Smoothed Analysis of Probabilistic Roadmaps

Siddhartha Chaudhuri and Vladlen Koltun*

Abstract

The probabilistic roadmap algorithm is a leading heuristic for robot motion planning. It is extremely efficient in practice, yet its worst case convergence time is *unbounded* as a function of the input's combinatorial complexity. We prove a smoothed polynomial upper bound on the number of samples required to produce an accurate probabilistic roadmap, and thus on the running time of the algorithm, in an environment of simplices. This sheds light on its widespread empirical success.

1 Introduction

Smoothed analysis. It is well-documented that many geometric algorithms that are extremely efficient in practice have exceedingly poor worst-case performance guarantees. Two approaches were put forth to address this issue. The first tries to formally model various classes of inputs that arise in practice and analyze the performance of algorithms on these models [16]. For example, it was proposed that in practice geometric objects are *fat* [1, 13, 32, 39], point sets have bounded *spread* [8,10,18,19], and geometric scenes have *low density*, are *uncluttered*, *sparse*, etc. [6,14,15,33].

The second approach stems from the observation that geometric inputs often contain a small amount of random noise, such as with point clouds generated by a laser scanner [30]. It can be argued that small degrees of randomness creep into geometric inputs even if they are created by a human modeler [36]. By this reasoning, finely tuned worst-case examples have a low probability of arising and should not disproportionately skew theoretical measures of algorithm performance. This is formalized in smoothed analysis [38], which measures the maximum over inputs of the expected running time of the algorithm under slight random perturbations of those inputs. For example, let $A \in \mathbb{R}^{n\times d}$ specify a set of n points in $\mathbb{R}^d$, and let $f_X(A)$, where $f_X : \mathbb{R}^{n\times d} \mapsto \mathbb{R}$, be a measure of the performance of algorithm X on A. Then the worst-case performance of X is

$$\max_{A\in\mathbb{R}^{n\times d}} f_X(A),$$

the average-case performance of X is

$$\mathop{\mathrm{E}}_{A\sim\mathcal{D}} \left[f_X(A)\right],$$

where $\mathcal{D} : \mathbb{R}^{n\times d} \mapsto \mathbb{R}$ is a suitable distribution, and the smoothed performance of X is

$$\max_{A\in\mathbb{R}^{n\times d}} \mathop{\mathrm{E}}_{R\sim\mathcal{N}} \left[f_X(A + \|A\|R)\right],$$

where $\|A\|$ denotes the Frobenius norm of A and $\mathcal{N} = N(0, \sigma^2 I_{n\times d})$ is a Gaussian distribution in $\mathbb{R}^{n\times d}$ with mean 0 and variance σ^2. The parameter σ controls the magnitude of the random perturbation, and as it varies from 0 to ∞ the smoothed performance measure interpolates between worst-case and average-case performance.

Smoothed analysis is a new framework that has already been applied to a wide variety of problems [3, 4, 7, 11, 12, 17, 37]. Its advantage compared to the above-described explicit formulation of realistic input models lies in its generality and immediate applicability across contexts, and its reliance on only one assumption, namely the presence of some degree of randomness in the input.

Probabilistic roadmaps. The probabilistic roadmap (PRM) algorithm revolutionized robot motion planning [23, 25, 27]. It is a simple heuristic that exhibits rapid performance and has become the standard algorithm in the field [20, 21, 35]. Yet its worst-case running time is *unbounded* as a function of the input's combinatorial complexity. The basic algorithm for constructing a probabilistic roadmap is as follows:

> Sample uniformly at random a set of points, called milestones, from the *configuration space* $\mathcal{C}$ of the robot. Keep only those milestones that

*Computer Science Department, 353 Serra Mall, Stanford University, Stanford, CA 94305, USA; {sidch,vladlen}@stanford.edu. The first author was supported by a PACCAR Inc Stanford Graduate Fellowship.

lie in the *free configuration space* $\mathcal{C}_{\mathsf{free}}$. [1] Let V be the resulting point set. For every $u, v \in V$, if the straight line segment between u and v lies entirely in $\mathcal{C}_{\mathsf{free}}$, add $\{u, v\}$ to the set of edges E, initially empty. The graph $G = (V, E)$ is the probabilistic roadmap.

Given such a roadmap G, a motion between two points p, q in $\mathcal{C}_{\text{free}}$ can be constructed as follows:

Find a milestone p' (resp., q') in V that is visible from p (resp., from q). If p' and q' lie in different connected components of G, report that there is no feasible motion between p and q. Otherwise plan the motion using a path in G that connects p' and q'.

The above PRM construction and query algorithms can be efficiently implemented in very general settings. The outstanding issue is what the number of samples should be to guarantee (in expectation) that G accurately represents the connectivity of $\mathcal{C}_{\text{free}}$. Clearly, for the algorithm to be accurate there should be a milestone visible from any point in $\mathcal{C}_{\text{free}}$, and there should be a bijective correspondence between the set of connected components of G and the set of connected components of $\mathcal{C}_{\text{free}}$. Unfortunately, the number of random samples required to guarantee this can be made arbitrarily large even for very simple configuration spaces [21].

A number of theoretical analyses provide bounds for the number of samples under assumptions on the structure of $\mathcal{C}_{\text{free}}$ such as goodness [5,26], expansiveness [22], and the existence of high-clearance paths [24]. However, none of these assumptions were justified in terms of realistic motion planning problems. In practice, the number of random samples is chosen ad hoc.

Contributions. This paper initiates the use of smoothed analysis to explain the success of PRM. We model the configuration space using a set of n simplices in $\mathbb{R}^d$ whose vertices are subject to Gaussian perturbation with variance σ^2. We prove a smoothed upper bound on the required number of milestones that is polynomial in n and $\frac{1}{\sigma}$. The result extends to all γ-smooth perturbations, see below.

In order to achieve this bound we define a space decomposition called the locally orthogonal decomposition. Previously known decompositions, like the vertical decomposition [9,28] and the "castles in the air" decomposition [2] turn out to be unsuitable for our purpose. We prove that for the roadmap to accurately represent the free configuration space it is sufficient that a milestone is sampled from every cell of this decomposition. We then prove a smoothed lower bound on the volume of every decomposition cell. This leads to the desired bound on the number of milestones.

Our result is only the first step towards a convincing theoretical justification of PRM. The analysis is quite challenging already for the simple representation of the configuration space using independently perturbed simplices. In Section 4 we outline directions for its extension to more general configuration space models.

2 Bounding the Number of Milestones

Notation. If V is a vector space of d dimensions then, for $0 \le k \le d$, a *k-subspace* of V is the set of linear combinations of k linearly independent vectors lying in V. A subspace necessarily contains the origin. A *k-flat* is any translation of a k-subspace. Points are 0-flats, straight lines are 1-flats, planes are 2-flats and hyperplanes are $(d-1)$-flats. Any lower-dimensional flat can be modelled as the intersection of hyperplanes.

A hyperplane divides the vector space V into two *halfspaces*. More generally, a set of hyperplanes $\mathcal{H}$ subdivides V into a number of disjoint, open, d-dimensional *cells*. Further, assume a subset $\mathcal{U}$ of $\mathcal{H}$ intersects in a k-dimensional flat F, and let $\mathcal{U}'$ be the set of hyperplanes in $\mathcal{H}$ which intersect F but do not contain it. $\mathcal{U}'$ subdivides F into disjoint, open, k-dimensional regions called k-faces (if $\mathcal{U}'$ is empty, F is a k-face on its own). A 0-face is called a *vertex*, and a $(d-1)$-face is called a *facet*. Extending the notation, a cell is considered a d-face. The entire structure is referred to as the *arrangement* of the set of hyperplanes $\mathcal{H}$. An arrangement of hyperplanes is a *convex subdivision*, since all its faces are convex sets.

A set of hyperplanes $\mathcal{H}$ (or their arrangement) is in *general position* if every pair of hyperplanes in $\mathcal{H}$ intersect, and the intersection of any k hyperplanes of $\mathcal{H}$ $(1 \le k \le d)$ is not contained in any other member of $\mathcal{H}$. We note that the precise meaning of "general position" we adopt here defines a suitable "general case" for our problem — other authors may use different notions.

The *affine span* of a set A (denoted $\mathrm{Aff}(A)$) is the set of all linear combinations of elements of the set with coefficients (which may be negative) summing to one. It is always a k-flat for some k. The *convex hull* of A is similar to the affine span, with the additional requirement that all the coefficients be non-negative —

[1] A robot's *configuration space* is the set of physical positions it may attain (which may or may not coincide with obstacles), parametrised by its degrees of freedom (so a robot with d degrees of freedom has a d-dimensional configuration space). The robot's *free configuration space* is the subset of these positions which do *not* coincide with obstacles, i.e. are possible in real life. These terms are standard in the motion planning literature [29].

it can be shown to be the smallest convex set that contains A.

For $0 \leq k \leq d$, a *k-simplex* in V is the convex hull of $k+1$ affinely independent points in V. For example, a point is a 0-simplex, a line segment is a 1-simplex, a triangle is a 2-simplex and a tetrahedron is a 3-simplex. Each k-simplex (for $k \geq 1$) is bounded by a collection of $(k-1)$-simplices — these are called the *facets of the simplex*.

The *shortest distance* between two flats X and Y is defined to be $\min_{x \in X, y \in Y} \|x - y\|$, where $\|.\|$ denotes the vector norm. (In $\mathbb{R}^d$, we will assume the norm is Euclidean.) Two flats are said to be *ε-close* if the shortest distance between them is at most ε; similarly, they are *ε-distant* if the shortest distance between them is at least ε.

$A \oplus B$ is the *Minkowski sum* of sets A and B, i.e. it is the set of all sums of the form $a + b$, where $a \in A$ and $b \in B$.

The *d-ball* of radius r, denoted $B_d(r)$, is the set of all points at a distance of at most r from the origin ($B_d(x, r)$ implies the centre is at the point x rather than at the origin). The boundary of the d-ball, i.e. the set of all points at a distance of precisely r from the origin, is the *$(d-1)$-sphere* of radius r, written $S_{d-1}(r)$. Omitting all arguments in the above notation implies the unit ball or sphere centred at the origin is being considered.

The *volume* of a k-dimensional object will refer to its k-dimensional measure. If this object is embedded in a space of higher dimension (such as the $(k-1)$-sphere, which is usually embedded in $\mathbb{R}^k$), we may also refer to this measure as the *area* of the object. The meaning of these terms should be clear from the context, and the Vol() and Area() predicates may be used.

