factor to polynomial. In fact, will show the existence of paths of length $(n/2)(\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}} + 1)$, however the paths we are able to construct will be longer.

For our path finding procedure, the base idea is rather straightforward, we simply attempt to follow the sequence of Voronoi cells on the straight line from the start point \mathbf{x} to the target \mathbf{t} . We dub this procedure the straight line algorithm. As we will show, the complexity of computing this path follows the same pattern as MV (under certain genericity assumptions), and hence the challenge is proving that the number of Voronoi cells the path crosses is polynomial. Unfortunately, we do not know how to analyze this procedure directly. In particular, we are unable to rule out the possibility that a “short” line segment (say of Voronoi length $O(1)$) may pass through exponentially many Voronoi cells in the worst case (though we do not have any examples of this).

To get around the problem of having “unexpectedly many” crossings, we will make use of randomization to perturb the starting point of the line segment. Specifically, we will use a *randomized straight line* path from $\mathbf{x} \in \mathcal{L}$ to \mathbf{t} which proceeds as follows (see Figure 3):

- (A) Move to $\mathbf{x} + Z$, where $Z \sim \text{Uniform}(\mathcal{V})$ is sampled

⁴The MV path is in fact built on a supergraph of the Voronoi graph, which has edges corresponding to 2^iVR , $i \geq 0$.

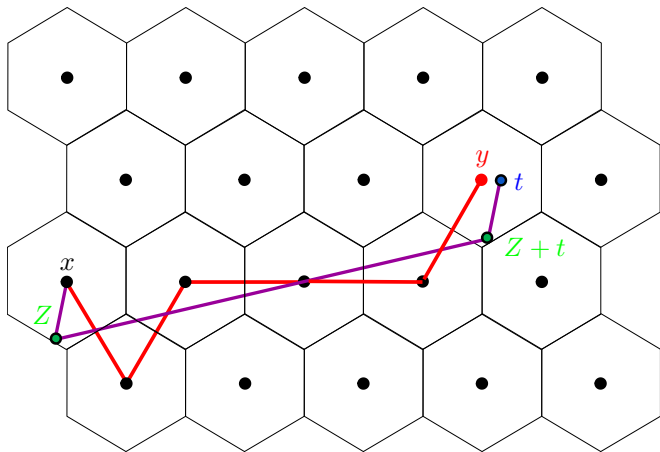


Figure 3: Randomized Straight Line algorithm

uniformly from the Voronoi cell.

(B) Follow the line from $\mathbf{x} + Z$ to $\mathbf{t} + Z$.

(C) Follow the line from $\mathbf{t} + Z$ to \mathbf{t} .

We briefly outline the analysis bounding the expected number of Voronoi cells this path crosses, which we claim achieves a polynomial proportionality factor with respect to $\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}}$.

To begin, note that in phase A, we stay entirely within $\mathbf{x} + Z$, and hence do not cross any Voronoi cells.

In phase B, at every time $\alpha \in [0, 1]$, the point $(1 - \alpha)\mathbf{x} + \alpha\mathbf{t} + Z$ is in a uniformly random coset of \mathbb{R}^n/\mathcal{L} since Z is uniform. Hence the probability that we cross a boundary between time α and $\alpha + \varepsilon$ is identical to the probability that we cross a boundary going from Z to $Z + \varepsilon(\mathbf{t} - \mathbf{x})$. Taking the limit as $\varepsilon \rightarrow 0$ and using linearity of expectation, we use the above invariance to show that the expected number of boundaries we cross is bounded by $(n/2)\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}}$, the Voronoi distance between \mathbf{x} and \mathbf{t} . In essence, we relate the number of crossings to the probability that a uniform sample from \mathcal{V} (equivalently, a uniform coset) is close under the Voronoi norm to the boundary $\partial\mathcal{V}$, which is a certain surface area to volume ratio.

Interestingly, as a consequence of our bound for phase B, we are able to give an optimal relationship between the Voronoi distance between two lattice points and their shortest path distance on \mathcal{G} , which we believe to be independent interest. In particular, for two lattice points $\mathbf{x}, \mathbf{y} \in \mathcal{L}$, we show in Theorem 3.3 that the shortest path distance on \mathcal{G} is at least $\|\mathbf{x} - \mathbf{y}\|_{\mathcal{V}}/2$ and at most $(n/2)\|\mathbf{x} - \mathbf{y}\|_{\mathcal{V}}$, which is tight for certain pairs of lattice points on \mathbb{Z}^n .

It remains now to bound the expected number of crossings in phase C. Here, the analysis is more difficult

than the second step, because the random shift is only on one side of the line segment from $\mathbf{t} + Z$ to \mathbf{t} . We will still be able to relate the expected number of crossings to “generalized” surface area to volume ratios, however the probability distributions at each time step will no longer be invariant modulo the lattice. In particular, the distributions become more concentrated as we move closer to \mathbf{t} , and hence we slowly lose the benefits of the randomness as we get closer to \mathbf{t} . Unfortunately, because of this phenomenon, we are unable to show in general that the number of crossings from $\mathbf{t} + Z$ to \mathbf{t} is polynomial. However, we will be able to bound the number of crossings from $\mathbf{t} + Z$ to $\mathbf{t} + \alpha Z$ by $O(n \ln(1/\alpha))$, that is, a very slow growing function of α as $\alpha \rightarrow 0$. Fortunately, for rational lattices and targets, we can show that for α not too small, in particular $\ln(1/\alpha)$ linear in the size of binary encoding of the basis and target suffices, that $\mathbf{t} + \alpha Z$ and \mathbf{t} lie in the same Voronoi cell. This yields the claimed (weakly) polynomial bound.

3 Analysis and Applications of Randomized Straight Line

In this section, we give the formal guarantees for the randomized straight line algorithm and its applications. The analysis here will rely on geometric estimates for the number of crossings, whose proofs are found in Section B.

To begin, we make formal the connection between Voronoi cells crossings, the length of the randomized straight line path, and the complexity of computing it.

LEMMA 3.1. (RANDOMIZED STRAIGHT LINE) *Let $\mathbf{x} \in \mathcal{L}$ be the starting point and let $\mathbf{t} \in \mathbb{R}^n$ be the target. Then using perturbation $Z \sim \text{Uniform}(\mathcal{V})$, the expected edge length of the path from \mathbf{x} to a closest lattice vector \mathbf{y} to \mathbf{t} on \mathcal{G} induced by the randomized straight line procedure is*

$$\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{t} + Z]|] + \mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{t} + Z, \mathbf{t}]|].$$

Furthermore, with probability 1, each edge of the path can be computed using $O(n|\text{VR}|)$ arithmetic operations.

While rather intuitive, the proof of this Lemma is somewhat tedious, and so we defer it to the appendix. Note that $(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{t} + Z]$ corresponds to the phase B crossings, and that $(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{t} + Z, \mathbf{t}]$ corresponds to the phase C crossings.

Our bound for the phase B crossings, which is proved in Section B.1, is as follows.

THEOREM 3.1. (PHASE B CROSSING BOUND) *Let \mathcal{L} be an n dimensional lattice. Then for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and*

$Z \sim \text{uniform}(\mathcal{V})$, we have that

$$\mathbb{E}_Z[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{y} + Z]|] \leq (n/2)\|\mathbf{y} - \mathbf{x}\|_{\mathcal{V}} .$$

For phase C, we give a bound on the number crossings for a truncation of the phase C path. That is, instead of going all the way from $\mathbf{t} + Z$ to \mathbf{t} , we stop at $\mathbf{t} + \alpha Z$, for $\alpha \in (0, 1]$. Its proof is given in Section B.2.

THEOREM 3.2. (PHASE C CROSSING BOUND) For $\alpha \in (0, 1]$, $Z \sim \text{Uniform}(\mathcal{V})$, $n \geq 2$, we have that

$$\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [Z + \mathbf{t}, \alpha Z + \mathbf{t}]|] \leq \frac{e^2}{\sqrt{2} - 1} n(2 + \ln(4/\alpha)) .$$

Using the crossing estimate for phase B, we now show that from the perspective of existence, one can improve the MV proportionality factor between geometric and path distance from exponential to linear in dimension.

THEOREM 3.3. For $\mathbf{x}, \mathbf{y} \in \mathcal{L}$, we have that

$$(1/2)\|\mathbf{x} - \mathbf{y}\|_{\mathcal{V}} \leq d_{\mathcal{G}}(\mathbf{x}, \mathbf{y}) \leq (n/2)\|\mathbf{x} - \mathbf{y}\|_{\mathcal{V}} .$$

Furthermore, the above is best possible, even when restricted to $\mathcal{L} = \mathbb{Z}^n$.

Proof. For the lower bound, note that $d_{\mathcal{G}}(\mathbf{x}, \mathbf{y})$ is the minimum $k \in \mathbb{Z}_+$ such that there exists $\mathbf{v}_1, \dots, \mathbf{v}_k \in \text{VR}$ satisfying $\mathbf{y} = \mathbf{x} + \sum_{i=1}^k \mathbf{v}_i$. Since $\forall \mathbf{v} \in \text{VR}$, $\|\mathbf{v}\|_{\mathcal{V}} = 2$, by the triangle inequality

$$\|\mathbf{y} - \mathbf{x}\|_{\mathcal{V}} = \left\| \sum_{i=1}^k \mathbf{v}_i \right\|_{\mathcal{V}} \leq \sum_{i=1}^k \|\mathbf{v}_i\|_{\mathcal{V}} = 2k,$$

as needed.

For the upper bound, we run the randomized straight line procedure from \mathbf{x} to \mathbf{y} , i.e. setting $\mathbf{t} = \mathbf{y}$. By Lemma 3.1, the expected path length on \mathcal{G} is

$$\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{y} + Z]|] + \mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{y} + Z, \mathbf{y}]|]$$

where $Z \sim \text{Uniform}(\mathcal{V})$. Since $\mathbf{y} \in \mathcal{L}$ and $Z \in \text{int}(\mathcal{V})$ with probability 1, note that

$$\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{y} + Z, \mathbf{y}]|] = 0 ,$$

i.e. the number of steps in phase C is 0. It therefore suffices to bound the number phase B steps. By Theorem 3.1, we have that

$$\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{y} + Z]|] \leq (n/2)\|\mathbf{x} - \mathbf{y}\|_{\mathcal{V}} ,$$

as needed. This shows the desired upper bound on the path length.

We now show that the above bounds are sharp. For the lower bound, note that it is tight for any two adjacent lattice vectors, since $\forall \mathbf{v} \in \text{VR}$, $\|\mathbf{v}\|_{\mathcal{V}} = 2$. For the upper bound, letting $\mathcal{L} = \mathbb{Z}^n$, $\mathcal{V} = [-1/2, 1/2]^n$, $\text{VR} = \{\pm \mathbf{e}_1, \dots, \pm \mathbf{e}_n\}$, the shortest path between $\mathbf{x} = \mathbf{0}$ and $\mathbf{y} = (1, \dots, 1)$ has length n , while $\|\mathbf{x} - \mathbf{y}\|_{\mathcal{V}} = 2\|\mathbf{x} - \mathbf{y}\|_{\infty} = 2$.

Since the Voronoi distance changes by at most 1 when switching from \mathbf{y} to $\mathbf{t} \in \mathbf{y} + \mathcal{V}$, we note that the above bound immediately yields a corresponding bound on the path length to a closest lattice vector to any target.

As the phase C bound in Theorem 3.2 only holds for the truncated path, it does yield a bound on the randomized straight line path length for general lattices. However, for rational lattices and targets, we now show that for α small enough, the truncated path in fact suffices.

We will derive this result from the following simple Lemmas.

LEMMA 3.2. (RATIONAL LATTICE BOUND) Let $\mathcal{L} \subseteq \mathbb{Q}^n$, and $\mathbf{t} \in \mathbb{Q}^n$. Let $\bar{q} \in \mathbb{N}$ be the smallest number such that $\bar{q}\mathcal{L} \subseteq \mathbb{Z}^n$ and $\bar{q}\mathbf{t} \in \mathbb{Z}^n$, and let $\mu = \mu(\mathcal{L})$ denote the covering radius of \mathcal{L} . For $\mathbf{y} \in \mathcal{L}$, if $\mathbf{t} \notin \mathbf{y} + \mathcal{V}$, then

$$\|\mathbf{t} - \mathbf{y}\|_{\mathcal{V}} \geq 1 + 1/(2\bar{q}\mu)^2 .$$

Proof. Note that $\mathbf{t} \notin \mathbf{y} + \mathcal{V}$ iff $\|\mathbf{t} - \mathbf{y}\|_{\mathcal{V}} > 1$. From here, we have that

$$1 < \|\mathbf{t} - \mathbf{y}\|_{\mathcal{V}} = 2 \frac{\langle \mathbf{v}, \mathbf{t} - \mathbf{y} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle},$$

for some $\mathbf{v} \in \text{VR}$. By our assumptions, we note that $\langle \mathbf{v}, \mathbf{t} - \mathbf{y} \rangle = a/q^2$, for $a \in \mathbb{N}$. Next, $\|\mathbf{v}\|_2 \leq 2\mu$ (see the end of section A.1.1 for details) and $\mathbf{v} \in \mathbb{Z}^n/q$, and hence we can write $\langle \mathbf{v}, \mathbf{v} \rangle = b/q^2$, $b \in \mathbb{N}$, for $b \leq (2\bar{q}\mu)^2$. Therefore $1 < \|\mathbf{t} - \mathbf{y}\|_{\mathcal{V}} = \frac{2a}{b}$ implies that $2a > b$. Since $a, b \in \mathbb{N}$, we must have that

$$\|\mathbf{t} - \mathbf{y}\|_{\mathcal{V}} = \frac{2a}{b} \geq \frac{b+1}{b} = 1 + \frac{1}{b} \geq 1 + \frac{1}{(2\bar{q}\mu)^2}$$

as needed.

