When Alpha-Complexes Collapse onto Codimension-1 Submanifolds

Dominique Attali
Université Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, Grenoble, France

Mattéo Clémot
Universite Claude Bernard Lyon 1, CNRS, INSA Lyon, LIRIS, Villeurbanne, France

Bianca B. Dornelas
Institute of Geometry, TU Graz, Austria
Institute for Medical Informatics, Statistics and Documentation, MedUni Graz, Austria

André Lieutier
No affiliation, Aix-en-Provence, France

Abstract

Given a finite set of points P sampling an unknown smooth surface $\mathcal{M} \subseteq \mathbb{R}^3$, our goal is to triangulate $\mathcal{M}$ based solely on P. Assuming $\mathcal{M}$ is a smooth orientable submanifold of codimension 1 in $\mathbb{R}^d$, we introduce a simple algorithm, *Naive Squash*, which simplifies the α-complex of P by repeatedly applying a new type of collapse called *vertical* relative to $\mathcal{M}$. Naive Squash also has a practical version that does not require knowledge of $\mathcal{M}$. We establish conditions under which both the naive and practical Squash algorithms output a triangulation of $\mathcal{M}$. We provide a bound on the angle formed by triangles in the α-complex with $\mathcal{M}$, yielding sampling conditions on P that are competitive with existing literature for smooth surfaces embedded in $\mathbb{R}^3$, while offering a more compartmentalized proof. As a by-product, we obtain that the restricted Delaunay complex of P triangulates $\mathcal{M}$ when $\mathcal{M}$ is a smooth surface in $\mathbb{R}^3$ under weaker conditions than existing ones.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Submanifold reconstruction, triangulation, abstract simplicial complexes, collapses, convexity

Digital Object Identifier 10.4230/LIPIcs.SoCG.2025.11

Related Version *Full Version*: `https://arxiv.org/pdf/2411.10388` [7]

Funding *Bianca B. Dornelas*: Funded by the Austrian Science Fund (FWF), grant W1230.

1 Introduction

Given a finite set of points P that sample an unknown smooth surface $\mathcal{M} \subseteq \mathbb{R}^3$ (example in Figure 1a), we aim to approximate $\mathcal{M}$ based solely on P. This problem, known as *surface reconstruction*, has been widely studied [2,11,14,17,36–38,41]. Several algorithms based on computational geometry have been developed, such as Crust [3], PowerCrust [5], Cocone [4], Wrap [27] and variants based on flow complexes [13,25,34,35]. These algorithms rely on the Delaunay complex of P and offer theoretical guarantees, summarized in [23].

The most desirable guarantee is that the reconstruction outputs a triangulation of $\mathcal{M}$, that is, a simplicial complex whose support is *homeomorphic* to $\mathcal{M}$, in which case we call the algorithm *topologically correct*. That has been established for many of the aforementioned algorithms, assuming that P is noiseless ($P \subseteq \mathcal{M}$) and sufficiently dense. Specifically, let $\mathcal{R} > 0$ be a lower bound on the reach of $\mathcal{M}$, and $\varepsilon \geq 0$ an upper bound on the distance between any point of $\mathcal{M}$ and its nearest point in P. Both Crust and Cocone are topologically correct under the condition $\frac{\varepsilon}{\mathcal{R}} \leq 0.06$ [24], which, to our knowledge, is the weakest such constraint guaranteeing topological correctness for surface reconstruction algorithms in $\mathbb{R}^3$.

Surface reconstruction generalizes to approximating an unknown smooth submanifold $\mathcal{M} \subseteq \mathbb{R}^d$ from a finite sample P. One approach in that case, similar to the Wrap algorithm in $\mathbb{R}^3$, involves *collapses*, which are typically applied to complexes like the *α-complex* [12]. The α-complex of P [28,29,31] includes simplices whose circumspheres have radius $\leq \alpha$ and enclose no other points of P [30]. For well-chosen α, the α-complex has the same homotopy type as $\mathcal{M}$ [8,18,19,40], provided that P is sufficiently dense and has low noise relative to the reach of $\mathcal{M}$. However, it may still fail to capture the topology of $\mathcal{M}$, as illustrated in Figure 1b: for $\mathcal{M} \subseteq \mathbb{R}^3$, the α-complex of P includes *slivers*, tetrahedra that have one dimension more than $\mathcal{M}$, preventing the existence of a homeomorphism. Slivers complicate reconstructing k-dimensional submanifolds in $\mathbb{R}^d$ for $k \geq 2$ for all Delaunay-based reconstruction attempts.