The volume of $B_d(r)$ will be written as $V_d(r)$. It is a standard result that

$$V_d(r) = \frac{\pi^{\frac{d}{2}} r^d}{\Gamma\left(\frac{d}{2} + 1\right)}$$

For fixed r this quantity diminishes to zero as d goes to infinity, and $V_d(1)$ is bounded by $8\pi^2/15$ for all d. Also

$$\text{Area}(S_{d-1}(r)) = \frac{d\ V_d(r)}{r}$$

The model. Let Σ be a fixed, convex, polyhedral bounding box for $\mathcal{C}_{\text{free}}$ in $\mathbb{R}^d$, where d is assumed to be constant. This is the domain from which the milestones are sampled by the PRM algorithm. Let D be the diameter and D_{in} be the inner diameter of Σ (the inner diameter of a region is the diameter of the largest ball contained completely within the region). Let $\mathcal{S}$ be a set of n $(d-1)$-simplices in Σ. These are the $\mathcal{C}$-space obstacles in our model. Thus $\mathcal{C}_{\text{free}} = \Sigma \setminus \bigcup_{s \in \mathcal{S}} s$.

A probability distribution $\mathcal{D}$ on $\mathbb{R}^d$ with density function $\mu(.)$ is said to be γ-smooth, for some $\gamma \in \mathbb{R}$, if

1. $\mu(x) \leq \gamma$ for all $x \in \mathbb{R}^d$, and
2. given any hyperplane H, a point distributed under $\mathcal{D}$ is almost surely not on H.

A symmetric d-variate Gaussian distribution with variance σ^2 (covariance matrix $\sigma^2 I_d$) is $\Theta\left(\frac{1}{\sigma^d}\right)$-smooth. We assume that each vertex of each simplex in $\mathcal{S}$ is independently perturbed according to a γ-smooth distribution within the domain.

We note that these simplices may also be thought of as boundary elements of full-dimensional polyhedral obstacles. Our upper bound on the the number of samples required to build an accurate roadmap applies verbatim, since we will discard those samples which fall in the interior of these polyhedra. However, our analysis is then not completely realistic because our perturbation model destroys the connectivity of these boundaries — an improved model and its analysis form a possible avenue of future work (see the Conclusion section).

The locally orthogonal decomposition. The locally orthogonal decomposition $\boxtimes(\mathcal{S})$ of $\mathcal{S}$ is the arrangement of the following two collections of hyperplanes:

- $\text{Aff}(s)$ for each $s \in \mathcal{S}$.
- The hyperplane orthogonal to s that is spanned by f, for each $s \in \mathcal{S}$ and each facet f of s.

Hyperplanes of the second type are called *walls*. A facet of $\boxtimes(\mathcal{S})$ is *bound* if it is contained in some $s \in \mathcal{S}$, otherwise it is *free*. In the following, the decomposition is assumed to be restricted to Σ. The second property of γ-smooth distributions ensures that under our perturbation model, $\boxtimes(\mathcal{S})$ is almost surely in general position.

LEMMA 2.1. *Let c_1 and c_2 be two cells of $\boxtimes(\mathcal{S})$ that are incident at a free facet. Then for any $p_1 \in c_1$ and $p_2 \in c_2$, the line segment between p_1 and p_2 is disjoint from all $s \in \mathcal{S}$.*

Proof. Let H be the hyperplane containing the facet that separates c_1 and c_2. H is part of $\boxtimes(\{s\})$ for some $s \in \mathcal{S}$. Let $\boxtimes(\{s\}) - H$ refer to the subdivision induced by the simplex s and all the hyperplanes of $\boxtimes(\{s\})$ other than H. It is easy to see that $\boxtimes(\{s\}) - H$ is a convex subdivision. Thus the overlay $\mathcal{O}$ of $\boxtimes(\mathcal{S} - \{s\})$ with $\boxtimes(\{s\}) - H$ is also a convex subdivision. The cells c_1 and c_2 lie in the same cell of $\mathcal{O}$. This implies the lemma.

COROLLARY 2.1. *If a milestone is placed in each cell of $\boxtimes(\mathcal{S})$ then any two points that can be connected by a path in $\mathcal{C}_{free}$ can also be connected by a piecewise linear path whose only internal vertices are milestones.*

Proof. Let p and q be points in $\mathcal{C}_{\text{free}}$ that can be connected by a feasible path Π. Let $\{c_1, c_2, \ldots, c_k\}$ be the sequence of cells of $\boxtimes(\mathcal{S})$ traversed by Π, and let m_i be a milestone in c_i. By Lemma 2.1, the piecewise linear path with vertices $\{p, m_1, m_2, \ldots, m_k, q\}$ is feasible. Figure 1 illustrates this.

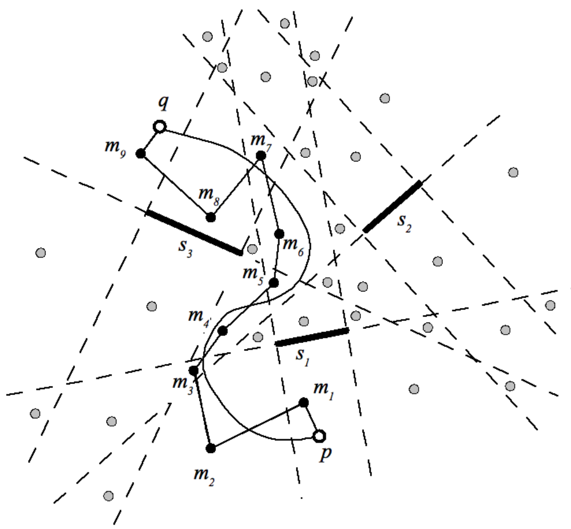

Figure 1: If two points p and q can be connected by a path in $\mathcal{C}_{\text{free}}$ they can also be connected by a linear interpolation of milestones $\{m_i\}$, as long as one is placed in each cell of the locally orthogonal decomposition of the obstacles s_1, s_2, s_3.

Volume bound. Corollary 2.1 implies that it suffices to place a milestone in every cell of $\boxtimes(\mathcal{S})$. To show that this can be accomplished with a polynomial number of samples we prove a high-probability lower bound on the volume of each cell of $\boxtimes(\mathcal{S})$. This is achieved with the help of the following simple lemma.

Lemma 2.2. *Let $\mathcal{A}(\mathcal{H})$ be the arrangement induced by a set of hyperplanes $\mathcal{H}$. If every vertex v of $\mathcal{A}(\mathcal{H})$ is ε-distant from every hyperplane $H \in \mathcal{H}$ for which $v \notin H$, then the volume (k-dimensional measure) of any k-face of the arrangement is at least $\varepsilon^k/k!$, for $1 \leq k \leq d$.*

Proof. Observe that for $0 \leq k < d$, the affine span of every k-face of the arrangement is defined by the intersection of $d-k$ hyperplanes in $\mathcal{H}$, and any vertex not contained in the span is ε-distant from at least one of these hyperplanes by assumption (the vertex cannot lie on all the $d-k$ hyperplanes, otherwise it would be contained in the span). So every vertex is ε-distant from the affine span of every k-face of the arrangement that does not contain it.

We will now prove the result by induction. The base case, for $k = 1$, is trivially proved: by our previous argument any two vertices are ε-distant from each other, hence the line segment (1-face) connecting them must have length at least ε. Now assume the hypothesis is true for $k-1$. Consider a k-face $f^{(k)}$, and let u be any vertex and $f^{(k-1)}$ a $(k-1)$-face of $f^{(k)}$. Connect u to every point of $f^{(k-1)}$ with straight line segments, forming a hyperpyramid T with vertex u and base $f^{(k-1)}$. By convexity of $f^{(k)}$ (proved above), each of these segments is entirely in $f^{(k)}$, so T itself is entirely in $f^{(k)}$. So the volume of T is at most the volume of $f^{(k)}$. It is a standard result that the volume of a k-dimensional hyperpyramid is hA/k, where h is the height of the vertex above the base and A is the area of the base. By the induction hypothesis A is at least $\varepsilon^{k-1}/(k-1)!$, and we have shown above that h must be at least ε, so we conclude that the volume of $f^{(k)}$ is at least $\varepsilon^k/k!$. The result is thus proved by induction.

Lemma 2.2 implies that volume bounds can be proved through vertex-hyperplane separation bounds. Accordingly, Section 3 is devoted to proving the following theorem:

Theorem 2.1. *Consider a vertex v and a hyperplane H of $\boxtimes(\mathcal{S})$ such that $v \notin H$, and let $\Delta := \min\{1, D_{\text{in}}\}$. Given $\varepsilon \in [0, \Delta)$, v is ε-close to H with probability at most*

$$O\left(\varepsilon^{1-\alpha} \max\{\gamma, \gamma^{d^2}\}\right)$$

for any $\alpha > 0$.

Note that all terms involving only the constants d and D are subsumed into the $O(.)$ notation. The number of hyperplanes in $\boxtimes(\mathcal{S})$ is $O(n)$ and the number of vertices of $\boxtimes(\mathcal{S})$ is $O(n^d)$. A union bound and an application of Lemma 2.2 thus yield the following corollary to Theorem 2.1.

Corollary 2.2. *Given $\varepsilon \in \left[0, \frac{\Delta^d}{d!}\right)$, the probability that some cell of $\boxtimes(\mathcal{S})$ has volume less than or equal to ε is*

$$O\left(n^{d+1} \varepsilon^{\frac{1-\alpha}{d}} \max\{\gamma, \gamma^{d^2}\}\right)$$

for any $\alpha > 0$. Hence each cell has volume at least ε with probability at least $1 - \omega$ if

$$\varepsilon \;\leq\; \min\left\{K\; \omega^{\frac{d}{1-\alpha}}\; n^{-\frac{d(d+1)}{1-\alpha}} \left(\max\{\gamma, \gamma^{d^2}\}\right)^{-\frac{d}{1-\alpha}}, \frac{\Delta^d}{d!}\right\}$$

for an appropriate constant K.

If each cell of $\boxtimes(\mathcal{S})$ has volume at least ε, standard probability theory implies that the expected number of samples sufficient for placing a milestone in every cell is $O\left(\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}\right)$ [31]. Applying Corollary 2.2, we conclude that with high probability, a set of $\text{Poly}(n,\gamma)$ samples from Σ is expected to place a milestone in every cell of $\boxtimes(\mathcal{S})$. This yields our main theorem, which we state in the special case of Gaussian perturbations.