The following shows that the relevant quantities in Lemma 3.2 can be bounded by the binary encoding length of the lattice basis and target. Since it is rather standard, we defer the proof to the appendix.

LEMMA 3.3. (BIT LENGTH BOUND) Let $B \in \mathbb{Q}^{n \times n}$ be a lattice basis matrix for an n dimensional lattice \mathcal{L} ,

with $B_{ij} = \frac{p_{ij}^B}{q_{ij}^B}$ where $p_{ij}^B \in \mathbb{Z}$ and $q_{ij}^B \in \mathbb{N}$. Let $\mathbf{t} \in \mathbb{Q}^n$, with $\mathbf{t}_i = \frac{p_i^t}{q_i^t}$, $p_i^t \in \mathbb{Z}$, $q_i^t \in \mathbb{N}$. Then for $\bar{q} \in \mathbb{N}$, the smallest number such that $\bar{q}\mathcal{L} \subseteq \mathbb{Z}^n$ and $\bar{q}\mathbf{t} \in \mathbb{Z}^n$, we have that $\log_2(\bar{q}\mu(\mathcal{L})) \leq \text{enc}(B) + \text{enc}(\mathbf{t})$ and $\log_2(\mu(\mathcal{L})/\lambda_1(\mathcal{L})) \leq \text{enc}(B)$.

We are now in a position to give our full CVPP algorithm.

THEOREM 3.4. (CVPP ALGORITHM) *Let \mathcal{L} be an n -dimensional lattice with basis $B \in \mathbb{Q}^{n \times n}$, let VR denote the set of Voronoi relevant vectors of \mathcal{L} . Given the set VR as preprocessing, for any target $\mathbf{t} \in \mathbb{Q}^n$, a closest lattice vector to \mathbf{t} can be computed using an expected $\text{poly}(n, \text{enc}(B), \text{enc}(\mathbf{t}))|\text{VR}|$ arithmetic operations.*

Proof. To start we pick linearly independent $\mathbf{v}_1, \dots, \mathbf{v}_n \in \text{VR}$. We then compute the coefficient representation of \mathbf{t} with respect to $\mathbf{v}_1, \dots, \mathbf{v}_n$, that is $\mathbf{t} = \sum_{i=1}^n a_i \mathbf{v}_i$. From here we compute the lattice vector $\mathbf{x} = \sum_{i=1}^n \lceil a_i \rceil \mathbf{v}_i$, i.e. the rounding of \mathbf{t} .

Next, using the convex body sampler (Theorem A.2), we compute a $(1/4)$ -uniform sample Z over \mathcal{V} . Note that a membership oracle for \mathcal{V} can be implemented using $O(n|\text{VR}|)$ arithmetic operations. Furthermore, letting $\lambda_1 = \lambda_1(\mathcal{L})$, $\mu = \mu(\mathcal{L})$, we have that

$$(\lambda_1/2)B_2^n \subseteq \mathcal{V} \subseteq \mu B_2^n,$$

where $\lambda_1 = \min_{\mathbf{v} \in \text{VR}} \|\mathbf{v}\|_2$, $(1/2)\max_{\mathbf{v} \in \text{VR}} \|\mathbf{v}\|_2 \leq \mu \leq (\sqrt{n}/2)\max_{\mathbf{v} \in \text{VR}} \|\mathbf{v}\|_2$ (see Lemma A.2 in the Appendix). Hence, nearly tight sandwiching estimates for \mathcal{V} can be easily computed using the set VR .

We now run the randomized straight line algorithm starting at lattice point \mathbf{x} , perturbation Z , and target \mathbf{t} . If the path produced by the algorithm becomes longer than $cn(n + (\text{enc}(B) + \text{enc}(\mathbf{t})))$ (for some $c \geq 1$ large enough), restart the algorithm, and otherwise return the found closest lattice vector.

The correctness of the algorithm follows directly from the correctness of the randomized straight line algorithm (Lemma 3.1), and hence we need only show a bound on the expected runtime.

Runtime. We first bound the number of operations performed in a single iteration. Computing $\mathbf{v}_1, \dots, \mathbf{v}_n, \mathbf{t}$, and the sandwiching estimates for \mathcal{V} , requires at most $O(n^3|\text{VR}|)$ arithmetic operations. By Lemma 3.3, the convex body sampler requires at most

$$\text{poly}(n, \log(\sqrt{n}\mu/\lambda_1))|\text{VR}| = \text{poly}(n, \text{enc}(B))|\text{VR}|$$

arithmetic operations. For the randomized straight line algorithm, each step requires at most $O(n|\text{VR}|)$ arithmetic operations by Lemma 3.1. Since we truncate

it at $O(n(n + (\text{enc}(B) + \text{enc}(\mathbf{t}))))$ iterations, this requires at most $O(n^2(n + (\text{enc}(B) + \text{enc}(\mathbf{t})))|\text{VR}|)$ arithmetic operations. Hence the total number of arithmetic operations per iteration is bounded by

$$\text{poly}(n, \text{enc}(B), \text{enc}(\mathbf{t}))|\text{VR}|.$$

We now show that the algorithm performs at most $O(1)$ iterations on expectation. For this it suffices to show that each iteration succeeds with constant probability. In particular, we will show that with constant probability, the length of the randomized straight line path is bounded by $O(n^2(\text{enc}(B) + \text{enc}(\mathbf{t})))$. To do this we will simply show that the expected path length is bounded by $O(n^2(\text{enc}(B) + \text{enc}(\mathbf{t})))$ under the assumption that Z is truly uniform. By Markov's inequality, the probability that the length is less than twice the expectation is at least $1/2$ for a truly uniform Z , and hence it will be at least $1/4$ for a $1/4$ -uniform Z .

To begin, we note that by the triangle inequality

$$\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}} \leq \sum_{i=1}^n |a_i - \lceil a_i \rceil| \|\mathbf{v}_i\|_{\mathcal{V}} \leq \sum_{i=1}^n (1/2)(2) = n.$$

Let \bar{q} be as in Lemma 3.3, and let $\alpha = \frac{1}{(4\bar{q}\mu)^2}$, where we have that $\ln(1/\alpha) = O(\text{enc}(B) + \text{enc}(\mathbf{t}))$. Let $\mathbf{y} \in \mathcal{L}$ denote the center of the first Voronoi cell containing $\mathbf{t} + \alpha Z$ found by the randomized straight line algorithm. We claim that \mathbf{y} is a closest lattice vector to \mathbf{t} , or equivalently that $\mathbf{t} \in \mathbf{y} + \mathcal{V}$. Assume not, then by Lemma 3.2, $\|\mathbf{t} - \mathbf{y}\|_{\mathcal{V}} \geq 1 + \frac{1}{(2\bar{q}\mu)^2}$. On the other hand, since $\mathbf{t} + \alpha Z \in \mathbf{y} + \mathcal{V}$ and $Z \in \mathcal{V}$, by the triangle inequality

$$\begin{aligned} \|\mathbf{t} - \mathbf{y}\|_{\mathcal{V}} &\leq \|\mathbf{t} - (\mathbf{t} + \alpha Z)\|_{\mathcal{V}} + \|(\mathbf{t} + \alpha Z) - \mathbf{y}\|_{\mathcal{V}} \\ &\leq \alpha + 1 = 1 + \frac{1}{(4\bar{q}\mu)^2}, \end{aligned}$$

a contradiction. Hence \mathbf{y} is a closest lattice vector to \mathbf{t} . If $Z \sim \text{Uniform}(\mathcal{V})$, then by Theorems 3.1 and 3.2 the expected length of the randomized straight line path until $\mathbf{t} + \alpha Z$ (i.e. until we find \mathbf{y}) is bounded by

$$\begin{aligned} (n/2)\|\mathbf{t} - \mathbf{x}\|_{\mathcal{V}} + \frac{e^2}{\sqrt{2}-1}n(2 + \ln(4/\alpha)) \\ = n^2/2 + \frac{e^2}{\sqrt{2}-1}n(2 + 2\ln(8\bar{q}\mu)) \\ = O(n(n + (\text{enc}(B) + \text{enc}(\mathbf{t})))) , \end{aligned}$$

as needed. The theorem thus follows.

4 Open Problems

Our work here raises a number of natural questions. Firstly, given the improvement for CVPP, it is natural

to wonder whether any of the insights developed here can be used to improve the complexity upper bound for CVP. As mentioned previously, this would seem to require new techniques, and we leave this as an open problem.

Secondly, while we now have a number of methods to navigate over the Voronoi graph, we have no lower bounds on the lengths of the path they create. In particular, it is entirely possible that either the MV path or the simple deterministic straight line path, also yield short paths on the Voronoi graph. Hence, showing either strong lower bounds for these methods or new upper bounds is an interesting open problem. In this vein, as mentioned previously, we do not know whether the expected number of iterations for the randomized straight line procedure is inherently weakly polynomial. We leave this as an open problem.

References

- [1] Dorit Aharonov and Oded Regev. Lattice problems in $NP \cap coNP$. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS 2004.
- [2] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610, 2001.
- [3] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. Sampling short lattice vectors and the closest lattice vector problem. In *IEEE Conference on Computational Complexity*, pages 53–57, 2002.
- [4] Mikhail Alekhovich, Subhash Khot, Guy Kindler, and Nisheeth K. Vishnoi. Hardness of approximating the closest vector problem with pre-processing. *Computational Complexity*, 20(4):741–753, 2011.
- [5] László Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. Preliminary version in STACS 1985.
- [6] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based fhe as secure as pke. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 1–12, 2014.
- [7] D. Dadush, N. Stephen Davidowitz, and O. Regev. On the closest vector problem with a distance guarantee. In *Proceedings of the Conference on Computational Complexity (CCC)*, 2014.
- [8] Rudi de Buda. Some optimal codes have structure. *IEEE Journal on Selected Areas in Communications*, 7(6):893–899, 1989.
- [9] M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991. Preliminary version in STOC 1989.
- [10] Uri Erez and Ram Zamir. Achieving $1/2 \log(1+SNR)$ on the AWGN channel with lattice encoding and decoding. *IEEE Transactions on Information Theory*, 50(10):2293–2314, 2004.
- [11] 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.
- [12] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [13] Craig Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *CRYPTO*, pages 116–137, 2010.
- [14] Guillaume Hanrot and Damien Stehlé. Improved analysis of Kannan’s shortest lattice vector algorithm. In *Proceedings of the 27th annual international cryptology conference on Advances in cryptology, CRYPTO’07*, pages 170–186, Berlin, Heidelberg, 2007. Springer-Verlag.
- [15] Bettina Helfrich. Algorithms to construct Minkowski reduced and Hermite reduced lattice bases. *Theor. Comput. Sci.*, 41:125–139, December 1985.
- [16] Ravi Kannan. Minkowski’s convex body theorem and integer programming. *Mathematics of operations research*, 12(3):415–440, August 1987. 1987.
- [17] Subhash Khot, Preyas Papat, and Nisheeth K. Vishnoi. $2^{\log^{1-\epsilon} n}$ hardness for the closest vector problem with preprocessing. In *STOC*, pages 277–288, 2012.
- [18] J. C. Lagarias, Hendrik W. Lenstra Jr., and Claus-Peter Schnorr. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990.
- [19] J. C. Lagarias and Andrew M. Odlyzko. Solving low-density subset sum problems. *J. ACM*, 32(1):229–246, 1985.
- [20] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.
- [21] Hendrik W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, November 1983.
- [22] L. Lovász and S. Vempala. Fast algorithms for log-concave functions: Sampling, rounding, integration and optimization. In *FOCS ’06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 57–68, Washington, DC, USA, 2006. IEEE Computer Society.
- [23] R.G. McKilliam, A. Grant, and I.V. Clarkson. Finding a closest point in lattices of Voronoi’s first kind. Arxiv report 1405.7014, 2014.
- [24] D. Micciancio and P. 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. Preliminary version in STOC 2010.
- [25] Daniele Micciancio. The hardness of the closest vector problem with preprocessing. *IEEE Transactions on Information Theory*, 47(3):1212–1215, 2001.
- [26] Daniele Micciancio and Michael Walter. Fast lattice point enumeration with minimal overhead. In *To appear in the Proceedings of the ACM-SIAM Symposium*

on *Discrete Algorithms*, 2015.

- [27] 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.
- [28] Naftali Sommer, Meir Feder, and Ofir Shalvi. Finding the closest lattice point by iterative slicing. *SIAM Journal on Discrete Mathematics*, 23(2):715–731, 2009.

A Preliminaries

Basics. For $n \geq 1$, we denote $\mathbb{R}^n, \mathbb{Q}^n, \mathbb{Z}^n$ to be the set of n dimensional real / rational / integral vectors respectively. We let \mathbb{N} denote the set of natural numbers, and \mathbb{Z}_+ denote the set of non-negative integers. For two sets $A, B \subseteq \mathbb{R}^n$, we denote their Minkowski sum $A + B = \{\mathbf{a} + \mathbf{b} : \mathbf{a} \in A, \mathbf{b} \in B\}$. We write ∂A to denote the topological boundary of A . For a set $A \subseteq \mathbb{R}^n$, its affine hull, $\text{affhull}(A)$, is the inclusion wise smallest linear affine space containing A . We denote the interior of A in \mathbb{R}^n as $\text{int}(A)$. We denote the relative interior of A by $\text{relint}(A)$, which is the interior of A with the respect to the subspace topology on $\text{affhull}(A)$.

For two n dimensional vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we denote their inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i$. The ℓ_2 (Euclidean) norm of a vector \mathbf{x} is denoted $\|\mathbf{x}\|_2 = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$. We let $B_2^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq 1\}$ denote the unit Euclidean ball, and let $S^{n-1} = \partial B_2^n$ denote the unit sphere. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we denote the closed line segment from \mathbf{x} to \mathbf{y} by $[\mathbf{x}, \mathbf{y}] \stackrel{\text{def}}{=} \{\alpha \mathbf{x} + (1 - \alpha) \mathbf{y} : \alpha \in [0, 1]\}$, and (\mathbf{x}, \mathbf{y}) the half open line segment not containing \mathbf{y} .