Contributions. We introduce a simple algorithm, `NaiveVerticalSimplification`, which takes as input a simplicial complex K and simplifies it by applying collapses guided by the knowledge of $\mathcal{M}$. We call it naive because this knowledge is non-realistic in practice. We find conditions under which the algorithm is topologically correct for smooth orientable submanifolds $\mathcal{M}$ of $\mathbb{R}^d$ with codimension one. Its variant, `PracticalVerticalSimplification`, does not rely on $\mathcal{M}$ and remains topologically correct, though it requires stricter conditions. When applying both algorithms to the α-complex of P and returning the result, we obtain two reconstruction algorithms which we refer to as `NaiveSquash` and `PracticalSquash`, respectively. We determine conditions on the inputs P and α that guarantee the topological correctness of these squash algorithms. Moreover, for $d = 3$, we show that `PracticalSquash` is correct under the sampling condition $\frac{\varepsilon}{\mathcal{R}} \leq 0.178$ (see Figure 1c for an example output), while `NaiveSquash` is correct for $\frac{\varepsilon}{\mathcal{R}} \leq 0.225$, assuming suitable choice of α. We also show that the *restricted Delaunay complex* [16] is generically homeomorphic to $\mathcal{M}$ when $\frac{\varepsilon}{\mathcal{R}} \leq 0.225$.

In addition, while proving these results, we derive an upper bound for when triangles with vertices on a smooth submanifold $\mathcal{M} \subseteq \mathbb{R}^d$ form a small angle with $\mathcal{M}$: for a triangle abc with $a, b, c \in \mathcal{M}$, longest edge bc, and circumradius ρ, we show that the angle between the affine space spanned by abc and the tangent space to $\mathcal{M}$ at a satisfies:

$$\sin \angle \operatorname{Aff}(abc), \mathbf{T}_a\mathcal{M} \leq \frac{\sqrt{3}\,\rho}{\mathcal{R}}. \tag{1}$$

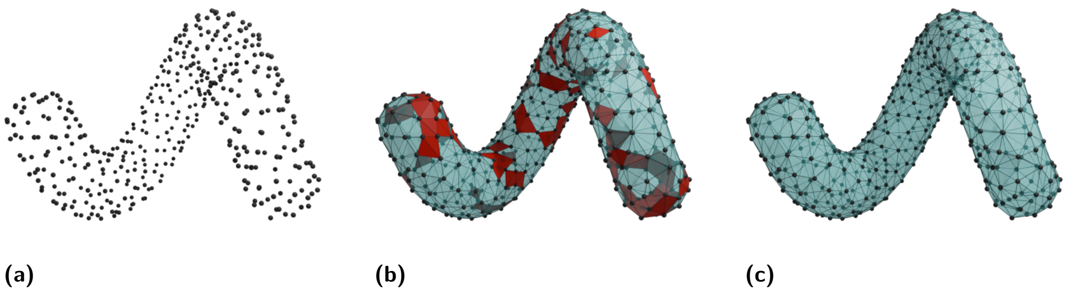

Figure 1 Points sampling a surface in $\mathbb{R}^3$ (a) with the corresponding α-complex, where tetrahedra are highlighted (b). Applying Practical Squash with parameter α outputs (c).

Techniques. Our proof of correctness for the squash algorithms is more compartmentalized than the ones present in the literature: we first consider a smooth orientable submanifold $\mathcal{M} \subseteq \mathbb{R}^d$ with codimension one and a general simplicial complex K embedded in $\mathbb{R}^d$ and contained within a small tubular neighborhood of $\mathcal{M}$. We introduce *vertical* collapses (relative to $\mathcal{M}$) in K which remove d-simplices of K that either have no d-simplices of K above them in directions normal to $\mathcal{M}$ or no d-simplices of K below them in directions normal to $\mathcal{M}$. `NaiveVerticalSimplification` iteratively applies vertical collapses relative to $\mathcal{M}$. `PracticalVerticalSimplification` does not depend on knowledge of $\mathcal{M}$ and applies vertical collapses relative to a hyperplane, constructed dynamically based on the simplex currently considered for collapse.