THEOREM 2.2. *Let a free configuration space be defined by n $(d-1)$-simplices in $\mathbb{R}^d$ within a fixed domain. If independent Gaussian perturbations of variance σ^2 are applied to the simplex vertices then the expected number of uniformly chosen random samples required to construct an accurate probabilistic roadmap is polynomial in n and $\frac{1}{\sigma}$.*

3 Distance Bounds

This section forms the technical bulk of the analysis and is devoted to proving Theorem 2.1, which upper-bounds the probability that a vertex v of $\boxtimes(\mathcal{S})$ and a hyperplane H of $\boxtimes(\mathcal{S})$ are ε-close. The one-dimensional case admits a simple proof which does not require the decomposition machinery, so we assume $d \geq 2$ in the balance of this paper. H can fall into three categories:

1. The affine span of $s \in \mathcal{S}$.
2. A wall spanned by a facet of $s \in \mathcal{S}$.
3. A hyperplane defining the boundary of Σ.

We analyze these cases separately, devoting a subsection to each.

3.1 Affine Spans of Simplices

THEOREM 3.1. *Consider a fixed point p in $\mathbb{R}^d$. Given $0 \leq k < d$, let $k+1$ points $U = \{u_1, u_2, \ldots, u_{k+1}\}$ be distributed independently and γ-smoothly in Σ. The probability that the affine span of U is ε-close to p is at most*

$$K\varepsilon^{d-k}\gamma^{k+1}$$

for $\varepsilon \geq 0$ and for a constant K depending on k, d and D.

Proof. For $k = 0$ the result is trivial. Assume $1 \leq k \leq \frac{d}{2}$. We will integrate over all k-flats formed by $(k+1)$-tuples of points. For a given u_1, the k-subspace $F - u_1$ of $\mathbb{R}^d$ can be represented as the span of k orthogonal unit vectors $v_1, v_2, \ldots, v_k$. These are chosen recursively as follows. Consider any $(d-k+1)$-subspace of $\mathbb{R}^d$ and let S_1 be its unit $(d-k)$-sphere. It's easy to check that any k-subspace of $\mathbb{R}^d$ necessarily intersects S_1 (at exactly two points under general position) and each point of S_1 lies in some such subspace. So v_1 can be chosen as the vector from the origin to a point in S_1. Let H_1 be the $(d-1)$-dimensional space orthogonal to v_1. We must now pick a $(k-1)$-subspace F_1 of G_1 spanned by $v_2, v_3, \ldots, v_k$ — the span of v_1 and F_1 will give us F. This is precisely the initial setting in one lower dimension: we can choose v_2 from a unit sphere S_2 of dimension $(d-1)-(k-1) = d-k$ in G_1, and recursively every vertex is selected from a unit $(d-k)$-sphere in an appropriate subspace.

This suggests that there is an onto mapping ϕ between k-tuples drawn from S_{d-k} and orthonormal bases for k-flats in $\mathbb{R}^d$. We will now attempt to define such a mapping precisely, starting from an auxiliary mapping ψ. Assume S_{d-k} is rigidly embedded in some $(d-k)$-subspace of $\mathbb{R}^d$ with center at the origin. Take some reference k-tuple $N^0 := (\hat{n}_1^0, \ldots, \hat{n}_k^0)$ from this embedding of S_{d-k} and map it to an arbitrary orthonormal k-basis $\psi(N^0) := (v_1^0, \ldots, v_k^0)$. For $1 \leq i \leq k$, let T_i be the rotation that maps $\hat{n}_i^0$ to v_i^0. Define a metric ρ_t on t-tuples of arbitrary vectors as follows: $\rho_t((v_1, \ldots, v_t), (v_1', \ldots, v_t')) := \sup_{i=1}^t \|v_i - v_i'\|$. Let N be some other k-tuple which is δ-close to N^0 under this metric, i.e., $\rho_k(N, N^0) \leq \delta$. We state the following lemma without proof.

LEMMA 3.1. *Given $\delta \geq 0$, let $(v_1, v_2, \ldots, v_t)$ and $(v_1', v_2', \ldots, v_t')$ be two t-tuples of orthonormal vectors in $\mathbb{R}^d$, such that $\|v_i - v_i'\| \leq \delta$ for $1 \leq i \leq t$. Then there exists a rotation R of $\mathbb{R}^d$ about the origin that maps v_i to v_i', for all i, and linearly displaces each unit vector by at most $K\delta$, where K is a constant that depends only on t. Such transformation R is called a $(K\delta)$-rotation.*

Now define $\psi(N) := (v_1, \ldots, v_k)$ as follows: let $v_1 := T_1(\hat{n}_1)$. Let R_2 be a $(K_2\delta)$-rotation that maps v_1^0 to v_1. Such a rotation exists by Lemma 3.1 because $\hat{n}_1$ and $\hat{n}_1^0$ are δ-close (so v_1 and v_1^0 are also δ-close). Define $v_2 := (R_2 \circ T_2)(\hat{n}_2)$. We have

$$\begin{aligned}
\|v_2 - v_2^0\| &= \|(R_2 \circ T_2)(\hat{n}_2) - T_I(\hat{n}_2^0)\| \\
&\leq \|(R_2 \circ T_2)(\hat{n}_2) - T_2(\hat{n}_2)\| \\
&\quad + \|T_2(\hat{n}_2) - T_2(\hat{n}_2^0)\|
\end{aligned}$$

Since R_2 is a $(K_2\delta)$-rotation and $\hat{n}_2$ and $\hat{n}_2^0$ are δ-close, we have $\|v_2 - v_2^0\| \leq K_2\delta + \delta = (K_2+1)\delta$. We continue the process recursively, maintaining the invariant that the partial basis we construct at any stage is $(K\delta)$-close to the appropriate prefix of $\psi(N^0)$ for a suitable constant K. Thus at the ith step, we can define a $(K_i\delta)$-rotation R_i that maps $(v_1^0, \ldots, v_{i-1}^0)$ to $(v_1, \ldots, v_{i-1})$ and then set $v_i = (R_i \circ T_i)(\hat{n}_i)$. It's simple to check that the basis thus constructed is indeed orthonormal. Taking K^* to be $\sup_i(K_i+1)$, $\rho_k(\psi(N), \psi(N^0)) \leq K^*\delta$.

Now we divide S_{d-k} into small differential elements $A_1, A_2, \ldots, A_m$ and choose a representative point $\hat{n}_i^0$ in each A_i. Given an index k-tuple $I := (i_1, \ldots, i_k)$, let

N_I^0 denote the k-tuple $(\hat{n}_{i_1}, \ldots, \hat{n}_{i_k})$. For every possible I, define the mapping described above over the set $A_{i_1} \times A_{i_2} \times \cdots \times A_{i_k}$ with N_I^0 as the reference k-tuple: call this mapping ψ_I. Let ϕ be the union of ψ_I over all possible index k-tuples: this is the required mapping.

We restrict the choice of auxiliary mappings as follows: if two index k-tuples I and J have the same first t elements ($0 \le t < k$), then the first $t+1$ transformations $T_1, \ldots, T_{t+1}$ associated with ψ_I and ψ_J must be identical. Further, if two k-tuples of points N_1 and N_2 have the same first t elements then the first $t+1$ transformations $R_1 :=$ identity$, R_2, \ldots, R_{t+1}$ employed in mapping them via ϕ must be identical. These restrictions imply that if N_t is any t-tuple of points from S_{d-k} and $\mathcal{N}$ is the set of all k-tuples with prefix N_t, then the set of $(t+1)$th elements of bases in $\phi(\mathcal{N})$ is precisely S_{d-k} (suitably embedded in $\mathbb{R}^d$ with center at the origin). In other words, ϕ is onto.

This complicated route was taken in order to ensure that k-tuples chosen from a particular sequence of differential elements of size δ in S_{d-k} are contained in a corresponding differential element of comparable size ($K^*\delta$) in the space of all orthonormal bases. We will use this fact to integrate over the latter space. Assume that the subdivision scheme for S_{d-k} has the following properties. (Diam(A_i) and Area(A_i) denote the diameter and area of A_i, respectively.)

Property 1: $\inf_i \text{Area}(A_i) \geq C \sup_i \text{Diam}(A_i)^{d-k}$ for all $i := 1 \ldots m$ and a positive constant C independent of m. That is, the differential elements are "round".

Property 2: $\sup_i \text{Area}(A_i) \to 0$ as m increases.

A proof of the existence of such a scheme is given in the appendix. Let $\delta = \sup_i \text{Diam}(A_i)$. For a particular I, consider N_I and N_I^0, with $\phi(N_I) := (v_1, \ldots, v_k)$ and $\phi(N_I^0) := (v_1^0, \ldots, v_k^0)$. Let $q := \alpha_1 v_1 + \alpha_2 v_2 + \cdots + \alpha_k v_k$ be any point on Span$(\phi(N_I))$ within distance D from the origin. The "neighbour" of q on Span$(\phi(N_I^0))$ is the point $q^0 := \alpha_1 v_{i_1}^0 + \alpha_2 v_{i_2}^0 + \cdots + \alpha_k v_{i_k}^0$. Now

$$\begin{aligned}
\|q - q^0\| &= \|\alpha_1(v_1 - v_1^0) + \cdots + \alpha_k(v_k - v_k^0)\| \\
&\le \|\alpha_1(v_1 - v_1^0)\| + \cdots + \|\alpha_k(v_k - v_k^0)\| \\
&\le kDK^*\delta.
\end{aligned}$$

So every point on N_I within Σ is $O(\delta)$-close to Span$(\phi(N_I^0))$. Now we can write

$$\Pr[F \text{ is } \varepsilon\text{-close to } p] = \sum_{(i_1,\ldots,i_k)\in\{1,\ldots,m\}^k} \Pr[A \text{ and } B]$$

where A := "F is ε-close to p", and B := "$F - u_1 =$ Span$(\phi((\hat{n}_1, \ldots, \hat{n}_k)))$", where $\hat{n}_j \in A_{i_j}$ for $1 \le j \le k$.

Write $F^0 :=$ Span$(\phi((\hat{n}_{i_1}, \ldots, \hat{n}_{i_k})))$. If B is satisfied then, within Σ, F must be contained in the set of points $(kDK^*\delta)$-close to $u_1 + F^0$. Let G be the ball of radius $kDK^*\delta$ in the linear space $F^\perp$ orthogonal to F^0. Then the required region is $G \oplus (u_1 + F^0)$, which has volume (within Σ) at most $V_{d-k}(kDK^*\delta)V_k(D) = K'\delta^{d-k}$. Each of $u_2, \ldots, u_{k+1}$ must lie within this region, so $\Pr[B] \le (\gamma K' \delta^{d-k})^k$.