We denote $\mathbf{e}_1, \dots, \mathbf{e}_n$ the vectors of the standard basis of \mathbb{R}^n , that is the vectors such that \mathbf{e}_i has a 1 in the i^{th} coordinate and 0’s elsewhere.

Binary encoding. For an integer $z \in \mathbb{Z}$, the standard binary encoding for z requires $1 + \lceil \log_2(|z| + 1) \rceil$ bits, which we denote $\text{enc}(z)$. For a rational number $\frac{p}{q} \in \mathbb{Q}$, $p \in \mathbb{Z}$, $q \in \mathbb{N}$, the encoding size of $\frac{p}{q}$ is $\text{enc}\left(\frac{p}{q}\right) = \text{enc}(p) + \text{enc}(q)$. For an $n \times m$ matrix $M \in \mathbb{Q}^{m \times n}$ or vector $\mathbf{a} \in \mathbb{Q}^n$, $\text{enc}(M)$, $\text{enc}(\mathbf{a})$ denotes the sum of encoding lengths of all the entries.

Integration. We denote the k -dimensional Lebesgue measure in \mathbb{R}^n by $\text{vol}_k(\cdot)$. Only $k = n$ and $k = n - 1$ will be used in this paper. For $k = n - 1$, we will only apply it to sets which can be written as a disjoint countable union of $n - 1$ dimensional flat pieces. When integrating a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ over a set $A \subseteq \mathbb{R}^n$ using the n dimensional Lebesgue measure, we use the notation $\int_A f(\mathbf{x}) d\mathbf{x}$. When integrating with respect to the $n - 1$ dimensional Lebesgue measure in \mathbb{R}^n , we write $\int_A f(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x})$.

Probability. For a random variable $X \in \mathbb{R}$, we define its expectation by $\mathbb{E}[X]$ and its variance by

$\text{VAR}[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$. For two random variables $X, Y \in \Omega$, we define their total variation distance to be

$$d_{TV}(X, Y) = \max_{A \subseteq \Omega} |\Pr[X \in A] - \Pr[Y \in A]|.$$

DEFINITION A.1. (UNIFORM DISTRIBUTION) For a set $A \subseteq \mathbb{R}^n$, we define the uniform distribution on A , denoted $\text{Uniform}(A)$, to have probability density function $1/\text{vol}_n(A)$ and 0 elsewhere. That is, for a uniform random variable $X \sim \text{Uniform}(A)$, we have that

$$\Pr[X \in B] = \text{vol}_n(A \cap B) / \text{vol}_n(A)$$

for any measurable set $B \subseteq \mathbb{R}^n$.

Complexity. We use the notation $\tilde{O}(T(n))$ to mean $O(T(n)\text{polylog}(T(n)))$.

A.1 Lattices An n dimensional lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is a discrete subgroup of \mathbb{R}^n whose linear span is \mathbb{R}^n . Equivalently, \mathcal{L} is generated by all integer combinations of some basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ of \mathbb{R}^n , i.e. $\mathcal{L} = B\mathbb{Z}^n$. For $k \in \mathbb{N}$, we define the quotient group $\mathcal{L}/k\mathcal{L} = \{\mathbf{y} + k\mathcal{L} : \mathbf{y} \in \mathcal{L}\}$. It is easy to check that the map $\mathbf{a} \rightarrow B\mathbf{a} + k\mathcal{L}$ from $(\mathbb{Z}/k\mathbb{Z})^n \stackrel{\text{def}}{=} \mathbb{Z}_k^n$ to $\mathcal{L}/(k\mathcal{L})$ is an isomorphism. In particular $|\mathcal{L}/(k\mathcal{L})| = k^n$.

A shift $\mathcal{L} + \mathbf{t}$ of \mathcal{L} is called a coset of \mathcal{L} . The set of cosets of \mathcal{L} form a group \mathbb{R}^n/\mathcal{L} under addition, i.e. the torus. We will use the notation $A \pmod{\mathcal{L}}$, for a set $A \subseteq \mathbb{R}^n$, to denote the set of cosets $\mathcal{L} + A$. Note that \mathbb{R}^n/\mathcal{L} is isomorphic to $[0, 1)^n$ under addition $\pmod{1}$ (coordinate wise), under the map $\mathbf{x} \rightarrow B\mathbf{x} + \mathcal{L}$ for any basis B of \mathcal{L} . We will need to make use of the uniform distribution over \mathbb{R}^n/\mathcal{L} , which we denote $\text{Uniform}(\mathbb{R}^n/\mathcal{L})$. To obtain a sample from $\text{Uniform}(\mathbb{R}^n/\mathcal{L})$, one can take $U \sim \text{Uniform}([0, 1)^n)$ and return $BU \pmod{\mathcal{L}}$.

We denote the length of the shortest non-zero vector (or minimum distance) of \mathcal{L} as $\lambda_1(\mathcal{L}) = \min_{\mathbf{y} \in \mathcal{L} \setminus \{0\}} \|\mathbf{y}\|_2$. We denote the covering radius of \mathcal{L} as $\mu(\mathcal{L}) = \max_{\mathbf{t} \in \mathbb{R}^n} \min_{\mathbf{y} \in \mathcal{L}} \|\mathbf{t} - \mathbf{y}\|_2$ to be the farthest distance between any point in space and the lattice.

The following standard lemma (see for instance [5]) allows us to bound the covering radius:

LEMMA A.1. Let \mathcal{L} be an n -dimensional lattice. If $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathcal{L}$ are linearly independent lattice vectors, then $\mu(\mathcal{L}) \leq \frac{1}{2} \sqrt{\sum_{i=1}^n \|\mathbf{v}_i\|^2}$.

A.1.1 Voronoi cell, tiling, and relevant vectors For a point $\mathbf{t} \in \mathbb{R}^n$, let $\text{CVP}(\mathcal{L}, \mathbf{t}) = \arg \min_{\mathbf{x} \in \mathcal{L}} \|\mathbf{t} - \mathbf{x}\|_2$, denote the set of closest lattice vectors to \mathbf{t} . For $\mathbf{y} \in \mathcal{L}$,

let

$$H_{\mathbf{y}} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq \|\mathbf{x} - \mathbf{y}\|_2\} \\ = \left\{ \mathbf{x} \in \mathbb{R}^n : \langle \mathbf{y}, \mathbf{x} \rangle \leq \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle \right\},$$

denote the halfspace defining the set of points closer to $\mathbf{0}$ than to \mathbf{y} .

DEFINITION A.2. (VORONOI CELL) *The Voronoi cell $\mathcal{V}(\mathcal{L})$ of \mathcal{L} is defined as*

$$\mathcal{V}(\mathcal{L}) = \bigcap_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} H_{\mathbf{y}},$$

the set of all points in \mathbb{R}^n closer or at equal distance to the origin than to any other lattice point.

Naturally, $\mathcal{V}(\mathcal{L})$ is the set of points of \mathcal{L} whose closest lattice vector is $\mathbf{0}$. We abbreviate $\mathcal{V}(\mathcal{L})$ to \mathcal{V} when the context is clear. It is easy to check from the definitions that a vector $\mathbf{y} \in \mathcal{L}$ is a closest lattice vector to a target $\mathbf{t} \in \mathbb{R}^n$ iff $\mathbf{t} - \mathbf{y} \in \mathcal{V}$. The CVP is then equivalent to finding a lattice shift of \mathcal{V} containing the target.

From this, we see that the Voronoi cell tiles space with respect to \mathbb{R}^n , that is, the set of shifts $\mathcal{L} + \mathcal{V}$ cover \mathbb{R}^n , and shifts $\mathbf{x} + \mathcal{V}$ and $\mathbf{y} + \mathcal{V}$, $\mathbf{x}, \mathbf{y} \in \mathcal{L}$, are interior disjoint if $\mathbf{x} \neq \mathbf{y}$. From the tiling property, we have the useful property that the distribution $\text{Uniform}(\mathcal{V}) \pmod{\mathcal{L}}$ is identical to $\text{Uniform}(\mathbb{R}^n/\mathcal{L})$.

We note that the problem of separating over the Voronoi cell reduces directly to CVP, since if $\mathbf{y} \in \mathcal{L}$ is closer to a target \mathbf{t} than $\mathbf{0}$, then $H_{\mathbf{y}}$ separates \mathbf{t} from \mathcal{V} . Also, if no such closer lattice vector exists, then $\mathbf{t} \in \mathcal{V}$.

DEFINITION A.3. (VORONOI RELEVANT VECTORS) *We define $\text{VR}(\mathcal{L})$, the Voronoi relevant vectors of \mathcal{L} , to be the minimal set of lattice vectors satisfying $\mathcal{V}(\mathcal{L}) = \bigcap_{\mathbf{v} \in \text{VR}(\mathcal{L})} H_{\mathbf{v}}$, which we abbreviate to VR when the context is clear.*

Since the Voronoi cell is a full dimensional centrally symmetric polytope, the set VR corresponds exactly to the set of lattice vectors inducing facets of \mathcal{V} (i.e. such that $\mathcal{V} \cap \partial H_{\mathbf{v}}$ is $n - 1$ dimensional).

DEFINITION A.4. (VORONOI CELL FACET) *For each $\mathbf{v} \in \text{VR}$, let*

$$F_{\mathbf{v}} = \mathcal{V} \cap \left\{ \mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{v} \rangle = \frac{1}{2} \langle \mathbf{v}, \mathbf{v} \rangle \right\},$$

denote the facet of \mathcal{V} induced by \mathbf{v} .

Here we have that

$$\partial \mathcal{V} = \bigcup_{\mathbf{v} \in \text{VR}} F_{\mathbf{v}} \quad \text{and} \quad \text{vol}_{n-1}(\partial \mathcal{V}) = \sum_{\mathbf{v} \in \text{VR}} \text{vol}_{n-1}(F_{\mathbf{v}})$$

since the intersection of distinct facets has affine dimension at most $n - 2$. Similarly,

$$\mathcal{V} = \bigcup_{\mathbf{v} \in \text{VR}} \text{conv}(\mathbf{0}, F_{\mathbf{v}}) \quad \text{and} \\ \text{vol}_n(\mathcal{V}) = \sum_{\mathbf{v} \in \text{VR}} \text{vol}_n(\text{conv}(\mathbf{0}, F_{\mathbf{v}})).$$

A central object in this paper will be $\mathcal{L} + \partial \mathcal{V}$, the boundary of the lattice tiling. We shall call $\mathbf{y} + F_{\mathbf{v}}$, for $\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}$, a facet of $\mathcal{L} + \partial \mathcal{V}$. Here, we see that

$$\mathcal{L} + \partial \mathcal{V} = \bigcup_{\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}} \mathbf{y} + F_{\mathbf{v}}.$$

Note that each facet is counted twice in the above union, i.e. $\mathbf{y} + F_{\mathbf{v}} = (\mathbf{y} + \mathbf{v}) + F_{-\mathbf{v}}$.

An important theorem of Voronoi classifies the set of Voronoi relevant vectors:

THEOREM A.1. (VORONOI) *For an n dimensional lattice \mathcal{L} , $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ is in $\text{VR}(\mathcal{L})$ if and only if*

$$\{\pm \mathbf{y}\} = \arg \min_{\mathbf{x} \in 2\mathcal{L} + \mathbf{y}} \|\mathbf{x}\|_2.$$

In particular, $|\text{VR}| \leq 2(2^n - 1)$.

Here the bound on $|\text{VR}|$ follows from the fact that the map $\mathbf{y} \mapsto \mathbf{y} + 2\mathcal{L}$ from VR to $\mathcal{L}/(2\mathcal{L}) \setminus \{2\mathcal{L}\}$ is 2-to-1. Furthermore, note that each Voronoi relevant vector can be recovered from solutions to CVPs over $2\mathcal{L}$. More precisely, given a basis B for \mathcal{L} , each vector in $\mathbf{v} \in \text{VR}$ can be expressed as $B\mathbf{p} - \mathbf{x}$, for some $\mathbf{p} \in \{0, 1\}^n \setminus \{\mathbf{0}\}$, and $\mathbf{x} \in \text{CVP}(2\mathcal{L}, B\mathbf{p})$ (we get a Voronoi relevant iff \mathbf{x} is unique up to reflection about $B\mathbf{p}$).

We now list several important and standard properties we will need about the Voronoi cell and relevant vectors. We give a proof for completeness.

LEMMA A.2. *For an n dimensional lattice \mathcal{L} :*

1. $\frac{\lambda_1(\mathcal{L})}{2} B_2^n \subseteq \mathcal{V} \subseteq \mu(\mathcal{L}) B_2^n$.
2. $\lambda_1(\mathcal{L}) = \min_{\mathbf{v} \in \text{VR}} \|\mathbf{v}\|_2$.
3. $2\mu(\mathcal{L})/\sqrt{n} \leq \max_{\mathbf{v} \in \text{VR}} \|\mathbf{v}\|_2 \leq 2\mu(\mathcal{L})$

Proof. We prove each of the above in order:

1. Since each vector $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ satisfies $\|\mathbf{y}\|_2 \geq \lambda_1(\mathcal{L})$, we clearly have that $\lambda_1(\mathcal{L})/2 B_2^n \subseteq H_{\mathbf{y}}$. The inner containment holds for \mathcal{V} since $\mathcal{V} = \bigcap_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} H_{\mathbf{y}}$. For the outer containment, note that for any $\mathbf{t} \in \mathcal{V}$, that $\mathbf{0}$ is a closest lattice vector to \mathbf{t} . Hence, by definition, $\|\mathbf{t}\|_2 = \|\mathbf{t} - \mathbf{0}\|_2 \leq \mu(\mathcal{L})$ as needed.