We examine conditions for the correctness of these algorithms. Apart from the requirement that K has no vertical i-simplices relative to $\mathcal{M}$ for $0 < i < d$ and that its support projects onto $\mathcal{M}$ and fully covers it, we require the *vertical convexity* of K relative to $\mathcal{M}$. This means that each normal line to $\mathcal{M}$ at a point m (restricted to a small ball around m) intersects the support of K in a convex set. For `PracticalVerticalSimplification`, an additional requirement is that the $(d-1)$-simplices of K must form an angle of at most $\frac{\pi}{4}$ with $\mathcal{M}$.

Afterwards, we present `PracticalSquash` and `NaiveSquash`, which initialize the previous algorithms with K as the α-complex of a point set $P \subseteq \mathbb{R}^d$ that samples $\mathcal{M}$. We show that correctness is guaranteed when the i-simplices in the α-complex form small angles with $\mathcal{M}$ for $0 < i < d$. We provide explicit upper bounds for these angles, expressed in terms of ε, δ, and α, where ε and δ control the sample density and noise in P.

We analyze the case $d = 3$ and provide numerical bounds on the ratios $\frac{\varepsilon}{\mathcal{R}}$ and $\frac{\alpha}{\mathcal{R}}$ that ensure the correctness of both squash algorithms. Instrumental to this step, we derive (1) which enables us to upper bound the angles of triangles in the α-complex relative to the manifold and is of independent interest.

Related work. The Squash algorithms (both practical and naive) are similar to Wrap [12,31] in that they compute a subcomplex K of the Delaunay complex of P and then perform a sequence of collapses. However, the selection of K and the nature of the collapses differ between the two methods: in the naive squash, definitions are relative to $\mathcal{M}$, unlike Wrap, which uses flow lines derived from P to guide the collapsing sequence. This distinction allows us to address the general case first and then focus on the specific case of the α-complex. Moreover, while we guarantee correctness for a larger interval of the ratio $\frac{\varepsilon}{\mathcal{R}}$ compared to previous literature, most existing work addresses non-uniform sampling cases, whereas our work focuses on uniform sampling.

The vertical convexity assumption, crucial for the correctness of our algorithms, has been employed in various forms to establish collapsibility of certain classes of simplicial complexes [1,10,20]. Similarly, bounding the angle between the manifold and the simplices used for reconstruction has been essential in prior work [6,9,21–23]. Our bound (1) remains true when replacing $\mathcal{R}$ with the local feature size of a, as explained in [7, App. A]. The thus modified bound improves upon the known bound [23, Lemma 3.5].

At last, for $d = 3$, we show in the full version [7] that the restricted Delaunay complex is generically homeomorphic to $\mathcal{M}$ for $\frac{\varepsilon}{\mathcal{R}} \leq 0.225$. In contrast, it is proven to be homeomorphic to $\mathcal{M}$ only if $\frac{\varepsilon}{\mathcal{R}} \leq 0.09$ [22, Theorem 13.16], a result based on the Topological Ball Theorem [22, Theorem 13.1]. Our proof bypasses this requirement, relying instead on `NaiveSquash`.

Outline. After the preliminaries in Section 2, Section 3 defines vertically convex simplicial complexes. We introduce the concepts of upper and lower skins for these complexes and prove that both are homeomorphic to their orthogonal projection onto $\mathcal{M}$. Section 4 presents general conditions under which a simplicial complex K can be transformed into a triangulation of $\mathcal{M}$ through either Naive or Practical vertical simplification. Section 5 provides conditions ensuring the topological correctness of both the naive and practical Squash algorithms and the restricted Delaunay complex. All missing proofs can be found in the full version [7].

2 Preliminaries

Subsets and submanifold

Given a subset $X \subseteq \mathbb{R}^d$, we define several important geometric concepts. The convex hull of X is denoted as $\mathrm{conv}(X)$ and the affine space spanned by X as $\mathrm{Aff}(X)$. The interior of X is denoted as X°. The *relative interior* of X, denoted as $\mathrm{relint}(X)$, represents the interior of X within $\mathrm{Aff}(X)$. For any point x and radius r, we denote the closed ball with center x and radius r as $B(x, r)$. The r-offset of X, denoted as $X^{\oplus r}$, is the union of closed balls centered at each point in X with radius r: $X^{\oplus r} = \bigcup_{x \in X} B(x, r)$. The *medial axis* of X, denoted as $\mathrm{axis}(X)$, is the set of points in $\mathbb{R}^d$ that have at least two nearest points in X. The *reach* of X, denoted as $\mathrm{Reach}(X)$, is the infimum of distances between X and $\mathrm{axis}(X)$. Furthermore, we define the projection map $\pi_X : \mathbb{R}^d \setminus \mathrm{axis}(X) \to X$, which associates each point x with its unique closest point in X. This projection map is well-defined on every subset of $\mathbb{R}^d$ that does not intersect $\mathrm{axis}(X)$, particularly on every r-offset of X with $r < \mathrm{Reach}(X)$.