Assume $\delta \ll \varepsilon$. Let G_ε be the ball of radius $\varepsilon + kDK\delta \approx \varepsilon$ in $F^\perp$. $\Pr[A \mid B]$ is 1 only if u_1 is in $G_\varepsilon \oplus (p + F^0)$, which has volume $K''\varepsilon^{d-k}$ within Σ, and is 0 otherwise. So integrating the indicator function over all possible locations of u_1, we get $\Pr[A \mid B] \le \gamma K'' \varepsilon^{d-k}$. Multiplying and applying Property 1:

$$\begin{aligned}
\Pr[A \text{ and } B] &= \Pr[A \mid B]\Pr[B] \\
&\le K'''\varepsilon^{d-k}\gamma^{k+1}\delta^{k(d-k)} \\
&\le \frac{K'''}{C^k}\varepsilon^{d-k}\gamma^{k+1}\prod_{j=1}^{k}\text{Area}(A_{i_j})
\end{aligned}$$

Summing over all possible k-tuples of indices, we have

$$\begin{aligned}
\Pr[F \text{ is } \varepsilon\text{-close to } p] &\le \frac{K'''}{C^k}\gamma^{k+1}\varepsilon^{d-k}\text{Area}(S_{d-k})^k \\
&= K\varepsilon^{d-k}\gamma^{k+1}
\end{aligned}$$

for a constant K depending only on d, k and D.

Lastly, we must handle the case $k > \frac{d}{2}$. Observe that if F is a k-flat for such k, then the orthogonal complement of $F - u_1$ is a $(d-k)$-flat which can be localized as above. Further, if two $(d-k)$-flats are defined by orthonormal bases that are pairwise $O(\delta)$-close, then their orthogonal complements must be $O(\delta)$-close within Σ (i.e., every point on one is $O(\delta)$-close to the other). Running through the above argument in this scenario yields an identical result.

The following corollary is immediate.

Corollary 3.1. *For nonnegative integers k, k' that satisfy $k + k' < d$, consider an arbitrarily distributed k'-flat F in $\mathbb{R}^d$, as well as a set $U = \{u_1, \ldots, u_{k+1}\}$ of γ-smoothly distributed points in Σ, independent of F and of each other. The shortest distance between F and Aff(U) is at most ε with probability at most*

$$K\varepsilon^{d-k-k'}\gamma^{k+1}$$

for $\varepsilon \ge 0$ and for a constant K depending on k, d and D.

Proof. For $k' = 0$, Theorem 3.1 immediately yields the result: since the point F is distributed independently of U, we can hold it fixed, apply the theorem and then integrate the result over the range of F — it is trivially verified that this last step does not change the probability bound from the second step. For $k' > 0$, fix F: by independence, the points in U retain their original distributions under this restriction. Let F_0 be the subspace of $\mathbb{R}^d$ identical to F except for translation, and let $F^\perp$ be the orthogonal complement of F_0. Evidently, the shortest distance of a point to F is preserved under orthogonal projection to $F^\perp$. F itself maps to a single point p of $F^\perp$. Further, the points in $\mathbb{R}^d$ mapping to a volume element $\mathrm{d}\sigma$ of $F^\perp$ are exactly those in $\mathrm{d}\sigma \oplus F_0$. The k'-area of any k'-flat, when restricted to Σ, is at most $V_{k'}(D)$, so the volume of the Minkowski sum (in Σ) is at most $V_{k'}(D)\mathrm{d}\sigma$. This is illustrated in Figure 2. Hence the projection $u_i^\perp$ of each u_i is $(\gamma V_{k'}(D))$-smoothly and independently distributed in $F^\perp$. Now we can apply Theorem 3.1 in the $(d-k')$-dimensional space $F^\perp$ to upper-bound the probability that p is ε-close to $\mathrm{Aff}(u_1^\perp, \ldots, u_{k+1}^\perp)$, and hence the probability that F is ε-close to $\mathrm{Aff}(U)$, by

$$K\varepsilon^{d-k-k'}\gamma^{k+1}$$

This has no dependence on F, so integrating over the distribution of F gives the same overall probability that F is ε-close to $\mathrm{Aff}(U)$. The formula simplifies to the required result.

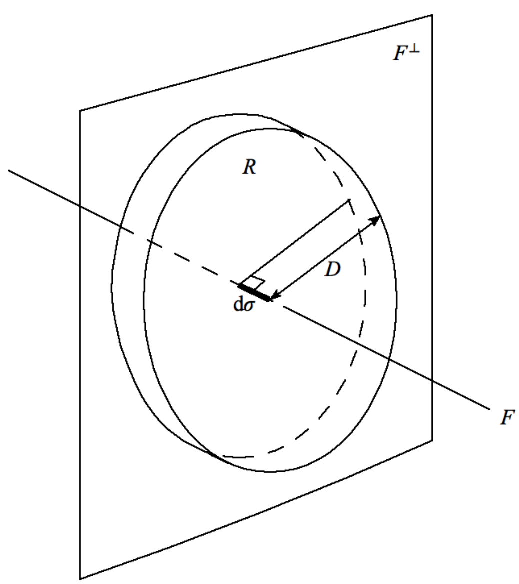

Figure 2: F is a (1D) subspace of $\mathbb{R}^3$, and $F^\perp$ is its (2D) orthogonal complement. $\mathrm{d}\sigma$ is a small volume (here, length) element of F and R is a ball of radius equal to the domain diameter D in $F^\perp$. $\mathrm{d}\sigma \oplus R$ (here, a cylinder) contains all points within the domain that orthogonally project onto $\mathrm{d}\sigma$.

From Theorem 3.1 we see that a hyperplane-vertex pair of $\boxtimes(\mathcal{S})$, in which the hyperplane is the affine span of a simplex s, and the vertex v is defined entirely by hyperplanes not associated with s, is ε-close with probability at most polynomial in ε and γ. Specifically, the bound is $K\varepsilon\gamma^d$ for a constant K depending on d and D.

Corollary 3.1 applies to the case when the vertex is formed by the intersection of one or more walls supporting s with hyperplanes not associated with s. We extend the use of the term "wall" as follows: The intersection of a number of walls of s is the wall W spanned by $\mathrm{Aff}(U)$, for a subset U of the vertices of s. Since W is orthogonal to s and contains v, we have

$$\mathrm{dist}(\mathrm{Aff}(s), v) = \mathrm{dist}(\mathrm{Aff}(U), v) \geq \mathrm{dist}(\mathrm{Aff}(U), Z)$$

where Z is the intersection of the hyperplanes unrelated to s. W and Z intersect at a point, so $\dim(W) + \dim(Z) = d$. Also, $\dim(\mathrm{Aff}(U)) = \dim(W) - 1$, and the points in U are distributed γ-smoothly and independently (of each other and of Z) in Σ. These are precisely the conditions required to apply Corollary 3.1, giving an upper bound on the probability that $\mathrm{dist}(\mathrm{Aff}(U), Z) \leq \varepsilon$, and hence on the probability that $\mathrm{dist}(\mathrm{Aff}(s), v) \leq \varepsilon$, that is again polynomial in ε and γ. Specifically, if k is the cardinality of U, then the bound is $K\varepsilon\gamma^k$ for a constant K depending on k, d and D.

3.2 Walls Supporting Simplices

When the hyperplane is a wall spanned by a simplex facet, the analysis is trickier. We will divide our work into three cases based on the interdependence of the wall and the vertex. These cases may be summarised as:

1. The wall and the vertex are entirely independent.

2. The wall and the vertex depend on the same simplex but the vertex does not lie in the affine span of that simplex.

3. The wall and the vertex depend on the same simplex and the vertex lies in the affine span of that simplex.

Case 1. We will assume that the simplex associated with the wall is entirely independent of the vertex and prove a rather general result.

THEOREM 3.2. *Consider a simplex $s \in \mathcal{S}$. For non-negative integers k, k' that satisfy $k + k' < d$, consider a subset $U = \{u_1, u_2, \ldots, u_k\}$ of the vertices of s and let W be the wall spanned by $Q := \mathit{Aff}(U)$. Let F be a k'-flat whose distribution is independent of s. The probability that W is ε-close to F is at most*

$$K\varepsilon^{d-k-k'}\gamma^{d-1}$$

for $\varepsilon \geq 0$ and for a constant K depending on k, d and D.

Proof. Let H be the affine span of the simplex. Fix F, and let F^H be the orthogonal projection of F to H. By orthogonality, it is immediate that $\mathrm{dist}(W, F) = \mathrm{dist}(Q, F^H)$. We assume a tessellation scheme of S_{d-1} into area elements $A_1, A_2, \ldots, A_m$ as in the proof of Theorem 3.1, satisfying Properties 1 and 2. Write

$$\begin{aligned}
&\Pr\left[W \text{ is } \varepsilon\text{-close to } F\right] \\
&\quad = \sum_{i=1}^{m} \Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } W \text{ is } \varepsilon\text{-close to } F\right] \\
&\quad = \sum_{i=1}^{m} \Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } Q \text{ is } \varepsilon\text{-close to } F^H\right]
\end{aligned}$$

Now fix an arbitrary normal $\hat{\mathrm{n}}_i$ in each A_i and let H_0 be the plane $\langle x, \hat{\mathrm{n}}_i\rangle = 0$. Another normal $\hat{\mathrm{n}}$ also in A_i satisfies $\|\hat{\mathrm{n}} - \hat{\mathrm{n}}_i\| \leq \mathrm{Diam}(A_i)$. We will show that when $\hat{\mathrm{n}}(H)$ is in A_i, projection to H_0 instead of to H will almost surely not change the shortest distance by "much". For this the following simple lemma is required.