2. Since the set $VR \subseteq \mathcal{L} \setminus \{\mathbf{0}\}$, the vectors in VR clearly have length greater than or equal to $\lambda_1(\mathcal{L})$. Next, let $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ denote a shortest non-zero vector of \mathcal{L} . We wish to show that $\mathbf{y} \in VR$. To do this, by Theorem A.1, we need only show that the only vectors of length $\lambda_1(\mathcal{L})$ in $\mathbf{y} + 2\mathcal{L}$ are $\pm\mathbf{y}$. Assume not, then there exists $\mathbf{z} \in \mathbf{y} + 2\mathcal{L}$, such that \mathbf{z} is not collinear with \mathbf{y} having $\|\mathbf{z}\|_2 = \lambda_1(\mathcal{L})$. But then note that $(\mathbf{y} + \mathbf{z})/2 \in \mathcal{L} \setminus \{\mathbf{0}\}$ and $\|(\mathbf{y} + \mathbf{z})/2\|_2 < \lambda_1(\mathcal{L})$, a contradiction.
3. For $\mathbf{v} \in VR$, we remember that $\mathbf{v} = \arg \min_{\mathbf{z} \in 2\mathcal{L} + \mathbf{v}} \|\mathbf{z}\|_2$. In particular, this implies that $\|\mathbf{v}\|_2 \leq \mu(2\mathcal{L}) = 2\mu(\mathcal{L})$ as needed. Since the VR vectors span \mathbb{R}^n , we can find linearly independent $\mathbf{v}_1, \dots, \mathbf{v}_n \in VR$. By Lemma A.1, we have that

$$\mu(\mathcal{L}) \leq (1/2) \sqrt{\sum_{i=1}^n \|\mathbf{v}_i\|_2^2} \leq \sqrt{n}/2 \max_{\mathbf{v} \in VR} \|\mathbf{v}\|_2,$$

as needed.

A.2 Convex geometry A set $K \subseteq \mathbb{R}^n$ is a convex body, if it is convex (i.e. $\mathbf{x}, \mathbf{y} \in K \Rightarrow [\mathbf{x}, \mathbf{y}] \subseteq K$), compact and has non-empty interior. K is symmetric if $K = -K$. For a symmetric convex body $K \subseteq \mathbb{R}^n$, we define the norm (or gauge function) with respect to K by $\|\mathbf{x}\|_K = \inf \{s \geq 0 : \mathbf{x} \in sK\}$, for any $\mathbf{x} \in \mathbb{R}^n$. A function $f : K \rightarrow \mathbb{R}$ is convex (concave) if for all $\mathbf{x}, \mathbf{y} \in K$, $\alpha \in [0, 1]$,

$$\alpha f(\mathbf{x}) + (1 - \alpha)f(\mathbf{y}) \geq (\leq) f(\alpha\mathbf{x} + (1 - \alpha)\mathbf{y}).$$

For a set $A \subseteq \mathbb{R}^n$, we define its convex hull $\text{conv}(A)$ to be the (inclusion wise) smallest convex set containing A . For two sets $A, B \subseteq \mathbb{R}^n$, we use the notation $\text{conv}(A, B) \stackrel{\text{def}}{=} \text{conv}(A \cup B)$.

For two non-empty measurable sets $A, B \subseteq \mathbb{R}^n$ such that $A + B$ is measurable, the Brunn-Minkowski inequality gives the following fundamental lower bound

$$(1.1) \quad \text{vol}_n(A + B)^{1/n} \geq \text{vol}_n(A)^{1/n} + \text{vol}_n(B)^{1/n}.$$

Laplace Distributions. We define the Gamma function, $\Gamma(k) = \int_0^\infty x^{k-1} e^{-x} dx$ for $k > 0$. For $k \in \mathbb{N}$, we note that $\Gamma(k) = (k - 1)!$. We define the two parameter distribution $\Gamma(k, \theta)$ on \mathbb{R} , $k, \theta \geq 0$, to have probability density function $\frac{1}{\theta^k \Gamma(k)} x^{k-1} e^{-x/\theta}$, for $x \in \mathbb{R}$. For $r \sim \Gamma(k, \theta)$, $k \in \mathbb{N}$, the moments of r are

$$\mathbb{E}[r^l] = \theta^l \frac{(k + l - 1)!}{(k - 1)!}, \text{ for } l \in \mathbb{N}.$$

In particular, $\mathbb{E}[r] = k\theta$ and $\text{VAR}[r] = k\theta^2$.

DEFINITION A.5. (LAPLACE DISTRIBUTION) We define the probability distribution $\text{Laplace}(K, \theta)$, with probability density function

$$f_K^\theta(\mathbf{x}) = \frac{\theta^n}{\text{vol}_n(K)n!} e^{-\|\mathbf{x}\|_K/\theta}, \text{ for } \mathbf{x} \in \mathbb{R}^n.$$

Equivalently, a well known and useful fact (which we state without proof) is:

LEMMA A.3. $X \sim \text{Laplace}(K, \theta)$ is identically distributed to rU , where $r \sim \Gamma(n + 1, \theta)$ and $U \sim \text{Uniform}(K)$ are sampled independently.

For our purposes, $\text{Laplace}(K, \theta)$ will serve as a “smoothed” out version of $\text{Uniform}(K)$. In particular, letting f denote the probability density function of $\text{Laplace}(K, \theta)$, for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, by the triangle inequality

$$(1.2) \quad \frac{f_K^\theta(\mathbf{x})}{f_K^\theta(\mathbf{y})} = \frac{e^{-\|\mathbf{x}\|_K/\theta}}{e^{-\|\mathbf{y}\|_K/\theta}} \in [e^{-\|\mathbf{y}-\mathbf{x}\|_K/\theta}, e^{\|\mathbf{y}-\mathbf{x}\|_K/\theta}].$$

Hence, the density varies smoothly as a function of $\|\cdot\|_K$ norm, avoiding the “sharp” boundaries of the uniform measure on K .

Algorithms. A membership oracle O_K for a convex body $K \subseteq \mathbb{R}^n$ is a function satisfying $O_K(\mathbf{x}) = 1$ if $\mathbf{x} \in K$, and $O_K(\mathbf{x}) = 0$ otherwise. Most algorithms over convex bodies can be implemented using only a membership oracle with some additional guarantees.

In our CVPP algorithm, we will need to sample nearly uniformly from the Voronoi cell. For this purpose, we will utilize the classic geometric random walk method of Dyer, Frieze, and Kannan [9], which allows for polynomial time near uniform sampling over any convex body.

THEOREM A.2. (CONVEX BODY SAMPLER [9]) Let $K \subseteq \mathbb{R}^n$ be a convex body, given my a membership oracle O_K , satisfying $rB_2^n \subseteq K \subseteq RB_2^n$. Then for $\varepsilon > 0$, a realisation of a random variable $X \in K$, having total variation distance at most ε from $\text{Uniform}(K)$, can be computed using $\text{poly}(n, \log(R/r), \log(1/\varepsilon))$ arithmetic operations and calls to the membership oracle.

B Bounding the Number of Crossings

In this section, we prove bounds on the number of crossings the randomized straight line algorithm induces on the tiling boundary $\mathcal{L} + \partial\mathcal{V}$. For a target \mathbf{t} , starting point $\mathbf{x} \in \mathcal{L}$, and perturbation $Z \sim \text{Uniform}(V)$, we need to bound the expected number of crossings in phases B and C, that is

$$(B) \mathbb{E}[|(\mathcal{L} + \partial) \cap [\mathbf{x} + Z, \mathbf{t} + Z]|]$$

$$(C) \mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{t} + Z, \mathbf{t}]|].$$

The phase B bound is given in Section B.1, and the phase C is given in Section B.2.

B.1 Phase B estimates The high level idea of the phase B bound is as follows. To count the number of crossings, we break the segment $[\mathbf{x} + Z, \mathbf{y} + Z]$ into k equal chunks (we will let $k \rightarrow \infty$), and simply count the number of chunks which cross at least 1 boundary. By our choice of perturbation, we can show that each point on the segment $[\mathbf{x} + Z, \mathbf{y} + Z]$ is uniformly distributed modulo the lattice, and hence the crossing probability will be identical on each chunk. In particular, we will get that each crossing probability is exactly the probability that Z “escapes” from \mathcal{V} after moving by $(\mathbf{y} - \mathbf{x})/k$. This measures how close Z tends to be to the boundary of \mathcal{V} , and hence corresponds to a certain “directional” surface area to volume ratio.

In the next lemma, we show that the escape probability is reasonably small for any symmetric convex body, when the size of the shift is measured using the norm induced by the body. We shall use this to prove the full phase B crossing bound in Theorem 3.1.

LEMMA B.1. *Let $K \subseteq \mathbb{R}^n$ be a centrally symmetric convex body. Then for $Z \sim \text{Uniform}(K)$ and $\mathbf{y} \in \mathbb{R}^n$, we have that*

$$\lim_{\varepsilon \rightarrow 0} \Pr[Z + \varepsilon \mathbf{y} \notin K] / \varepsilon \leq (n/2) \|\mathbf{y}\|_K$$

Proof. By applying a linear transformation to K and \mathbf{y} , we may assume that $\mathbf{y} = \mathbf{e}_n$. Let $\pi_{n-1} : \mathbb{R}^n \rightarrow \mathbb{R}^{n-1}$ denote the projection onto the first $n - 1$ coordinates. Define $l : \pi_{n-1}(K) \rightarrow \mathbb{R}_+$ as $l(\mathbf{x}) = \text{vol}_1(\{(\mathbf{x}, x_n) : x_n \in \mathbb{R}, (\mathbf{x}, x_n) \in K\})$, i.e. the length of the chord of K passing through $(\mathbf{x}, 0)$ in direction \mathbf{e}_n .

For $\mathbf{x} \in \pi_{n-1}(K)$, let $\{(\mathbf{x}, x_n) : x_n \in \mathbb{R}, (\mathbf{x}, x_n) \in K\} = [(\mathbf{x}, a), (\mathbf{x}, b)]$, $a \leq b$, denote its associated chord, where we note that $|b - a| = l(\mathbf{x})$. From here, conditioned on Z landing on this chord, note that $Z + \varepsilon \mathbf{e}_n \notin K$ if and only if Z lies in the half open line segment $((\mathbf{x}, b - \varepsilon), (\mathbf{x}, b)]$. Given this, we have that

$$\begin{aligned} & \lim_{\varepsilon \rightarrow 0} \Pr[Z + \varepsilon \mathbf{e}_n \notin K] / \varepsilon \\ &= \lim_{\varepsilon \rightarrow 0} (1/\varepsilon) \int_{\pi_{n-1}(K)} \min\{\varepsilon, l(\mathbf{x})\} \frac{d\mathbf{x}}{\text{vol}_n(K)} \\ (2.3) \quad &= \lim_{\varepsilon \rightarrow 0} \int_{\pi_{n-1}(K)} \min\{1, l(\mathbf{x})/\varepsilon\} \frac{d\mathbf{x}}{\text{vol}_n(K)} \\ &= \int_{\pi_{n-1}(K)} \frac{d\mathbf{x}}{\text{vol}_n(K)} = \frac{\text{vol}_{n-1}(\pi_{n-1}(K))}{\text{vol}_n(K)}. \end{aligned}$$

Let $s = 1/\|\mathbf{e}_n\|_K$. Since K is centrally symmetric, $\mathbb{R}\mathbf{e}_n \cap K = [-s\mathbf{e}_n, s\mathbf{e}_n]$ and hence $l(\mathbf{0}) = 2s$. Note that by central symmetry of K , for all $\mathbf{x} \in \pi_{n-1}(K)$,

$l(\mathbf{x}) = l(-\mathbf{x})$. Since K is convex, the function l is concave on $\pi_{n-1}(K)$, and hence

$$\begin{aligned} \max_{\mathbf{x} \in \pi_{n-1}(K)} l(\mathbf{x}) &= \max_{\mathbf{x} \in \pi_{n-1}(K)} \frac{1}{2}l(\mathbf{x}) + \frac{1}{2}l(-\mathbf{x}) \\ &\leq \max_{\mathbf{x} \in \pi_{n-1}(K)} l(\mathbf{0}) = 2s. \end{aligned}$$

Let $K' = \{(\mathbf{x}, x_n) : \mathbf{x} \in \pi_{n-1}(K), 0 \leq x_n \leq l(\mathbf{x})\}$. By concavity of l , it is easy to see that K' is also a convex set. Furthermore, note that K' has exactly the same chord lengths as K in direction \mathbf{e}_n , and hence $\text{vol}_n(K') = \text{vol}_n(K)$. For $a \in \mathbb{R}$, let $K'_a = K' \cap \{(\mathbf{x}, a) : \mathbf{x} \in \mathbb{R}^{n-1}\}$. Here $\mathbb{R}\mathbf{e}_n \cap K' = [0, 2s\mathbf{e}_n]$, and hence $K_a \neq \emptyset, \forall a \in [0, 2s]$. Therefore $K_a = \emptyset$ for $a > 2s$, since the maximum chord length is $l(\mathbf{0}) = 2s$, as well as for $a < 0$. By construction of K' , we see that $K'_0 = \pi_{n-1}(K) \times \{0\}$, and hence $\text{vol}_{n-1}(K'_0) = \text{vol}_{n-1}(\pi_{n-1}(K))$.

Given (2.3), to prove the Lemma, it now suffices to show that

$$\frac{\text{vol}_{n-1}(\pi_{n-1}(K))}{\text{vol}_n(K)} = \frac{\text{vol}_{n-1}(\pi_{n-1}(K))}{\text{vol}_n(K')} \leq (n/2) \|\mathbf{e}_n\|_K.$$

For $a \in [0, 2s]$, by convexity of K' and the Brunn-Minkowski inequality on \mathbb{R}^{n-1} , we have that

$$\begin{aligned} & \text{vol}_{n-1}(K'_a)^{\frac{1}{n-1}} \\ & \geq \text{vol}_{n-1}\left(\left(1 - \frac{a}{2s}\right)K'_0 + \frac{a}{2s}K'_{2s}\right)^{\frac{1}{n-1}} \\ & \geq \left(1 - \frac{a}{2s}\right) \text{vol}_{n-1}(K'_0)^{\frac{1}{n-1}} + \frac{a}{2s} \text{vol}_{n-1}(K'_{2s})^{\frac{1}{n-1}} \\ & \geq \left(1 - \frac{a}{2s}\right) \text{vol}_{n-1}(\pi_{n-1}(K))^{\frac{1}{n-1}}. \end{aligned}$$

Therefore, we have that

$$\begin{aligned} \text{vol}_n(K') &= \int_0^{2s} \text{vol}_{n-1}(K'_a) da \\ &\geq \text{vol}_{n-1}(\pi_{n-1}(K)) \int_0^{2s} \left(1 - \frac{a}{2s}\right)^{n-1} da \\ &= \text{vol}_{n-1}(\pi_{n-1}(K))(2s) \int_0^1 (1 - a)^{n-1} da \\ &= \text{vol}_{n-1}(\pi_{n-1}(K))(2s)/n \\ &= \frac{2\text{vol}_{n-1}(\pi_{n-1}(K))}{n\|\mathbf{e}_n\|_n}, \end{aligned}$$

as needed.

B.1.1 Proof of Theorem 3.1 (Phase B crossing bound)

Proof. Note first that the sets $(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{y} + Z]$ and $(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{y} + Z]$ agree unless $\mathbf{y} + Z \in \mathcal{L} + \partial\mathcal{V}$.

Given that this event happens with probability 0 (as $\mathcal{L} + \partial\mathcal{V}$ has n dimensional Lebesgue measure 0), we get that

$$\begin{aligned} & \mathbb{E}_{\mathcal{Z}}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{y} + Z]|] \\ &= \mathbb{E}_{\mathcal{Z}}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{y} + Z]|] . \end{aligned}$$

We now bound the expectation on the right hand side. For $s \in [0, 1]$, define the random variable $\ell(s) = (1 - s)\mathbf{x} + s\mathbf{y} + Z$. Let $A_j^k, 0 \leq j < 2^k$, denote the event that

$$\begin{aligned} & |(\mathcal{L} + \partial\mathcal{V}) \cap [\ell(j/2^k), \ell((j+1)/2^k)]| \geq 1 \\ \Leftrightarrow & |(\mathcal{L} + \partial\mathcal{V}) \cap [\ell(j/2^k), \ell(j/2^k) + (\mathbf{y} - \mathbf{x})/2^k]| \geq 1 . \end{aligned}$$

Clearly, we have that

$$|(\mathcal{L} + \partial\mathcal{V}) \cap [\ell(0), \ell(1)]| = \lim_{k \rightarrow \infty} \sum_{j=0}^{2^k-1} A_j^k .$$

By the monotone convergence theorem, we get that

$$\mathbb{E}_{\mathcal{Z}}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\ell(0), \ell(1)]|] = \lim_{k \rightarrow \infty} \sum_{j=0}^{2^k-1} \Pr[A_j^k = 1] .$$

Since $\mathcal{L} + \partial\mathcal{V}$ is by definition invariant under lattice shifts, we see that $\Pr[A_j^k = 1]$ depends only on the distribution of $\ell(j/2^k) \pmod{\mathcal{L}}$. Given that $Z \pmod{\mathcal{L}} \sim \text{uniform}(\mathbb{R}^n/\mathcal{L})$ and that \mathbb{R}^n/\mathcal{L} is shift invariant, we have that $\ell(j/2^k) \pmod{\mathcal{L}} \sim \text{uniform}(\mathbb{R}^n/\mathcal{L})$. In particular, this implies that $\Pr[A_0^k] = \dots = \Pr[A_{2^k}^k]$, and hence by Lemma B.1

$$\begin{aligned} (2.5) \quad \lim_{k \rightarrow \infty} \sum_{j=0}^{2^k-1} \Pr[A_j^k = 1] &= \lim_{k \rightarrow \infty} 2^k \Pr[A_0^k = 1] \\ &= \lim_{k \rightarrow \infty} 2^k \Pr[Z + (\mathbf{y} - \mathbf{x})/2^k \notin \mathcal{V}] \\ &\leq (n/2) \|\mathbf{y} - \mathbf{x}\|_{\mathcal{V}} , \end{aligned}$$

as needed. The result follows by combining (2.4) and (2.5).

B.2 Phase C estimates As mentioned previously in the paper, our techniques will not be sufficient to fully bound the number of phase C crossings. However, we will use be able to give bounds for a truncation of the phase C path, that is for $\alpha \in (0, 1]$, we will bound

$$\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{t} + Z, \mathbf{t} + \alpha Z]|] .$$

We will give a bound of $O(n \ln 1/\alpha)$ for the above crossings in Theorem 3.2.

For the proof strategy, we follow the approach as phase B in terms of bounding the crossing probability on infinitesimal chunks of $[\mathbf{t} + Z, \mathbf{t} + \alpha Z]$. However, the implementation of this strategy will be completely different here, since the points along the segment no longer have the same distribution modulo the lattice. In particular, as $\alpha \rightarrow 0$, the distributions get more concentrated, and hence we lose the effects of the randomness. This loss will be surprisingly mild however, as evidenced by the claimed $\ln(1/\alpha)$ dependence.

For the infinitesimal probabilities, it will be convenient to parametrize the segment $[\mathbf{t} + Z, \mathbf{t} + \alpha Z]$ differently than in phase B. In particular, we use $\mathbf{t} + Z/s$, for $s \in [1, 1/\alpha]$. From here, note that

$$\begin{aligned} (2.6) \quad & \Pr[(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{t} + Z/s, \mathbf{t} + Z/(s + \varepsilon)] \neq \emptyset] \\ &= \Pr[Z \in \cup_{\gamma \in [s, s+\varepsilon]} \gamma(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})] . \end{aligned}$$

Taking the limit as $\varepsilon \rightarrow 0$, we express the infinitesimal probability as a certain weighted surface area integral over $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ (see Lemma B.1).

In the same spirit as phase B, we will attempt to relate the surface integral to a nicely bounded integral over all of space. To help us in this task, we will rely on a technical trick to “smooth out” the distribution of Z . More precisely, we will replace the perturbation $Z \sim \text{Uniform}(\mathcal{V})$ by the perturbation $X \sim \text{Laplace}(\mathcal{V}, \theta)$, for an appropriate choice of θ . For the relationship between both types of perturbations, we will use the representation of X as rZ , where $r \sim \Gamma(n+1, \theta)$. We will choose θ so that r is concentrated in the interval $[1, 1 + \frac{1}{n}]$, which will insure that the number of crossings for X and Z are roughly the same.

The benefit of the Laplace perturbation for us will be as follows. Since it varies much more smoothly than the uniform distribution (which has “sharp boundaries”), it will allow us to make the analysis of the surface integral entirely local by using the periodic structure of $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$. In particular, we will be able to relate the surface integral over each tile $s(\mathbf{y} - \mathbf{t} + \partial\mathcal{V})$, $\mathbf{y} \in \mathcal{L}$, to a specific integral over each cone $s(\mathbf{y} + \text{conv}(\mathbf{0}, F_{\mathbf{v}}))$, $\forall \mathbf{v} \in \text{VR}$, making up the tile. Under the uniform distribution, the probability density over each tile can be challenging to analyze, since the tile may only be partially contained in \mathcal{V} . However, under the Laplace distribution, we know that over $s(\mathbf{y} + \mathcal{V})$ the density can vary by at most $e^{\pm s/\theta}$ (see Equation 1.2 in the Preliminaries).

The integral over \mathbb{R}^n we end up using to control the surface integral over $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ turns out to be rather natural. At all scales, we simply use the integral representation of $\mathbb{E}[\|X\|_{\mathcal{V}}] = n\theta$ (see Lemma B.6). In particular, as $s \rightarrow \infty$, for the appropriate choice of θ , this will allow us to bound the surface integral over

$s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ by $O(n/s)$. Integrating this bound from 1 to $1/\alpha$ yields the claimed $O(n \ln(1/\alpha))$ bound on the number of crossings.

This section is organized as follows. In subsection B.2.1, we relate the number of crossings for uniform and Laplace perturbations. In subsection B.2.2, we bound the number of crossings for Laplace perturbations. Lastly, in subsection B.2.3, we combine the estimates from the previous subsections to give the full phase C in Theorem 3.2.

B.2.1 Converting Uniform Perturbations to Laplace Perturbations In this section, we show that the number of crossings induced by uniform perturbations can be controlled by the number of crossings induced by Laplace perturbation.

We define $\theta_n = \frac{1}{(n+1) - \sqrt{2(n+1)}}$, $\gamma_n = (1 + \frac{2\sqrt{2}}{\sqrt{n+1} - \sqrt{2}})^{-1}$ for use in the rest of this section.

The following Lemma shows the $\Gamma(n+1, \theta)$ distribution is concentrated in a small interval above 1 for the appropriate choice of θ . This will be used in Lemma B.3 to relate the number of crossings between the uniform and Laplace perturbations.

LEMMA B.2. For $r \sim \Gamma(n+1, \theta_n)$, $n \geq 2$, we have that

$$\Pr[r \in [1, 1 + \frac{2\sqrt{2}}{\sqrt{n+1} - \sqrt{2}}]] \geq \frac{1}{2}$$

Proof. Remember that $\mathbb{E}[r] = (n+1)\theta_n$ and that $\text{VAR}[r] = (n+1)\theta_n^2$. Letting $\sigma = \sqrt{\text{VAR}[r]}$, by Chebyshev's inequality

$$\Pr[|r - \mathbb{E}[r]| \geq \sqrt{2}\sigma] \leq \frac{\text{VAR}[r]}{2\sigma^2} = \frac{1}{2}$$

The result now follows from the identities

$$\begin{aligned} \mathbb{E}[r] - \sqrt{2}\sigma &= ((n+1) - \sqrt{2(n+1)})\theta_n = 1 \\ \mathbb{E}[r] + \sqrt{2}\sigma &= ((n+1) + \sqrt{2(n+1)})\theta_n \\ &= 1 + \frac{2\sqrt{2}}{\sqrt{n+1} - \sqrt{2}} \end{aligned}$$

LEMMA B.3. Let \mathcal{L} be an n -dimensional lattice, $n \geq 2$, and $\mathbf{t} \in \mathbb{R}^n$. Then for $\alpha \in [0, 1]$, $Z \sim \text{Uniform}(\mathcal{V})$ and $X \sim \text{Laplace}(\mathcal{V}, \theta_n)$, we have that

$$\begin{aligned} &\mathbb{E}_Z[(\mathcal{L} + \partial\mathcal{V}) \cap [Z + \mathbf{t}, \alpha Z + \mathbf{t}]] \\ &\leq 2\mathbb{E}_X[(\mathcal{L} + \partial\mathcal{V}) \cap [X + \mathbf{t}, \gamma_n \alpha X + \mathbf{t}]] \end{aligned}$$

where $\gamma_n = (1 + \frac{2\sqrt{2}}{\sqrt{n+1} - \sqrt{2}})^{-1}$.

Proof. We shall use the fact that X is identically distributed to rZ where $r \sim \Gamma(n+1, \theta_n)$ is sampled independently from Z . Conditioned on any value of Z , the following inclusion holds

$$[Z + \mathbf{t}, \alpha Z + \mathbf{t}] \subseteq [rZ + \mathbf{t}, \gamma_n \alpha rZ + \mathbf{t}]$$

as long as $r \in [1, \gamma_n^{-1}] = [1, 1 + \frac{2\sqrt{2}}{\sqrt{n+1} - \sqrt{2}}]$. By Lemma B.2, we get that

$$\begin{aligned} &\mathbb{E}_X[(\mathcal{L} + \partial\mathcal{V}) \cap [X + \mathbf{t}, \gamma_n \alpha X + \mathbf{t}]] \\ &= \mathbb{E}_Z[\mathbb{E}_r[(\mathcal{L} + \partial\mathcal{V}) \cap [rZ + \mathbf{t}, \gamma_n \alpha rZ + \mathbf{t}]]] \\ &\geq \mathbb{E}_Z[(\mathcal{L} + \partial\mathcal{V}) \cap [Z + \mathbf{t}, \alpha Z + \mathbf{t}]] \Pr[r \in [1, \gamma_n^{-1}]] \\ &\geq \frac{1}{2} \mathbb{E}_Z[(\mathcal{L} + \partial\mathcal{V}) \cap [Z + \mathbf{t}, \alpha Z + \mathbf{t}]], \end{aligned}$$

as needed.

B.2.2 Bounding the Number of Crossing for Laplace Perturbations In this section, we bound the number of crossings induced by Laplace perturbations. The expression for the infinitesimal crossing probabilities is given in Lemma B.4, the bound on the surface area integral over $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ to $\mathbb{E}[\|X\|_{\mathcal{V}}]$ is given in Lemma B.6, and the full phase C Laplace crossing bound is given in Theorem B.1.

For $\mathbf{t} \in \mathbb{R}^n$, and $s > 0$, the set $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ is a shifted and scaled version of the tiling boundary $\mathcal{L} + \partial\mathcal{V}$. For $\mathbf{y} \in \mathcal{L}$, and $\mathbf{v} \in \text{VR}$, we will call $s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})$ a facet of $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$.

DEFINITION B.1. (TILING BOUNDARY NORMALS) We define the function $\eta : (\mathcal{L} - \mathbf{t} + \partial\mathcal{V}) \rightarrow S^{n-1}$ as follows. For $\mathbf{x} \in (\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$, choose the lexicographically minimal $\mathbf{v} \in \text{VR}$ such that $\exists \mathbf{y} \in \mathcal{L}$ satisfying $\mathbf{x} \in (\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})$. Finally, define $\eta(\mathbf{x}) = \mathbf{v}/\|\mathbf{v}\|_2$.