Lemma 3.2. *Let v_1 and v_2 be the orthogonal projections of vector v onto hyperplanes H_1 and H_2, respectively, and assume that the normals of H_1 and H_2 are δ-close, i.e., $\|\hat{\mathrm{n}}_2 - \hat{\mathrm{n}}_1\| \leq \delta$. Then $\|v_2\| - \|v_1\| \ \leq\ 2\delta\|v\|$.*

Proof. Write $\Delta n := \hat{\mathrm{n}}_2 - \hat{\mathrm{n}}_1$. We have $v_1 = v - \langle\hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1$ and $v_2 = v - \langle\hat{\mathrm{n}}_2, v\rangle\hat{\mathrm{n}}_2$. So

$$\begin{aligned}
\|v_2\| - \|v_1\| &\leq \|v_2 - v_1\| \\
&\leq \|(v - \langle\hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1) - (v - \langle\hat{\mathrm{n}}_2, v\rangle\hat{\mathrm{n}}_2)\| \\
&= \|\langle\hat{\mathrm{n}}_2, v\rangle\hat{\mathrm{n}}_2 - \langle\hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1)\| \\
&= \|\langle\hat{\mathrm{n}}_1 + \Delta n, v\rangle(\hat{\mathrm{n}}_1 + \Delta n) - \langle\hat{\mathrm{n}}_1, v\rangle\hat{\mathrm{n}}_1)\| \\
&= \|\langle\hat{\mathrm{n}}_2, v\rangle\Delta n + \langle\Delta n, v\rangle\hat{\mathrm{n}}_1)\| \\
&\leq 2\|\Delta n\|\|v\| \ \leq\ 2\delta\|v\|
\end{aligned}$$

Let (q, f^H) be a pair in $Q \times F^H$ such that $\mathrm{dist}(q, f^H) = \mathrm{dist}(Q, F^H)$ and let f be the pre-image of f^H under the projection—if there are multiple pre-images we choose the one closest to q. Let Q^{H_0}, F^{H_0}, q^{H_0} and f^{H_0} be the orthogonal projections of Q, F, q and f respectively to H_0. By the above lemma, $\mathrm{dist}(Q^{H_0}, F^{H_0}) \leq \mathrm{dist}(q^{H_0}, f^{H_0}) \leq \mathrm{dist}(q, f^H) + 2\delta_i \mathrm{dist}(q, f) = \mathrm{dist}(Q, F^H) + 2\delta_i \mathrm{dist}(q, f)$, where $\delta_i := \mathrm{Diam}(A_i)$ (see Figure 3). For every possible configuration of s and F, $\mathrm{dist}(q, f)$ is a finite positive quantity, so given $\omega \in (0, 1)$ we can always find a large

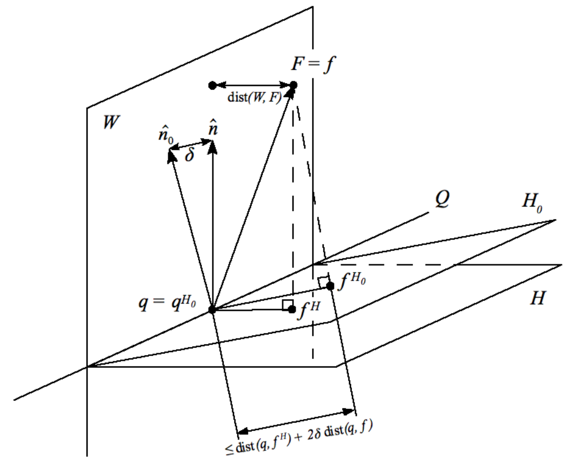

Figure 3: The projections of vector $\overrightarrow{qf}$ onto two hyperplanes H and H_0 differ by at most $2\delta\ \mathrm{dist}(q, f)$, where δ is the length of the difference of the normals of the hyperplanes. The notation is that of Theorem 3.2.

enough constant M such that $\mathrm{dist}(q, f) \leq M$ with probability at least ω. This implies

$$\begin{aligned}
&\Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } Q \text{ is } \varepsilon\text{-close to } F^H\right] \\
&\quad = \Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } \mathrm{dist}(q, f) \leq M \\
&\qquad\qquad \text{and } Q \text{ is } \varepsilon\text{-close to } F^H] \\
&\qquad + \Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } \mathrm{dist}(q, f) > M \\
&\qquad\qquad \text{and } Q \text{ is } \varepsilon\text{-close to } F^H] \\
&\quad \leq \Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } \mathrm{dist}(q, f) \leq M \\
&\qquad\qquad \text{and } Q^{H_0} \text{ is } (\varepsilon + 2\delta_i M)\text{-close to } F^{H_0}] \\
&\qquad + \Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } \mathrm{dist}(q, f) > M] \\
&\quad \leq \Pr[\hat{\mathrm{n}}(H) \in A_i \\
&\qquad\qquad \text{and } Q^{H_0} \text{ is } (\varepsilon + 2\delta_i M)\text{-close to } F^{H_0}] \\
&\qquad + \Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } \mathrm{dist}(q, f) > M]
\end{aligned}$$

Let u_d be a vertex of s not in U. For the first term, observe that $\hat{\mathrm{n}}(H)$ is in A_i only if every vertex of s is in the slab T between the parallel planes $\langle x - u_d, \hat{\mathrm{n}}_i\rangle = \pm\delta_i D$. Let $u_i^{H_0}$ be the orthogonal projection of each $u_i \in U$ on H_0 — under the above restriction, the $u_i^{H_0}$'s are $(2\gamma\delta_i D)$-smoothly and independently distributed on H_0. Corollary 3.1 can now be directly applied in H_0 to obtain, for some constants K_1 and K_2,

$$\begin{aligned}
&\Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } Q^{H_0} \text{ is } (\varepsilon + 2\delta_i M)\text{-close to } F^{H_0}\right] \\
&\quad\le \underbrace{(2\gamma\delta_i D V_{d-1}(D))^{d-k-1}}_{\text{vertices not in } U \cup \{u_d\}} \\
&\qquad \times \underbrace{K_1(\varepsilon + 2\delta_i M)^{d-k-k'}(2\gamma\delta_i D)^k}_{\text{vertices in } U} \\
&\quad\le K_2(\varepsilon + 2\delta_i M)^{d-k-k'}\gamma^{d-1}\mathrm{Area}(A_i).
\end{aligned}$$

Summing over all i,

$$\begin{aligned}
&\Pr\left[W \text{ is } \varepsilon\text{-close to } F\right] \\
&\quad\le \sum_i K_2(\varepsilon + 2\delta_i M)^{d-k-k'}\gamma^{d-1}\mathrm{Area}(A_i) \\
&\qquad\qquad + \sum_i \Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } \mathrm{dist}(q, F) > M\right] \\
&\quad\le K_2(\varepsilon + 2\sup_i \delta_i M)^{d-k-k'}\gamma^{d-1}\mathrm{Area}(S_{d-1}) \\
&\qquad\qquad\qquad + \Pr\left[\mathrm{dist}(q, F) > M\right] \\
&\quad\le (8\pi^2 K_2/15)(\varepsilon + 2\sup_i \delta_i M)^{d-k-k'}\gamma^{d-1} + (1-\omega)
\end{aligned}$$

Make ω arbitrarily close to 1 and choose small enough area elements so that $\sup_i \delta_i M \ll \varepsilon$, thus obtaining

$$\Pr\left[W \text{ is } \varepsilon\text{-close to } F\right] \le K\varepsilon^{d-k-k'}\gamma^{d-1}$$

for a constant K. By independence, integrating over the range of F does not change the expression.

Setting $k = d-1$ and $k' = 0$ yields the required vertex-wall separation result: The probability of ε-closeness is at most $K\varepsilon\gamma^{d-1}$.

Case 2. The next case to be treated is when the simplex associated with the wall is involved in the definition of the vertex but does not contain it in its affine span.

THEOREM 3.3. *Consider a simplex $s \in \mathcal{S}$. Given a set $U = \{u_1, u_2, \ldots, u_k\}$ of k vertices of s, for $1 \le k < d$, define $Q := \mathrm{Aff}(U)$. Let Z be the wall spanned by Q, let F be a $(d-k)$-flat whose distribution is independent of s, and define $v := Z \cap F$. Let W be a wall of $\boxtimes(\{s\})$ that does not contain Z. Given $\varepsilon \in [0, 1)$, the probability that v is ε-close to W is at most*

$$K\varepsilon^{1-\alpha}\gamma^{d-1}$$

for any $\alpha > 0$ and a constant K depending on α, k, d and D.

Proof. For $d \le 2$ the proof is trivial. Assume $d > 2$. Let H be the affine span of the simplex. Assume, without loss of generality, that W is spanned by $W_b := \mathrm{Aff}(u_2, u_3, \ldots, u_d)$. The intersection of W and Z is the wall Y spanned by $Y_b := \mathrm{Aff}(u_2, u_3, \ldots, u_k)$. Consider the $(d-k+1)$-dimensional linear space $Y^\perp$ orthogonal to Y, in which Y itself orthogonally projects to a point y and Z projects to a line L. Note that $Y^\perp$ must contain $\hat{\mathrm{n}}(W)$. Let $\vec{a}$ be the vector $v^\perp - y$, where $v^\perp$ is the orthogonal projection of v to $Y^\perp$, lying on L. If v is ε-close to W, then $|\langle \vec{a}, \hat{\mathrm{n}}(W)\rangle| \le \varepsilon$, i.e., $v^\perp$ is $(\varepsilon \sec\Theta)$-close to y, where Θ is the (acute) angle between L and $\hat{\mathrm{n}}(W)$. This means that $\mathrm{dist}(F, Y) \le \mathrm{dist}(v, Y) = \mathrm{dist}(v^\perp, y) \le \varepsilon\sec\Theta$. Hence for infinitesimally small $\mathrm{d}\theta$,

$$\begin{aligned}
&\Pr\left[v \text{ is } \varepsilon\text{-close to } W \text{ and } \Theta \in [\theta, \theta + \mathrm{d}\theta)\right] \\
&\quad\le \Pr\left[F \text{ is } (\varepsilon\sec\theta)\text{-close to } Y \text{ and } \Theta \in [\theta, \theta + \mathrm{d}\theta)\right]
\end{aligned}$$

We now localize the normal of H and follow the proof of Theorem 3.2 with a few changes. Specifically, the probability

$$\Pr\left[\hat{\mathrm{n}}(H) \in A_i \text{ and } Q^{H_0} \text{ is } (\varepsilon + 2\delta_i M)\text{-close to } F^{H_0}\right]$$

is replaced with the probability $P = \int_0^{\pi/2} P_\theta$, where

$$\begin{aligned}
P_\theta \;:=\; &\Pr[\hat{\mathrm{n}}(H) \in A_i \text{ and } \Theta \in [\theta, \theta + \mathrm{d}\theta)] \\
&\text{and } Y_b^{H_0} \text{ is } ((\varepsilon + 2\delta_i M)\sec\theta)\text{-close to } F^{H_0}]
\end{aligned}$$

(The differential $\mathrm{d}\theta$ will be shown to be present as a factor in P_θ.)

The fixed point is taken to be u_d as before. Note that if W_b is fixed then L, and hence the angle Θ, depends only on the position of u_1. In $Y^\perp$, let R be the region between the double cones with vertex y, axis $\hat{\mathrm{n}}(W)$ and half-angles θ and $\theta + \mathrm{d}\theta$. Evidently, Θ lies in the required range if and only if u_1 lies in the extruded region $R_Y := R \oplus Y_b$. This is illustrated in Figure 4.

Even if W_b and hence Y_b are fixed, R depends on $Y^\perp$ and thus on $\hat{\mathrm{n}}(W)$ and on u_1. This yields a circularity. However, $\hat{\mathrm{n}}(W)$ can be approximated by a single normal $\hat{\mathrm{n}}_i$ in A_i — the approximation improves as A_i shrinks. This fixes Y independently of u_1 (denote this value by Y_0) and places R in $Y_0^\perp$ as the region R^0, which extrudes to R_Y^0 in H_0. The required probability can now be approximated as

$$\Pr\left[u_1 \text{ lies in } R_Y\right] \approx \Pr\left[u_1^{H_0} \text{ lies in } R_Y^0\right],$$

where $u_1^{H_0}$ is the orthogonal projection of u_1 to H_0. R^0, in domain H_0, has measure at most

$$\frac{2\mathrm{Area}(S_{d-k-1}(D\sin\theta)) \times D^2\mathrm{d}\theta}{d-k+1} \le \frac{2(d-k)DV_{d-k}(D)\mathrm{d}\theta}{d-k+1}$$

as may be verified by picturing R^0 as a rotational sweep of a 2D double-cone. This implies that the volume of

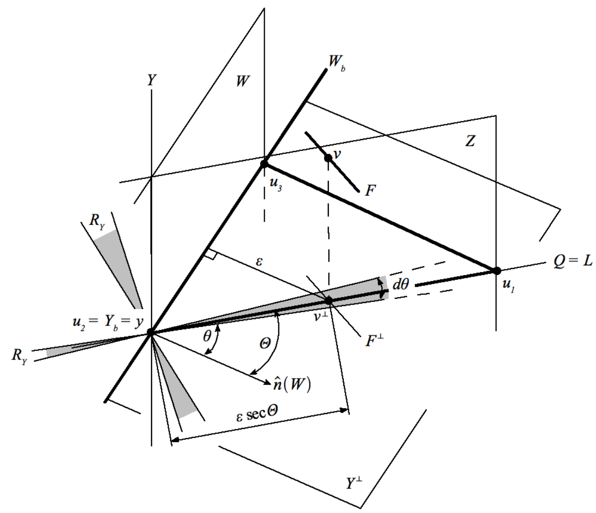

Figure 4: A vertex v is formed by the intersection of flat F and wall Z spanned by the simplex $u_1u_2u_3$. Using the notation of Theorem 3.3, if L is at an angle $\Theta \in [\theta, \theta + \mathrm{d}\theta)$ to $\hat{n}(W)$ in $Y^\perp$, and v is at distance ε from W, then the projection $v^\perp$ of v to $Y^\perp$ is at distance $\approx \varepsilon \sec\theta$ from y, the projection of Y to $Y^\perp$. L is at the required angle if and only if u_1 is in the shaded region R_Y.

R_Y^0 within Σ is at most

$$\frac{2(d-k)\ D^{k-1}\ V_{d-k}(D)\ \mathrm{d}\theta}{d-k+1}.$$

We can now evaluate the required probability using

$$\begin{aligned} P_\theta \quad &\leq \quad \Pr[\hat{n}(H) \in A_i \text{ and } \Theta \in [\theta, \theta + \mathrm{d}\theta) \ \mid\ B_1] \\ &\qquad \times\ \Pr[\hat{n}(H) \in A_i \\ &\qquad\quad \text{and } Y_b^{H_0} \text{ is } ((\varepsilon + 2\delta_i M)\sec\theta)\text{-close to } F^{H_0}] \\ &\leq \quad \Pr[u_1 \in T \text{ and } u_1^{H_0} \text{ is in } R_Y^0 \ \mid\ B_2] \\ &\qquad \times\ \Pr[\{u_2, \ldots, u_{d-1}\} \subset T \\ &\qquad\quad \text{and } Y_b^{H_0} \text{ is } ((\varepsilon + 2\delta_i M)\sec\theta)\text{-close to } F^{H_0}] \end{aligned}$$

where B_1 and B_2 are the conditions in the corresponding second factors in the two lines and T is the usual $2\delta_i D$-thick slab for localizing the normal to the differential element. Note that the first factor in the last line depends only on u_1 and the second only on $u_2, \ldots, u_d$. The first factor is

$$\begin{aligned} &\Pr\left[u_1 \in T \text{ and } u_1^{H_0} \text{ is in } R_Y^0 \ \mid\ B_2\right] \\ &\qquad \leq \quad 2\gamma\delta_i D \frac{2(d-k)\ D^{k-1}\ V_{d-k}(D)\ \mathrm{d}\theta}{d-k+1} \\ &\qquad = \quad K_1\gamma\delta_i \mathrm{d}\theta \end{aligned}$$

for an appropriate constant K_1. The second factor is as in Theorem 3.2 (minus the vertex u_1 of U, and with an extra $\sec\theta$ factor), i.e. it is at most

$$K_2(\varepsilon + 2\delta_i M)\sec\theta\gamma^{d-2}\delta_i^{d-2},$$

for another constant K_2. Multiplying the bounds yields

$$\begin{aligned} P_\theta \quad &\leq \quad K_1K_2(\varepsilon + 2\delta_i M)\sec\theta\gamma^{d-1}\delta_i^{d-1}\mathrm{d}\theta \\ &\leq \quad \frac{K_1K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\mathrm{Area}(A_i)\sec\theta\mathrm{d}\theta \end{aligned}$$

The probability P is thus $\int_0^{\pi/2} P_\theta$, which is at most

$$\begin{aligned} &\frac{K_1K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\mathrm{Area}(A_i)\int_0^{\pi/2} \sec\theta\mathrm{d}\theta \\ &= \frac{K_1K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\mathrm{Area}(A_i)[\log(\sec\theta + \tan\theta)]_0^{\pi/2} \end{aligned}$$

Unfortunately this integral is unbounded. To circumvent this problem we write $P = P_a + P_b$, for

$$P_a = \int_{\pi/2-(\varepsilon+2\delta_i M)}^{\pi/2} P_\theta \quad \text{and} \quad P_b = \int_0^{\pi/2-(\varepsilon+2\delta_i M)} P_\theta$$

Then

$$\begin{aligned} P_a \ &\leq \ \int_{\pi/2-(\varepsilon+2\delta_i M)}^{\pi/2} \Pr\left[u_1 \in T \text{ and } u_1^{H_0} \text{ is in } R_Y^0 \mid B_2\right] \\ &\qquad\qquad \times \Pr\left[\{u_2, \ldots u_{d-1}\} \subset T\right] \\ &\leq \ K_1\gamma\delta_i(\varepsilon + 2\delta_i M) \times K_3\gamma^{d-2}\delta_i^{d-2} \\ &\leq \ \frac{K_1K_3}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\mathrm{Area}(A_i) \end{aligned}$$

for some constant K_3, and

$$\begin{aligned} P_b \ &\leq \ \frac{K_1K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\mathrm{Area}(A_i) \\ &\qquad \times [\log(\sec\theta + \tan\theta)]_0^{\pi/2-(\varepsilon+2\delta_i M)} \end{aligned}$$

Now, for $0 < x < \pi/2$,

$$\begin{aligned} \log(\sec\theta + \tan\theta)\ |_{\pi/2-x} \ &= \ \log(\csc x + \cot x) \\ &= \ \log\frac{1 + \cos x}{\sin x} \\ &\leq \ \log\frac{2}{2x/\pi} \\ &= \ \log\frac{\pi}{x} \ \leq \ K_\alpha\left(\frac{\pi}{x}\right)^\alpha \end{aligned}$$

for any $\alpha > 0$ and a constant K_α depending on α. Thus

$$\begin{aligned} P_b \ &\leq \ \frac{K_1K_2}{C}(\varepsilon + 2\delta_i M)\gamma^{d-1}\mathrm{Area}(A_i)K_\alpha\left(\frac{\pi}{\varepsilon + 2\delta_i M}\right)^\alpha \\ &= \ \frac{K_1K_2K_\alpha\pi^\alpha}{C}\ (\varepsilon + 2\delta_i M)^{1-\alpha}\ \gamma^{d-1}\mathrm{Area}(A_i) \end{aligned}$$

The previous arguments imply that we can choose small enough area elements so that $\varepsilon + 2\sup_i \delta_i M \to \varepsilon < 1$. Therefore,

$$P \;=\; P_a + P_b \;\leq\; K_4\varepsilon^{1-\alpha}\,\gamma^{d-1}\text{Area}(A_i)$$

for another constant K_4. Summing over all i,

$$\begin{aligned}\Pr[v \text{ is } \varepsilon\text{-close to } W] \;&\leq\; K_4\varepsilon^{1-\alpha}\,\gamma^{d-1}\,\text{Area}(S_{d-1})\\ &\leq\; \frac{8\pi^2 dK_4}{15}\varepsilon^{1-\alpha}\,\gamma^{d-1}\end{aligned}$$

Case 3. The third case is when the affine span of the simplex associated with the wall is one of the hyperplanes defining the vertex.

THEOREM 3.4. *Consider a simplex $s \in \mathcal{S}$. Given a set $U = \{u_1, u_2, \ldots, u_k\}$ of k vertices of s, for $1 \leq k \leq d$, define $Q := Aff(U)$. Let F be a $(d-k+1)$-flat whose distribution is independent of s, and define $v := Q \cap F$. Let W be a wall of $\boxtimes(\{s\})$ that does not contain Q. Given $\varepsilon \in [0,1)$, the probability that v is ε-close to W is at most*

$$K\varepsilon^{1-\alpha}\max\{\gamma, \gamma^k\}$$

for any $\alpha > 0$ and a constant K depending on α, k, d and D.