Note that for $\mathbf{x} \in s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$, $\eta(\mathbf{x}/s)$ is a unit normal to a facet of $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ containing \mathbf{x} . Furthermore, the subset of points in $s(\mathcal{L} - \mathbf{t} + F_{\mathbf{v}})$ having more than one containing facet has $n-1$ dimensional measure 0, and hence can be ignored from the perspective of integration over $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$.

The following lemma gives the expression for the infinitesimal crossing probabilities.

LEMMA B.4. For $\alpha \in (0, 1]$, and $X \sim \text{Laplace}(\mathcal{V}, \theta)$, we have that

$$\begin{aligned} &\mathbb{E}[(\mathcal{L} + \partial\mathcal{V}) \cap [X + \mathbf{t}, \alpha X + \mathbf{t}]] \\ &= \int_1^{1/\alpha} \int_{s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})} |\langle \eta(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) \text{dvol}_{n-1}(\mathbf{x}) \text{d}s. \end{aligned}$$

Proof. Firstly, shifting by $-\mathbf{t}$ on both sides, we get that $\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [X + \mathbf{t}, \alpha X + \mathbf{t}]|] = \mathbb{E}[|(\mathcal{L} - \mathbf{t} + \partial\mathcal{V}) \cap [X, \alpha X]|]$.

From here, we first decompose the expected number of intersections by summing over all facets. This yields

$$(2.7) \quad \begin{aligned} & \mathbb{E}[|(\mathcal{L} - \mathbf{t} + \partial\mathcal{V}) \cap [X, \alpha X]|] \\ &= \frac{1}{2} \sum_{\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}} \mathbb{E}[|(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}}) \cap [X, \alpha X]|]. \end{aligned}$$

The factor $1/2$ above accounts for the fact that we count each facet twice, i.e. $\mathbf{y} - \mathbf{t} + F_{\mathbf{v}}$ and $(\mathbf{y} + \mathbf{v}) - \mathbf{t} + F_{-\mathbf{v}}$. Secondly, note that the intersections we count more than twice in the above decomposition correspond to a countable number of lines passing through at most $n - 2$ dimensional faces, and hence have n dimensional Lebesgue measure 0. The equality in Equation (2.7) thus follows.

If we restrict to one facet $\mathbf{y} - \mathbf{t} + F_{\mathbf{v}}$, for some $\mathbf{y} \in \mathcal{L}$, $\mathbf{v} \in \text{VR}$, we note that the line segment $[X, \alpha X]$ crosses the facet $\mathbf{y} - \mathbf{t} + F_{\mathbf{v}}$ at most once with probability 1. Hence, we get that

$$(2.8) \quad \begin{aligned} & \mathbb{E}[|(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}}) \cap [X, \alpha X]|] \\ &= \Pr[(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}}) \cap [X, \alpha X] \neq \emptyset] \\ &= \Pr[X \in \cup_{s \in [1, \frac{1}{\alpha}]} s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})] \\ &= \int_{\cup_{s \in [1, \frac{1}{\alpha}]} s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})} f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\mathbf{x}. \end{aligned}$$

Let $r = \langle \mathbf{v} / \|\mathbf{v}\|_2, \mathbf{y} - \mathbf{t} + F_{\mathbf{v}} \rangle$, noting the inner product with \mathbf{v} (and hence $\hat{\mathbf{v}}$) is constant over $F_{\mathbf{v}}$. By possibly switching \mathbf{v} to $-\mathbf{v}$ and \mathbf{y} to $\mathbf{y} + \mathbf{v}$ (which maintains the facet), we may assume that $r \geq 0$. Notice that by construction, for any \mathbf{x} in the (relative) interior of $s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})$, we get that $r = |\langle \eta(\mathbf{x}/s), \mathbf{x}/s \rangle|$, since then there is a unique facet of $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ containing \mathbf{x} . Integrating first in the in the direction \mathbf{v} , we get that

$$(2.9) \quad \begin{aligned} & \int_{\cup_{s \in [1, \frac{1}{\alpha}]} s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})} f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\mathbf{x} \\ &= \int_r^{r/\alpha} \int_{(s/r)(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})} f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}) ds \\ &= \int_1^{1/\alpha} \int_{s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})} r f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}) ds \\ &= \int_1^{1/\alpha} \int_{s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})} |\langle \eta(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}) ds. \end{aligned}$$

Note that we use the $n - 1$ dimensional Lebesgue measure to integrate over $s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})$ since it is embedded

in \mathbb{R}^n . If $r = 0$, note that the set $\cup_{s \in [1, \frac{1}{\alpha}]} s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})$ is $n - 1$ dimensional and hence has measure (and probability) 0. This is still satisfied by the last expression in (2.9), and hence the identity is still valid in this case.

Putting everything together, combining Equation (2.7),(2.9), we get that

$$\begin{aligned} & \mathbb{E}[|(\mathcal{L} - \mathbf{t} + \partial\mathcal{V}) \cap [X, \alpha X]|] \\ &= \frac{1}{2} \sum_{\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}} \int_1^{1/\alpha} \int_{s(\mathbf{y} - \mathbf{t} + F_{\mathbf{v}})} |\langle \eta(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) \\ & \quad d\text{vol}_{n-1}(\mathbf{x}) ds \\ &= \int_1^{1/\alpha} \int_{s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})} |\langle \eta(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}) ds, \end{aligned}$$

as needed.

The lower bound given in the following Lemma will be needed in the proof of Lemma B.6.

LEMMA B.5. For $a, b, c, d \in \mathbb{R}$, $c \leq d$, we have that

$$\int_c^d |a + bh| dh \geq (\sqrt{2} - 1)(d - c) \max\{|(a + bc)|, |a + bd|\}.$$

Proof. Firstly, we note that

$$\begin{aligned} & \int_c^d |a + bh| dh \\ &= (d - c) \int_0^1 |a + b(c + (d - c)h)| dh \\ &= (d - c) \int_0^1 |(a + bc) + b(d - c)h| dh, \end{aligned}$$

hence it suffices to prove the inequality when $c = 0, d = 1$. After this reduction, by possibly applying the change of variables $h \leftarrow 1 - h$, we may assume that $|a| \geq |a + b|$. Next, by changing the signs of a, b , we may assume that $a \geq 0$. Hence, it remains to prove the inequality

$$(2.10) \quad \int_0^1 |a + bh| dh \geq a(\sqrt{2} - 1)$$

under the assumption that $a \geq |a + b|$, or equivalently $a \geq 0$ and $-2a \leq b \leq 0$. Notice that if $a = 0$ or $b = 0$, the above inequality is trivially true. If $a, b \neq 0$, then dividing inequality 2.10 by a , we reduce to the case where $a = 1, -2 \leq b < 0$. Letting $\alpha = -1/b$, we have that $\alpha \in [1/2, \infty)$. From here, we get that

$$\int_0^1 |1 + hb| dh = \int_0^1 |1 - h/\alpha| dh = (1/2)(\alpha + (1 - \alpha)^2/\alpha)$$

The derivative of the above expression is $1 - \frac{1}{2\alpha^2}$. The expression is thus minimized for $\alpha = \frac{1}{\sqrt{2}} > 1/2$, and the result follows by plugging in this value.

We now prove the bound on the surface integral in terms of the expectation $\mathbb{E}[\|X\|_{\mathcal{V}}]$.

LEMMA B.6. For $s \geq 1$ and $X \sim \text{Laplace}(\mathcal{V}, \theta)$, we have that

$$\int_{s(\mathcal{L}-\mathbf{t}+\partial\mathcal{V})} |\langle \eta(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) \text{dvol}_{n-1}(\mathbf{x}) \leq c \max \left\{ \frac{n}{s^2}, \frac{1}{\theta s} \right\} \mathbb{E}[\|X\|_{\mathcal{V}}] = c \max \left\{ \frac{n^2\theta}{s^2}, \frac{n}{s} \right\},$$

for $c = \frac{e^2}{2(\sqrt{2}-1)} \leq 9$.

Proof. We first prove the equality on the right hand side. We remember that X is identically distributed to rZ where $r \sim \Gamma(n+1, \theta)$ and $Z \sim \text{Uniform}(\mathcal{V})$. From here, we have that

$$\begin{aligned} \mathbb{E}[\|X\|_{\mathcal{V}}] &= \mathbb{E}[\|rZ\|_{\mathcal{V}}] = \mathbb{E}[r] \mathbb{E}[\|Z\|_{\mathcal{V}}] \\ &= (n+1)\theta \int_0^1 \Pr[\|Z\|_{\mathcal{V}} \geq s] \text{d}s \\ &= (n+1)\theta \int_0^1 (1-s^n) \text{d}s = (n+1)\theta \left(\frac{n}{n+1} \right) \\ &= n\theta, \end{aligned}$$

as needed.

We now prove the first inequality. To prove the bound, we write the integral expressing $\mathbb{E}[\|X\|_{\mathcal{V}}]$ over the cells of $s(\mathcal{L}-\mathbf{t}+\partial\mathcal{V})$, and compare the integral over each cell to the corresponding boundary integral. To begin

$$\begin{aligned} \mathbb{E}[\|X\|_{\mathcal{V}}] &= \int_{\mathbb{R}^n} \|\mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(\mathbf{x}) \text{d}\mathbf{x} \\ &= \sum_{\mathbf{y} \in \mathcal{L}} \int_{s(\mathbf{y}-\mathbf{t})+s\mathcal{V}} \|\mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(\mathbf{x}) \text{d}\mathbf{x} \\ &= \sum_{\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}} \int_{s(\mathbf{y}-\mathbf{t})+\text{conv}(\mathbf{0}, sF_{\mathbf{v}})} \|\mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(\mathbf{x}) \text{d}\mathbf{x}. \end{aligned} \tag{2.11}$$

Fix $\mathbf{y} \in \mathcal{L}$ and $\mathbf{v} \in \mathcal{V}$ in the above sum. Noting that $\langle \mathbf{v}/\|\mathbf{v}\|_2, sF_{\mathbf{v}} \rangle = s\|\mathbf{v}\|_2/2$ by construction, and integrating first in the direction \mathbf{v} , we get that

$$\begin{aligned} \int_{s(\mathbf{y}-\mathbf{t})+\text{conv}(\mathbf{0}, sF_{\mathbf{v}})} \|\mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(\mathbf{x}) \text{d}\mathbf{x} &= \int_0^{s\|\mathbf{v}\|_2/2} \int_{\frac{2h}{s\|\mathbf{v}\|_2}(sF_{\mathbf{v}})} \|s(\mathbf{y}-\mathbf{t}) + \mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(s(\mathbf{y}-\mathbf{t}) + \mathbf{x}) \text{dvol}_{n-1}(\mathbf{x}) \text{d}h. \end{aligned} \tag{2.12}$$

In the above, we use the $n-1$ dimensional Lebesgue measure to integrate over $\frac{2h}{s\|\mathbf{v}\|_2}F_{\mathbf{v}}$ since it is embedded in \mathbb{R}^n (we also do this for ease of notation). Setting $\beta = \frac{2h}{s\|\mathbf{v}\|_2}$, note that $\beta \in [0, 1]$. In equation (2.12), β represents the convex combination between $\mathbf{0}$ and $sF_{\mathbf{v}}$, that is $\text{conv}(\mathbf{0}, F_{\mathbf{v}}) = \bigcup_{\beta \in [0,1]} \beta sF_{\mathbf{v}}$. Performing a change of variables, Equation (2.12) simplifies to

$$\begin{aligned} & \int_0^1 \int_{h(sF_{\mathbf{v}})} (s\|\mathbf{v}\|_2/2) \|s(\mathbf{y}-\mathbf{t}) + \mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(s(\mathbf{y}-\mathbf{t}) + \mathbf{x}) \text{dvol}_{n-1}(\mathbf{x}) \text{d}h \\ &= \int_{sF_{\mathbf{v}}} \int_0^1 (s\|\mathbf{v}\|_2/2) \|s(\mathbf{y}-\mathbf{t}) + h\mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(s(\mathbf{y}-\mathbf{t}) + h\mathbf{x}) h^{n-1} \text{d}h \text{dvol}_{n-1}(\mathbf{x}) \\ &= \int_{sF_{\mathbf{v}}} \int_0^1 (s\|\mathbf{v}\|_2/2) \|s(\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(s(\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}) (1-h)^{n-1} \text{d}h \text{dvol}_{n-1}(\mathbf{x}). \end{aligned} \tag{2.13}$$

From here we note that

$$\begin{aligned} & (s\|\mathbf{v}\|_2/2) \|s(\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}\|_{\mathcal{V}} \\ &= (s\|\mathbf{v}\|_2/2) \max_{\mathbf{w} \in \text{VR}} \left| \frac{2\langle \mathbf{w}, s(\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x} \rangle}{\langle \mathbf{w}, \mathbf{w} \rangle} \right| \\ &\geq (s\|\mathbf{v}\|_2/2) \left| \frac{2\langle \mathbf{v}, s(\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} \right| \\ &= s^2 |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}/s \rangle|. \end{aligned} \tag{2.14}$$

From inequality (2.14), we have that the expression in equation (2.13) is greater than or equal to

$$\begin{aligned} & \int_{sF_{\mathbf{v}}} s^2 \int_0^1 |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(s(\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}) (1-h)^{n-1} \text{d}h \text{dvol}_{n-1}(\mathbf{x}). \end{aligned} \tag{2.15}$$

To compare to the surface integral, we now lower bound the inner integral.