Proof. The proof is similar to, but simpler than, that of Theorem 3.3. Because the intersection point lies in the affine span H of the simplex, the localization of the normal and subsequent projection onto this span is unnecessary. Assume, without loss of generality, that W stands on $W_b := \text{Aff}(u_2, u_3, \ldots, u_d)$. The intersection of W_b and Q is $Y := \text{Aff}(u_2, u_3, \ldots, u_k)$. Since the simplex and the wall are orthogonal, $\text{dist}(W, v) = \text{dist}(W_b, v)$ and we can restrict our attention to the hyperplane H: our next few comments will pertain strictly to this domain. Let $Y^\perp$ be the orthogonal complement of Y (w.r.t. H). The orthogonal projection of Y to $Y^\perp$ is the single point y, that of Q is the line L, and that of v is a point $v^\perp$ lying on L. Let the normal to W_b be $\hat{\text{n}}(W_b)$. $\text{dist}(W_b, v) = \text{dist}(y, v^\perp)\cos\theta \geq \text{dist}(Y, F)\cos\theta$, where θ is the measure of the angle Θ between $\hat{\text{n}}(W_b)$ and $v^\perp - y$. Let R be the region between the double cones with vertex y, axis $\hat{\text{n}}(W_b)$ and half-angles θ and $\theta + \text{d}\theta$. Evidently, Θ lies in the required range iff u_1 lies in the extruded region $R_Y := R \oplus Y$, which has volume at most

$$\frac{2(d-k)\;D^{k-1}\;V_{d-k}(D)\;\text{d}\theta}{d-k+1}$$

Now we shift our attention back to the full-dimensional domain. Given $\{u_2, u_3, \ldots, u_d\}$, Θ is in the required range iff u_1 lies in the region R' swept out by rotating R_Y around the axis W_b. By a rough estimate, the volume of this region is at most $2\pi D\;\text{Vol}(R_Y)$, so

$$\Pr[\Theta \in [\theta, \theta + \text{d}\theta)] \;=\; \gamma\text{Vol}(R') \;=\; K_1\gamma\text{d}\theta$$

for a suitable constant K_1 depending on k, d and D, and

$$\begin{aligned}P_\theta \;:=&\; \Pr[\Theta \in [\theta, \theta + \text{d}\theta) \text{ and } Y \text{ is } (\varepsilon\sec\theta)\text{-close to } F]\\ \leq&\; K_1\gamma\text{d}\theta \times K_2\varepsilon\sec\theta\gamma^{k-1}\end{aligned}$$

where K_2 is another constant depending on k, d and D, from Corollary 3.1. Now as in the proof of Theorem 3.3, integrating this upper bound for P_θ over $[0, \pi/2]$ gives an unbounded result. We reuse our earlier hack to solve this problem. First,

$$\int_{\pi/2-\varepsilon}^{\pi/2} P_\theta \;\leq\; \int_{\pi/2-\varepsilon}^{\pi/2} \Pr[\Theta \in [\theta, \theta + \text{d}\theta)] \;\leq\; K_1\varepsilon\gamma$$

Next, for any $\alpha > 0$,

$$\int_0^{\pi/2-\varepsilon} P_\theta \;\leq\; K_1K_2K_\alpha\varepsilon^{1-\alpha}\gamma^k$$

for a constant K_α depending on α. Putting everything together,

$$\int_0^{\pi/2} P_\theta \;\leq\; K\varepsilon^{1-\alpha}\max\{\gamma, \gamma^k\}$$

for a constant K depending on α, k, d and D.

3.3 The Boundary of the Domain

For this final case, we must bound the probability that a vertex v of $\boxtimes(\mathcal{S})$ not contained in a hyperplane H constituting the boundary of the domain Σ is ε-close to it. Since this hyperplane is fixed, we must consider the distribution of the vertex instead. A non-boundary vertex of $\boxtimes(\mathcal{S})$ is defined by the intersection of h_1 hyperplanes associated with one simplex, h_2 hyperplanes associated with another and so on, where $\sum_i h_i = d$. If v lies in a small region of volume $\text{d}\sigma$ with centre p and diameter δ, then *all* of these hyperplanes must pass through that region. There are two possible cases for the set of h hyperplanes associated with a particular simplex s:

Case 1: The hyperplanes are all walls supporting the simplex. Their intersection is a $(d-h)$-dimensional "wall" Z standing on the affine span of $d-h$ vertices of s. Theorem 3.2, with $k = d-h$ and $k' = 0$, tells us that the probability that Z is δ-close to p is at most $K\delta^h\gamma^{d-1}$.

Case 2: The hyperplanes include the affine span of the simplex itself. Then their intersection is simply the affine span of a $(d-h)$-face of the simplex and Theorem 3.1, with $k = d-h$, tells us that the probability that Z is δ-close to p is at most $K\delta^h\gamma^{d-h+1}$.

So the probability that all the hyperplanes pass through $\mathrm{d}\sigma$ is at most

$$\prod_i K_i \delta^{h_i} \max\{\gamma, \gamma^d\} \leq K' \delta^d \max\{\gamma, \gamma^{d^2}\}$$
$$\leq K'' \mathrm{d}\sigma \max\{\gamma, \gamma^{d^2}\}$$

(The last step assumes that the small region is "round", i.e. $\mathrm{d}\sigma = \Theta(\delta^d)$.) In other words, the vertex v follows a $K'' \max\{\gamma, \gamma^{d^2}\}$-smooth distribution. The portion of the domain Σ within ε distance of the hyperplane H has volume at most $\varepsilon V_{d-1}(D)$, so the probability that v is ε-close to H is at most $K'' \varepsilon \max\{\gamma, \gamma^{d^2}\} V_{d-1}(D)$.

If the vertex lies on the boundary, we must modify our analysis only slightly. Constrain ε to be less than D_{in}, and assume b hyperplanes from the boundary contain v. If $h_1, h_2, \ldots$ hyperplanes associated with simplices also contain v as before, then it must be the case that $\sum_i h_i = d - b$. The b hyperplanes on the boundary intersect in a $(d-b)$-flat B. Carry out the previous analysis assuming that the differential region is a subset of B and not of the full-dimensional space. Then $\mathrm{d}\sigma = \Theta(\delta^{d-b})$ and we obtain the result that v is $K''' \max\{\gamma, \gamma^{d(d-b)}\}$-smoothly distributed on B. The boundary is fixed, so B and H are at a constant angle. If this angle is zero (B and H are parallel), then B is D_{in}-distant, and hence ε-distant, from H. Else, the $(d-b)$-measure of the region of B within Σ ε-close to H is at most $C \varepsilon V_{d-b-1}(D)$ for some constant C. The probability that v lies in this region is at most $K''' C \varepsilon \max\{\gamma, \gamma^{d(d-b)}\} V_{d-b-1}(D)$.

This concludes the proof of Theorem 2.1.

4 Conclusion

The following extensions of the presented analysis naturally suggest themselves and are left for future work.

- **Translational motion planning.** The free configuration space for translational motion planning of a polyhedral robot among polyhedral obstacles is the complement of the union of Minkowski sums of the obstacles and the antipode of the robot [34]. Independently perturbed simplices cannot model this setting since the connectivity of the Minkowski sums is not preserved. Extending our results to translational motion planning necessitates relieving the reliance of the current analysis on complete independence of the perturbations. Preliminary derivations suggest that the limited amount of independence present in the free configuration space of translational motion planning is sufficient to obtain a polynomial bound on the number of milestones.
- **General motion planning and curved $\mathcal{C}$-space obstacles.** General motion planning, such as holonomic or articulated motion with translations and rotations, gives rise to configuration spaces with curved $\mathcal{C}$-space obstacles, generally represented as semi-algebraic sets. In order to do smoothed analysis in this setting, a convincing perturbation model for semi-algebraic sets needs to be defined. Building on this definition, a polynomial number of random samples has to be shown to yield an accurate roadmap.
- **Connection to previous theoretical models.** It is reasonable to conjecture that smoothly perturbed free configuration spaces are $(\varepsilon, \alpha, \beta)$-expansive, for appropriate values of ε, α, and β [22]. This would imply that results previously obtained for expansive configuration spaces carry over to the smoothed setting.

Acknowledgements

We are grateful to Dan Halperin for crucial early discussions and to Jean-Claude Latombe for his support throughout this project. We also acknowledge the suggestions of anonymous referees.

References

[1] P. K. Agarwal, M. J. Katz, and M. Sharir. Computing depth orders for fat objects and related problems. *Computational Geometry Theory and Applications*, 5:187–206, 1995.

[2] B. Aronov and M. Sharir. Castles in the air revisited. *Discrete and Computational Geometry*, 12:119–150, 1994.

[3] D. Arthur and S. Vassilvitskii. Worst-case and smoothed analyses of the ICP algorithm, with an application to the k-means method. In *Proc. 47th IEEE Symposium on Foundations of Computer Science*, 2006.

[4] C. Banderier, R. Beier, and K. Mehlhorn. Smoothed analysis of three combinatorial problems. In *Proc. 28th Symposium on Mathematical Foundations of Computer Science*, pages 198–207, 2003.

[5] J. Barraquand, L. E. Kavraki, J.-C. Latombe, T.-Y. Li, R. Motwani, and P. Raghavan. A random sampling scheme for path planning. *International Journal of Robotics Research*, 16:759–774, 1997.

[6] R.-P. Berretty, M. Overmars, and A. F. van der Stappen. Dynamic motion planning in low obstacle density environments. In *Proc. Workshop Algorithms Data Struct.*, volume 1272 of *Lecture Notes in Computer Science*, pages 3–16. Springer-Verlag, 1997.

[7] A. Blum and J. Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms*, pages 905–914, 2002.

[8] D. E. Cardoze and L. J. Schulman. Pattern matching for spatial point sets. In *Proc. 39th IEEE Symposium on Foundations of Computer Science*, pages 156–165, 1998.

[9] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoretical Computer Science*, 84:77–105, 1991. Also in *Proc. International Colloquium on Automata, Languages and Programming*, pages 179–193, 1989.

[10] K. L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete and Computational Geometry*, 22:63–93, 1999.

[11] V. Damerow, F. M. auf der Heide, H. Räcke, C. Scheideler, and C. Sohler. Smoothed motion complexity. In *Proc. 11th European Symposium on Algorithms*, pages 161–171, 2003.

[12] V. Damerow and C. Sohler. Smoothed number of extreme points under uniform noise. In *Proc. 20st European Workshop on Computational Geometry*, 2004.

[13] M. de Berg. Linear size binary space partitions for fat objects. In *Proc. 3rd Annu. European Sympos. Algorithms*, volume 979 of *Lecture Notes in Computer Science*, pages 252–263. Springer-Verlag, 1995.

[14] M. de Berg. Linear size binary space partitions for uncluttered scenes. *Algorithmica*, 28:353–366, 2000.

[15] M. de Berg, D. Halperin, M. Overmars, and M. van Kreveld. Sparse arrangements and the number of views of polyhedral scenes. *International Journal of Computational Geometry and Applications*, 7:175–195, 1997.

[16] M. de Berg, M. J. Katz, A. F. van der Stappen, and J. Vleugels. Realistic input models for geometric algorithms. In *Proc. 13th Symposium on Computational Geometry*, pages 294–303, 1997.

[17] A. Deshpande and D. A. Spielman. Improved smoothed analysis of the shadow vertex simplex method. In *Proc. 46th IEEE Symposium on Foundations of Computer Science*, pages 349–356, 2005.

[18] J. Erickson. Nice point sets can have nasty delaunay triangulations. *Discrete and Computational Geometry*, 30:109–132, 2003.

[19] M. Gavrilov, P. Indyk, R. Motwani, and S. Venkatasubramanian:. Geometric pattern matching: A performance study. In *Proc. 15th Symposium on Computational Geometry*, pages 79–85, 1999.