CLAIM B.1. For $\|\mathbf{x}\|_{\mathcal{V}} \leq s$, we have that

$$\begin{aligned} & \int_0^1 |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(s(\mathbf{y}-\mathbf{t}) + (1-h)\mathbf{x}) (1-h)^{n-1} \text{d}h \\ &\geq e^{-2}(\sqrt{2}-1) \min \left\{ \frac{1}{n}, \frac{\theta}{s} \right\} |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y}-\mathbf{t}) + \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(s(\mathbf{y}-\mathbf{t}) + \mathbf{x}) \end{aligned}$$

Proof. Note that for $0 \leq h \leq \min\{\frac{1}{n}, \frac{\theta}{s}\}$, we have that

$$\begin{aligned} & f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + (1 - h)\mathbf{x})(1 - h)^{n-1} \\ & \geq f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + \mathbf{x})e^{-\|h\mathbf{x}\|_{\mathcal{V}}/\theta}(1 - h)^{n-1} \\ & \geq f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + \mathbf{x})e^{-1}(1 - 1/n)^{n-1} \\ & \geq e^{-2}f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + \mathbf{x}). \end{aligned}$$

Hence, using the above and Lemma B.5, we have that

$$\begin{aligned} & \int_0^1 |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y} - \mathbf{t}) + (1 - h)\mathbf{x}/s \rangle| \\ & \quad f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + (1 - h)\mathbf{x})(1 - h)^{n-1} dh \\ & \geq e^{-2}f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + \mathbf{x}) \\ & \quad \int_0^{\min\{\frac{1}{n}, \frac{\theta}{s}\}} |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y} - \mathbf{t}) + (1 - h)\mathbf{x}/s \rangle| dh \\ & \geq e^{-2}(\sqrt{2} - 1) \min\left\{\frac{1}{n}, \frac{\theta}{s}\right\} \\ & \quad |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y} - \mathbf{t}) + \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + \mathbf{x}), \end{aligned}$$

as needed.

Given Claim B.1, we get that expression (2.15) is greater than or equal to

$$\begin{aligned} & \int_{sF_{\mathbf{v}}} s^2 e^{-2}(\sqrt{2} - 1) \min\left\{\frac{1}{n}, \frac{\theta}{s}\right\} \\ & |\langle \mathbf{v}/\|\mathbf{v}\|_2, (\mathbf{y} - \mathbf{t}) + \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(s(\mathbf{y} - \mathbf{t}) + \mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}) \\ & = e^{-2}(\sqrt{2} - 1) \min\left\{\frac{s^2}{n}, s\theta\right\} \int_{s(\mathbf{y}-\mathbf{t}+F_{\mathbf{v}})} \\ & \quad |\langle \eta_s(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}). \end{aligned}$$

Putting everything together, combining the above with equation (2.11), we get that

$$\begin{aligned} & \mathbb{E}[\|X\|_{\mathcal{V}}] \\ & = \sum_{\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}} \int_{s(\mathbf{y}-\mathbf{t})+\text{conv}(\mathbf{0}, sF_{\mathbf{v}})} \|\mathbf{x}\|_{\mathcal{V}} f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\mathbf{x} \\ & \geq e^{-2}(\sqrt{2} - 1) \min\left\{\frac{s^2}{n}, s\theta\right\} \sum_{\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}} \int_{s(\mathbf{y}-\mathbf{t}+F_{\mathbf{v}})} \\ & \quad |\langle \eta_s(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}) \\ & = 2e^{-2}(\sqrt{2} - 1) \min\left\{\frac{s^2}{n}, s\theta\right\} \int_{s(\mathcal{L}-\mathbf{t}+\partial\mathcal{V})} \\ & \quad |\langle \eta_s(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}), \end{aligned} \tag{A}$$

where the last equality follows since each facet in $s(\mathcal{L} - \mathbf{t} + \partial\mathcal{V})$ is counted twice. The lemma thus follows.

The following gives the full phase C bound for Laplace perturbations.

THEOREM B.1. For $\alpha \in (0, 1]$ and $X \sim \text{Laplace}(\mathcal{V}, \theta)$, we have that

$$\begin{aligned} & \mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [X + \mathbf{t}, \alpha X + \mathbf{t}]|] \\ & \leq cn \left(\frac{n\theta}{2} \left(1 - \frac{1}{s^*}\right) + \ln\left(\frac{1}{\alpha s^*}\right) \right) \end{aligned}$$

where $c = \frac{e^2}{2(\sqrt{2}-1)} \leq 9$ and $s^* = \max\{1, n\theta\}$.

Proof. Using Lemmas B.4 and B.6, we have that

$$\begin{aligned} & \mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [X + \mathbf{t}, \alpha X + \mathbf{t}]|] \\ & = \int_1^{1/\alpha} \int_{s(\mathcal{L}-\mathbf{t}+\partial\mathcal{V})} |\langle \eta(\mathbf{x}/s), \mathbf{x}/s \rangle| f_{\mathcal{V}}^{\theta}(\mathbf{x}) d\text{vol}_{n-1}(\mathbf{x}) ds \\ & \leq c \int_1^{1/\alpha} \max\left\{\frac{n^2\theta}{s^2}, \frac{n}{s}\right\} ds = c \int_1^{s^*} \frac{n^2\theta}{s^2} ds + c \int_{s^*}^{1/\alpha} \frac{n}{s} ds \\ & = cn^2\theta \left(1/2 - \frac{1}{s^*}\right) + cn \ln\left(\frac{1}{\alpha s^*}\right) \\ & = cn \left(\frac{n\theta}{2} \left(1 - \frac{1}{s^*}\right) + \ln\left(\frac{1}{\alpha s^*}\right)\right), \end{aligned}$$

as needed.

B.2.3 Proof of Theorem 3.2 (Phase C crossing bound)

Proof. We recall that $\theta_n = \frac{1}{(n+1)-\sqrt{2(n+1)}}$ and $\gamma_n = \left(1 + \frac{2\sqrt{2}}{\sqrt{n+1}-\sqrt{2}}\right)^{-1}$. Note that $n\theta_n > 1$.

By Lemma B.3 and Theorem B.1, for $X \sim \text{Laplace}(\mathcal{V}, \theta_n)$, we have that

$$\begin{aligned} & \mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [Z + \mathbf{t}, \alpha Z + \mathbf{t}]|] \\ & \leq 2\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [X + \mathbf{t}, \gamma_n \alpha X + \mathbf{t}]|] \\ & \leq 2cn \left(\frac{n\theta_n}{2} \left(1 - \frac{1}{n\theta_n}\right) + \ln\left(\frac{1}{\alpha \gamma_n n\theta_n}\right) \right) \\ & \leq \frac{e^2}{\sqrt{2}-1} n(2 + \ln(4/\alpha)), \text{ for } n \geq 2, \end{aligned}$$

as needed.

C Missings proofs from Section 3

Proof. [Proof of Lemma 3.1 (Randomized Straight Line Complexity)] We recall the three phases of the randomized straight line algorithm:

- (A) Move from \mathbf{x} to $\mathbf{x} + Z$.
- (B) Follow the sequence of Voronoi cells from $\mathbf{x} + Z$ to $\mathbf{y} + Z$.
- (C) Follow the sequence of Voronoi cells from $\mathbf{y} + Z$ to \mathbf{y} .

Characterizing the Path Length: We will show that with probability 1, the length of the path on \mathcal{G} induced by the three phases is

$$|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{t} + Z]| + |(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{t} + Z, \mathbf{t}]|.$$

Firstly, note that since $Z \in \mathcal{V}$ and $\mathbf{x} \in \mathcal{L}$, $\mathbf{x} + Z$ and \mathbf{x} lie in the same Voronoi cell $\mathbf{x} + Z$, and hence phase A corresponds to the trivial path \mathbf{x} . Hence, we need only worry about the number of edges induced by phases B and C.

The following claim will establish the structure of a generic intersection pattern with the tiling boundary, which will be necessary for establishing the basic properties of the path.

CLAIM C.1. *With probability 1, the path $[\mathbf{x} + Z, \mathbf{t} + Z] \cup [\mathbf{t} + Z, \mathbf{t}]$ only intersects $\mathcal{L} + \partial\mathcal{V}$ in the relative interior of its facets. Furthermore, with probability 1, the intersection consists of isolated points, and $\mathbf{x} + Z, \mathbf{t} + Z \notin \mathcal{L} + \partial\mathcal{V}$.*

Proof. We prove the first part. Let C_1, \dots, C_k denote the $n - 2$ dimensional faces of \mathcal{V} . Note that the probability of not hitting $\mathcal{L} + \partial\mathcal{V}$ in the relative interior of its facets, can be expressed as

$$\begin{aligned} (3.16) \quad & \Pr[(\mathbf{x} + Z, \mathbf{t} + Z] \cup [\mathbf{t} + Z, \mathbf{t})) \cap (\cup_{i \in [k]} \mathcal{L} + C_i) \neq \emptyset] \\ & \leq \sum_{\mathbf{y} \in \mathcal{L}, i \in [k]} \Pr[[\mathbf{x} + Z, \mathbf{t} + Z] \cap (\mathbf{y} + C_i) \neq \emptyset] + \\ & \Pr[[\mathbf{t} + Z, \mathbf{t}] \cap (\mathbf{y} + C_i) \neq \emptyset]. \end{aligned}$$

Here the last inequality is valid since \mathcal{L} is countable. Analyzing each term separately, we see that

$$\Pr[[\mathbf{x} + Z, \mathbf{t} + Z] \cap (\mathbf{y} + C_i) \neq \emptyset] = \Pr[Z \in \mathbf{y} + C_i - [\mathbf{x}, \mathbf{t}]] = 0.$$

To justify the last equality, note that since C_i is $n - 2$ dimensional, $\mathbf{y} + C_i - [\mathbf{x}, \mathbf{t}]$ is at most $n - 1$ dimensional (since the line segment can only add 1 dimension). Therefore $\mathbf{y} + C_i - [\mathbf{x}, \mathbf{t}]$ has n dimensional Lebesgue measure 0, and in particular probability 0 with respect to $\text{Uniform}(\mathcal{V})$. Next, we have that

$$\Pr[[\mathbf{t} + Z, \mathbf{t}] \cap (\mathbf{y} + C_i) \neq \emptyset] = \Pr[Z \in \cup_{s > 1} s(\mathbf{y} + C_i - \mathbf{t})] = 0,$$

where the last equality follows since $\cup_{s > 1} s(\mathbf{y} + C_i - \mathbf{t})$ is at most $n - 1$ dimensional. Hence the probability in (3.16) is 0, as needed.

We now prove the second part. Note that if the path $[\mathbf{x} + Z, \mathbf{t} + Z] \cup [\mathbf{t} + Z, \mathbf{t}]$ does not intersect $\mathcal{L} + \partial\mathcal{V}$ in isolated points (i.e. the intersection contains a non-trivial interval), then either $[\mathbf{x} + Z, \mathbf{t} + Z]$ or $[\mathbf{t} + Z, \mathbf{t}]$

must intersect some facet of $\mathcal{L} + \partial\mathcal{V}$ in a least 2 distinct points.

Let $F_{\mathbf{v}}$ be the facet of \mathcal{V} induced by $\mathbf{v} \in \text{VR}$. If $[\mathbf{t} + Z, \mathbf{t}]$ intersects $\mathbf{y} + F_{\mathbf{v}}$, for some $\mathbf{y} \in \mathcal{L}$, in two distinct points then we must have that $\langle \mathbf{v}, Z \rangle = 0$. Since $\Pr[\cup_{\mathbf{v} \in \text{VR}} \{\langle \mathbf{v}, Z \rangle = 0\}] = 0$, this event happens with probability 0. Next, note that $[\mathbf{x} + Z, \mathbf{t} + Z]$ intersects $\mathbf{y} + F_{\mathbf{v}}$ in two distinct points, if and only if $\langle \mathbf{v}, \mathbf{x} - \mathbf{t} \rangle = 0$ and $\langle \mathbf{v}, Z \rangle = \langle \mathbf{v}, \mathbf{y} + \mathbf{v}/2 \rangle$. But then, the probability of this happening for any facet can be bounded by

$$\Pr[\cup_{\mathbf{y} \in \mathcal{L}, \mathbf{v} \in \text{VR}} \{\langle \mathbf{v}, Z \rangle = \langle \mathbf{v}, \mathbf{y} + \mathbf{v}/2 \rangle\}] = 0$$

since $\mathcal{L} \times \text{VR}$ is countable.

For the last part, note that since $\mathcal{L} + \partial\mathcal{V}$ is the union of $n - 1$ dimensional pieces, $\Pr[\mathbf{x} + Z \in \mathcal{L} + \partial\mathcal{V}] + \Pr[\mathbf{t} + Z \in \mathcal{L} + \partial\mathcal{V}] = 0$.

The claim thus follows.

Conditioning on the intersection structure given in claim C.1, we now describe the associated path on \mathcal{G} . Let $\mathbf{p}_1, \dots, \mathbf{p}_k$ denote the points in $([\mathbf{x} + Z, \mathbf{t} + Z] \cup [\mathbf{t} + Z, \mathbf{t}]) \cap (\mathcal{L} + \partial\mathcal{V})$ ordered in order of appearance on the path $[\mathbf{x} + Z, \mathbf{t} + Z] \cup [\mathbf{t} + Z, \mathbf{t}]$ from left to right. Letting $\mathbf{p}_{k+1} = \mathbf{t}$, let $\mathbf{y}_i \in \mathcal{L}$, $1 \leq i \leq k$, denote the center of the unique Voronoi cell in $\mathcal{L} + \mathcal{V}$ containing the interval $[\mathbf{p}_i, \mathbf{p}_{i+1}]$. Note that the existence of \mathbf{y}_i is guaranteed since the Voronoi cells in the tiling $\mathcal{L} + \mathcal{V}$ are interior disjoint, and the open segment $(\mathbf{p}_i, \mathbf{p}_{i+1})$ lies in the interior of some Voronoi cell by convexity of the cells.