[20] R. Geraerts and M. Overmars. A comparative study of probabilistic roadmap planners. In *Proc. 5th Workshop on the Algorithmic Foundations of Robotics*, pages 43–59, 2002.

[21] D. Hsu, J.-C. Latombe, and H. Kurniawati. On the probabilistic foundations of probabilistic roadmap planning. *International Journal of Robotics Research*. to appear.

[22] D. Hsu, J.-C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. *International Journal of Computational Geometry and Applications*, 9:495–512, 1999.

[23] D. Hsu, J.-C. Latombe, R. Motwani, and L. E. Kavraki. Capturing the connectivity of high-dimensional geometric spaces by parallelizable random sampling techniques. In P. Pardalos and S. Rajasekaran, editors, *Advances in Randomized Parallel Computing*, pages 159–182. Kluwer Academic Publishers, Boston, MA, 1999.

[24] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Robotics and Automation*, 14:166–171, 1998.

[25] L. E. Kavraki and J.-C. Latombe. Probabilistic roadmaps for robot path planning. In K. Gupta and A. del Pobil, editors, *Practical Motion Planning in Robotics: Current Approaches and Future Directions*, pages 33–53. John Wiley, 1998.

[26] L. E. Kavraki, J.-C. Latombe, R. Motwani, and P. Raghavan. Randomized query processing in robot path planning. In *Proc. 27th ACM Symposium on Theory of Computing*, pages 353–362, 1995.

[27] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 11:566–580, 1996.

[28] V. Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. *Journal of the ACM*, 51:699–730, 2004.

[29] J.-C. Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, Boston, 1991.

[30] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. E. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital Michelangelo project: 3D scanning of large statues. In *SIGGRAPH*, pages 131–144, 2000.

[31] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, New York, NY, 1995.

[32] M. H. Overmars and A. F. van der Stappen. Range searching and point location among fat objects. *Journal of Algorithms*, 21:629–656, 1996.

[33] O. Schwarzkopf and J. Vleugels. Range searching in low-density environments. *Information Processing Letters*, 60:121–127, 1996.

[34] M. Sharir. Algorithmic motion planning. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 40, pages 733–754. CRC Press LLC, Boca Raton, FL, 1997.

[35] G. Song, S. L. Thomas, and N. M. Amato. A general framework for PRM motion planning. In *Proc. International Conference on Robotics and Automation*, pages 4445–4450, 2003.

[36] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms. In *Proc. of the International Congress of Mathematicians 2002*.

[37] D. A. Spielman and S.-H. Teng. Smoothed analysis of termination of linear programming algorithms. *Mathematical Programming, Series B*, 97, 2003.

[38] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51:385–463, 2004.

[39] A. F. van der Stappen and M. H. Overmars. Motion planning amidst fat obstacles. In *Proc. 10th Symposium on Computational Geometry*, pages 31–40, 1994.

Appendix

In this section we prove that there is an infinite sequence of tessellations of the unit sphere S_{k-1} ($k > 1$) into a set of regions $\{A_1, \ldots, A_m\}$, each tessellation finer than the last, which satisfies Properties 1 and 2. To recap, there exists a positive constant C independent of m such that

Property 1: $\inf_i \text{Area}(A_i) \geq C \sup_i \text{Diam}(A_i)^{k-1}$ for all $i := 1 \ldots m$.

Property 2: $\sup_i \text{Area}(A_i) \to 0$ as m increases.

Consider the hypercube in $\mathbb{R}^k$ with vertices $\{(\pm\frac{1}{\sqrt{k}}, \pm\frac{1}{\sqrt{k}}, \ldots, \pm\frac{1}{\sqrt{k}})\}$. This cube is inscribed in the unit sphere. Now, given an arbitrarily large integer t, subdivide each side of the cube into a grid of t^{k-1} equally-sized $(k-1)$-cubes (we'll call them "squares") — this gives $m = 2kt^{k-1}$ squares in all. Project each square onto the unit sphere by drawing rays from the origin through all points of the square and taking the set of points where these rays intersect the sphere (we'll call the projection the "shadow" of the square on the unit sphere). This obviously defines a tessellation of the sphere into a set of non-overlapping shadows. Also, the area of the shadow of a square tends to zero as that of the square itself does, so Property 2 is immediately satisfied. The scheme is illustrated in Figure 5.

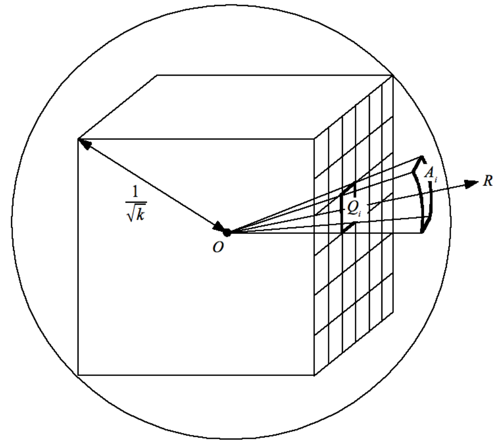

Figure 5: A tessellation scheme for the unit sphere, by radially projecting an uniform tessellation of the inscribed cube onto it.

To prove that Property 1 holds as m increases, observe, first, that it holds for each (unprojected) square. So, given square Q_i, $\text{Area}(Q_i) \geq K\ \text{Diam}(Q_i)^{k-1}$ for some constant K independent of i and m. What can we say about the area and diameter of its shadow A_i? Let $J_S^R(Q)$ denote the orthogonal projection of square Q along ray R on surface S. A_i is the intersection of a "pencil" of rays from the origin through Q_i with the unit sphere, so for any ray R in this pencil, A_i contains $J_{S_{k-1}}^R(Q_i)$ (since the pencil's cross-sectional area increases after passing through the square and hence contains the cylinder with base Q_i and axis R). Also, if T is the tangent to the unit sphere at the point where R intersects it, then $\text{Area}(J_{S_{k-1}}^R(Q_i)) \geq \text{Area}(J_T^R(Q_i))$ (a cylinder subtends its smallest cross-sectional area on an orthogonal plane). Assume, without loss of generality, that Q_i lies on the side of the hypercube $x_k = \frac{1}{\sqrt{k}}$, which has unit normal $\hat{x}_k$. If $\hat{R}$ is the unit vector along R, then $\text{Area}(J_T^R(Q_i)) = |\langle \hat{R}, \hat{x}_k \rangle| \text{Area}(Q_i)$. The inner product is numerically smallest if R passes through a vertex of the cube, so let's assume $\hat{R} = (\frac{1}{\sqrt{k}}, \frac{1}{\sqrt{k}}, \ldots, \frac{1}{\sqrt{k}})$. This gives, in general:

$$
\begin{aligned}
\text{Area}(A_i) &\geq \text{Area}(J_T^R(Q_i)) \\
&\geq \left| \left\langle \left(\frac{1}{\sqrt{k}}, \frac{1}{\sqrt{k}}, \ldots, \frac{1}{\sqrt{k}} \right), \hat{x}_k \right\rangle \right| \text{Area}(Q_i) \\
&= \frac{1}{\sqrt{k}} \text{Area}(Q_i)
\end{aligned}
$$

Now for the diameter. Consider any two distinct points u and v in $\mathbb{R}^k$ other than the origin, projected onto the unit sphere along rays through the origin. The projections are $u' = \frac{u}{\|u\|}$ and $v' = \frac{v}{\|v\|}$. The squared distance between the projections is

$$\begin{aligned}
\|u' - v'\|^2 &= \left\| \frac{u}{\|u\|} - \frac{v}{\|v\|} \right\|^2 \\
&= \frac{\langle u, u\rangle}{\|u\|^2} + \frac{\langle v, v\rangle}{\|v\|^2} - 2\frac{\langle u, v\rangle}{\|u\|\|v\|} \\
&= 2\left(1 - \frac{\langle u, v\rangle}{\|u\|\|v\|}\right) \\
&= \frac{2}{\|u\|\|v\|}(\|u\|\|v\| - \langle u, v\rangle)
\end{aligned}$$

Now the old squared distance was $\|u - v\|^2 = \|u\|^2 + \|v\|^2 - 2\langle u, v\rangle$. Projection changes the squared distance by a factor of

$$\frac{\|u' - v'\|^2}{\|u - v\|^2} = \frac{2}{\|u\|\|v\|}\left(\frac{\|u\|\|v\| - \langle u, v\rangle}{\|u\|^2 + \|v\|^2 - 2\langle u, v\rangle}\right)$$

The numerator and the denominator of the bracketed fraction are both evidently positive (positive multiples of squared distances), and their difference is:

$$\begin{aligned}
&\|u\|^2 + \|v\|^2 - 2\langle u, v\rangle - \|u\|\|v\| + \langle u, v\rangle \\
&\quad = \|u\|^2 + \|v\|^2 - \langle u, v\rangle - \|u\|\|v\| \\
&\quad \geq \|u\|^2 + \|v\|^2 - 2\|u\|\|v\| \qquad \text{(Cauchy-Schwarz)} \\
&\quad = (\|u\| - \|v\|)^2 \geq 0
\end{aligned}$$

Therefore the numerator is less than the denominator and the fraction lies in $[0, 1]$. So we have

$$\|u' - v'\| \leq \left(\frac{2}{\|u\|\|v\|}\right)^{1/2} \|u - v\|$$

Now assume u and v lie on our hypercube. They must both be at least $1/\sqrt{k}$ units from the origin, so we have

$$\|u' - v'\| \leq \left(\frac{2}{1/\sqrt{k}\sqrt{k}}\right)^{1/2} \|u - v\| = \sqrt{2k}\, \|u - v\|$$

So it immediately follows that $\mathrm{Diam}(A_i) \leq \sqrt{2k}\mathrm{Diam}(Q_i)$.

Now, let Q_i and Q_j be any two squares (they may be the same), and A_i and A_j be the corresponding shadows. Q_i and Q_j are of equal size, so they have equal diameter. From the inequalities,

$$\begin{aligned}
\frac{\mathrm{Area}(A_i)}{\mathrm{Diam}(A_j)^{k-1}} &\geq \frac{\frac{1}{\sqrt{k}}\mathrm{Area}(Q_i)}{(\sqrt{2k}\mathrm{Diam}(Q_j))^{k-1}} \\
&= \frac{1}{2^{(k-1)/2}k^{k/2}}\left(\frac{\mathrm{Area}(Q_i)}{\mathrm{Diam}(Q_i)^{k-1}}\right) \\
&\geq \frac{K}{2^{(k-1)/2}k^{k/2}}
\end{aligned}$$

Since this holds for *any* pair of shadows, by setting $C = K/2^{(k-1)/2}k^{k/2}$ we see that Property 1 is satisfied.