Letting $\mathbf{y}_0 = \mathbf{x}$, we now claim that $\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_k$ form a valid path in \mathcal{G} . To begin, we first establish that $\mathbf{y}_i \neq \mathbf{y}_{i+1}$, $0 \leq i \leq k$. Firstly, since $\mathbf{x} + Z \notin \mathcal{L} + \partial\mathcal{V}$, we have that Z is in the interior of \mathcal{V} , and hence the ray starting at Z in the direction of \mathbf{p}_1 exits $\mathbf{x} + \mathcal{V}$ at \mathbf{p}_1 and never returns (by convexity of \mathcal{V}). Furthermore, since $\mathbf{p}_1 \neq \mathbf{t} + Z$, the Voronoi cell $\mathbf{y}_1 + \mathcal{V}$ must contain a non-trivial interval on this ray starting at \mathbf{p}_1 , i.e. $[\mathbf{p}_1, \mathbf{p}_2]$, and hence $\mathbf{y}_1 \neq \mathbf{x}$. Indeed, for the remaining cases, the argument follows in the same way as long as the Voronoi cell $\mathbf{y}_{i+1} + \mathcal{V}$ contains a non-trivial interval of the ray exiting $\mathbf{y}_i + \mathcal{V}$. Note that this is guaranteed by the assumption that $\mathbf{t} + Z \notin \partial\mathcal{V}$ and by the fact that none of the \mathbf{p}_i s equals \mathbf{t} . Hence $\mathbf{y}_i \neq \mathbf{y}_{i+1}$, $0 \leq i \leq k$, as needed.

Next, note that each \mathbf{p}_i , $i \in [k]$, belongs to the relative interior of some facet of $\mathcal{L} + \partial\mathcal{V}$. Furthermore, by construction $\mathbf{p}_i \in \mathbf{y}_{i-1} + \partial\mathcal{V}$ and $\mathbf{p}_i \in \mathbf{y}_i + \partial\mathcal{V}$. Since the relative interior of facets of $\mathcal{L} + \partial\mathcal{V}$ touch exactly two adjacent Voronoi cells, and since $\mathbf{y}_{i-1} \neq \mathbf{y}_i$, we must have that $\mathbf{p}_i \in \mathbf{y}_{i-1} + F_{\mathbf{v}}$, where $\mathbf{v} = \mathbf{y}_i - \mathbf{y}_{i-1} \in \text{VR}$. Hence the path $\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_k$ is valid in \mathcal{G} as claimed.

From here, note that the length of is indeed $k = |([\mathbf{x} + Z, \mathbf{t} + Z] \cup [\mathbf{t} + Z, \mathbf{t}]) \cap (\mathcal{L} + \partial\mathcal{V})|$. Since this holds

with probability 1, we get that the expected path length is

$$\mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{x} + Z, \mathbf{t} + Z]|] + \mathbb{E}[|(\mathcal{L} + \partial\mathcal{V}) \cap [\mathbf{t} + Z, \mathbf{t}]|] .$$

as needed.

Computing the Path: We now explain how to compute each edge of the path using $O(n|\text{VR}|)$ arithmetic operations, conditioning on the conclusions of Claim C.1.

In constructing the path, we will in fact compute the intersection points $\mathbf{p}_1, \dots, \mathbf{p}_k$ as above, and the lattice points $\mathbf{y}_1, \dots, \mathbf{y}_k$. As one would expect, this computation is broken up in phase B and C, corresponding to computing the intersection / lattice points for $[\mathbf{x} + Z, \mathbf{t} + Z]$ in phase B, followed by the intersection / lattice points from $[\mathbf{t} + Z, \mathbf{t}]$ in Phase C.

For each phase, we will use a generic line following procedure that given vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, and a starting lattice point $\mathbf{z} \in \mathcal{L}$, such that $\mathbf{a} \in \mathbf{z} + \mathcal{V}$, follows the path of Voronoi cells along the line segment $[\mathbf{a}, \mathbf{b}]$, and outputs a lattice vector $\mathbf{w} \in \mathcal{L}$ satisfying $\mathbf{b} \in \mathbf{w} + \mathcal{V}$. To implement phase B, we initialize the procedure with $\mathbf{x} + Z, \mathbf{t} + Z$ and starting point \mathbf{x} . For phase C, we give it $\mathbf{t} + Z, \mathbf{t}$ and the output of phase B as the starting point.

We describe the line following procedure. Let $\ell(\alpha) = (1 - \alpha)\mathbf{a} + \alpha\mathbf{b}$, for $\alpha \in [0, 1]$, i.e. the parametrization of $[\mathbf{a} + Z, \mathbf{b} + Z]$ as a function of time. The procedure will have a variable for α , which will be set at its bounds at the beginning and end of the procedure, starting at 0 ending at ≥ 1 , and in intermediate steps will correspond to an intersection point. We will also have a variable $\mathbf{w} \in \mathcal{L}$, corresponding to the current Voronoi cell center. We will maintain the invariant that $\ell(\alpha) \in \mathbf{w} + \mathcal{V}$, and furthermore that $\ell(\alpha) \in \mathbf{w} + \partial\mathcal{V}$ for $\alpha \in (0, 1)$.

The line following algorithm is as follows:

Data: $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n, \mathbf{z} \in \mathcal{L}, \mathbf{a} \in \mathbf{z} + \mathcal{V}$
Result: $\mathbf{w} \in \mathcal{L}$ such that $\mathbf{b} \in \mathbf{w} + \mathcal{V}$
 $\mathbf{w} \leftarrow \mathbf{z}, \mathbf{e} \leftarrow 0, \alpha \leftarrow 0$
 $\text{VR}' \leftarrow \{\mathbf{v} \in \text{VR} : \langle \mathbf{v}, \mathbf{b} - \mathbf{a} \rangle > 0\}$
repeat
 $\mathbf{w} \leftarrow \mathbf{w} + \mathbf{e}$
 $\mathbf{e} \leftarrow \arg \min_{\mathbf{v} \in \text{VR}'} \frac{\langle \mathbf{v}, \mathbf{v}/2 + \mathbf{w} - \mathbf{a} \rangle}{\langle \mathbf{b} - \mathbf{a}, \mathbf{v} \rangle}$
 $\alpha \leftarrow \frac{\langle \mathbf{e}, \mathbf{e}/2 + \mathbf{w} - \mathbf{a} \rangle}{\langle \mathbf{b} - \mathbf{a}, \mathbf{e} \rangle}$
until $\alpha \geq 1$
return \mathbf{w}

Described in words, each loop iteration does the following: given the current Voronoi cell $\mathbf{w} + \mathcal{V}$, and the entering intersection point $\ell(\alpha)$ of the line segment

$[\mathbf{a}, \mathbf{b}]$ with respect to $\mathbf{w} + \mathcal{V}$, we first compute the exiting intersection point $\ell(\alpha')$, $\alpha' > \alpha$, and the exiting facet $\mathbf{w} + F_{\mathbf{e}}$. If $\alpha' \geq 1$, we know that $\mathbf{b} \in [\ell(\alpha), \ell(\alpha')] \subseteq \mathbf{w} + \mathcal{V}$, and hence we may return \mathbf{w} . Otherwise, we move to the center of the Voronoi cell sharing the facet $\mathbf{w} + F_{\mathbf{e}}$ opposite \mathbf{w} .

To verify the correctness, we need only show that the line $[\mathbf{a}, \mathbf{b}]$ indeed exits $\mathbf{w} + \mathcal{V}$ through the facet $\mathbf{w} + F_{\mathbf{e}}$ at the end of each iteration. Note that by our invariant $\ell(\alpha) \in \mathbf{w} + \mathcal{V}$ at the beginning of the iteration, and hence

$$(3.17) \quad \begin{aligned} \langle \mathbf{v}, \ell(\alpha) - \mathbf{w} \rangle &\leq \frac{1}{2} \langle \mathbf{v}, \mathbf{v} \rangle, \quad \forall \mathbf{v} \in \text{VR} \quad \Leftrightarrow \\ \langle \mathbf{v}, (1 - \alpha)\mathbf{a} + \alpha\mathbf{b} - \mathbf{w} \rangle &\leq \frac{1}{2} \langle \mathbf{v}, \mathbf{v} \rangle, \quad \forall \mathbf{v} \in \text{VR} \quad \Leftrightarrow \\ \alpha \langle \mathbf{v}, \mathbf{b} - \mathbf{a} \rangle &\leq \langle \mathbf{v}, \mathbf{v}/2 - \mathbf{a} + \mathbf{w} \rangle, \quad \forall \mathbf{v} \in \text{VR} \end{aligned}$$

Since we move along the line segment $[\mathbf{a}, \mathbf{b}]$ by increasing α , i.e. going from \mathbf{a} to \mathbf{b} , note that the only constraints that can be eventually violated as we increase α are those for which $\langle \mathbf{v}, \mathbf{b} - \mathbf{a} \rangle > 0$. Hence, in finding the first violated constraint (i.e. exiting facet), we may restrict our attention to the subset of Voronoi relevant vectors $\text{VR}' = \{\mathbf{v} \in \text{VR} : \langle \mathbf{v}, \mathbf{b} - \mathbf{a} \rangle > 0\}$ as done in the algorithm.

From (3.17), we see that we do not to cross any facet $\mathbf{w} + F_{\mathbf{v}}, \mathbf{v} \in \text{VR}'$, as long as

$$\alpha \leq \frac{\langle \mathbf{v}, \mathbf{v}/2 - \mathbf{a} + \mathbf{w} \rangle}{\langle \mathbf{v}, \mathbf{b} - \mathbf{a} \rangle}, \quad \forall \mathbf{v} \in \text{VR}' .$$

Hence the first facet we violate must be induced by

$$(3.18) \quad \mathbf{e} = \arg \min_{\mathbf{v} \in \text{VR}'} \frac{\langle \mathbf{v}, \mathbf{v}/2 - \mathbf{a} + \mathbf{w} \rangle}{\langle \mathbf{v}, \mathbf{b} - \mathbf{a} \rangle} .$$

Letting $\alpha' = \frac{\langle \mathbf{e}, \mathbf{e}/2 - \mathbf{a} + \mathbf{w} \rangle}{\langle \mathbf{e}, \mathbf{b} - \mathbf{a} \rangle}$, we see that $\ell(\alpha') \in \mathbf{w} + F_{\mathbf{e}}$ is the correctly computed exiting point (corresponding to $\ell(\alpha)$ at the end of the loop iteration), and that $\mathbf{w} + F_{\mathbf{e}}$ is the exiting facet. Since the facet $\mathbf{w} + F_{\mathbf{e}}$ is shared by $(\mathbf{w} + \mathbf{e}) + \mathcal{V}$, we see that $\ell(\alpha') \in (\mathbf{w} + \mathbf{e}) + \partial\mathcal{V}$, and hence the invariant is maintained in the next iteration. The line following algorithm is thus correct.

Notice that each iteration of the line following procedure clearly requires at most $O(n|\text{VR}|)$ arithmetic operations. We note that the conclusions of Claim C.1 are only needed to ensure that each iteration of the path finding procedure can be associated with exactly one intersection point in $([\mathbf{x} + Z, \mathbf{t} + Z] \cup [\mathbf{t} + Z, \mathbf{t}]) \cap (\mathcal{L} + \partial\mathcal{V})$. In particular, it assures that the minimizer in (3.18) is unique. This concludes the proof of the Lemma.

Proof. [Proof of Lemma 3.3 (Bit length bound)] Clearly, (3.19)

$$\bar{q} \leq \left(\prod_{ij} q_{ij}^B \right) \left(\prod_i q_i^{\mathbf{t}} \right) \Rightarrow \log_2 \bar{q} \leq \sum_{ij} \log_2(q_{ij}^B) + \sum_i q_i^{\mathbf{t}}.$$

Hence $\log_2 \bar{q}$ is smaller than the sum of encoding sizes of the denominators of the entries of B and \mathbf{t} . Next, it is well known that $\mu(\mathcal{L}) \leq \frac{1}{2} \sqrt{\sum_{ij} B_{ij}^2}$ (see for example [5]). From here, we get that

$$\begin{aligned} (3.20) \quad \log_2 \mu(\mathcal{L}) &\leq \log_2 \left(\sqrt{\sum_{ij} B_{ij}^2} \right) \leq \log_2 \left(\sqrt{\sum_{ij} (p_{ij}^B)^2} \right) \\ &\leq \log_2 \left(\sqrt{\prod_{ij} ((p_{ij}^B)^2 + 1)} \right) \\ &= \sum_{ij} \log_2 \left(\sqrt{(p_{ij}^B)^2 + 1} \right) \leq \sum_{ij} \log_2 (|p_{ij}^B| + 1) \end{aligned}$$

Hence $\log_2 \mu(\mathcal{L})$ is less than the sum of encoding sizes of the numerators of the entries in B . The bound $\log_2(\bar{q}\mu(\mathcal{L})) \leq \text{enc}(B) + \text{enc}(\mathbf{t})$ now follows by adding (3.19), (3.20).

We now bound $\log_2(\mu(\mathcal{L})/\lambda_1(\mathcal{L}))$. Letting $\tilde{q} = \prod_{ij} q_{ij}^B$, note that

$$\tilde{q}\lambda_1(\mathcal{L}) = \lambda_1(\tilde{q}\mathcal{L}) \geq \lambda_1(\mathbb{Z}^n) = 1.$$

Therefore $1/\lambda_1(\mathcal{L}) \leq \tilde{q}$. Since $\log_2 \tilde{q} \leq \sum_{ij} \log_2(q_{ij}^B)$ and $\log_2(\mu(\mathcal{L})/\lambda_1(\mathcal{L})) \leq \log_2(\tilde{q}\mu(\mathcal{L}))$, combining with (3.20) we get that $\log_2(\mu(\mathcal{L})/\lambda_1(\mathcal{L})) \leq \text{enc}(B)$ as needed.