Throughout the paper, we designate $\mathcal{M}$ as a compact C^2 submanifold of $\mathbb{R}^d$ of codimension one, and, therefore, orientable (see e.g. [42]).

Given $m \in \mathcal{M}$, we denote the affine tangent space to $\mathcal{M}$ at m as $\mathbf{T}_m\mathcal{M}$ and the affine normal space as $\mathbf{N}_m\mathcal{M}$. As $\mathcal{M}$ has codimension one, $\mathbf{T}_m\mathcal{M}$ is a hyperplane and $\mathbf{N}_m\mathcal{M}$ is a line. Additionally, since $\mathcal{M}$ is C^2, it has a positive reach [43]. For all real numbers r such that $0 < r < \mathrm{Reach}(\mathcal{M})$, the r-offset of $\mathcal{M}$ can be partitioned into the set of normal segments $\{\mathbf{N}_m\mathcal{M} \cap B(m, r)\}_{m \in \mathcal{M}}$ [26], that is,

$$\mathcal{M}^{\oplus r} = \dot{\bigcup}_{m \in \mathcal{M}} \mathbf{N}_m\mathcal{M} \cap B(m, r).$$

We define $\mathbf{n} : \mathcal{M} \to \mathbb{R}^d$ as a differentiable field of unit normal vectors of $\mathcal{M}$ [26]. We let $\mathcal{R}$ be a finite arbitrary number such that $0 < \mathcal{R} \leq \mathrm{Reach}(\mathcal{M})$, fixed throughout.

Abstract simplicial complexes and collapses

We recall some classical definitions of algebraic topology [29,39]. An *abstract simplicial complex* is a collection K of finite non-empty sets with the property that if σ belongs to K, so does every non-empty subset of σ. Each element σ of K is called an *abstract simplex* and its *dimension* is one less than its cardinality: $\dim \sigma = \mathrm{card}\, \sigma - 1$. A simplex of dimension i is called an i-simplex and the set of i-simplices of K is denoted as $K^{[i]}$. If τ and σ are two simplices such that $\tau \subseteq \sigma$, then τ is called a *face* of σ, and σ is called a *coface* of τ. The $(d-1)$-dimensional faces of σ are the *facets* of σ. The *vertex set* of K is $\mathrm{Vert}\, K = \bigcup_{\sigma \in K} \sigma$. A *subcomplex* L of K is a simplicial complex whose elements belong to K. The *link* of σ in K, denoted $\mathrm{Lk}(\sigma, K)$, is the set of simplices τ in K such that $\tau \cup \sigma \in K$ and $\tau \cap \sigma = \emptyset$. It is a subcomplex of K. The *star* of σ in K, denoted as $\mathrm{St}(\sigma, K)$, is the set of cofaces of σ. The simplicial complex formed by all the faces of σ is the *closure* of σ, $\mathrm{Cl}\, \sigma$.

Consider next an abstract simplex $\sigma \subseteq \mathbb{R}^d$. One can associate it to the geometric simplex $\mathrm{conv}(\sigma) \subseteq \mathbb{R}^d$, called the *support* of σ. In general, $\dim(\mathrm{Aff}(\sigma)) \leq \dim(\sigma)$ and we say that σ is *non-degenerate* whenever $\dim(\mathrm{Aff}(\sigma)) = \dim \sigma$. Given a simplicial complex K with vertices in $\mathbb{R}^d$, we say that K is *canonically embedded* if the following two conditions are satisfied:

1. $\dim \sigma = \dim(\mathrm{Aff}(\sigma))$ for all $\sigma \in K$;
2. $\mathrm{conv}(\alpha \cap \beta) = \mathrm{conv}(\alpha) \cap \mathrm{conv}(\beta)$ for all $\alpha, \beta \in K$.

In this paper we consider exclusively abstract simplicial complexes K with vertex sets in $\mathbb{R}^d$ and which are canonically embedded.

Given such a simplicial complex, its *underlying space* (or *support*) is the point set $|K| = \bigcup_{\sigma \in K} \operatorname{conv}(\sigma)$. If $|K|$ is homeomorphic to $\mathcal{M}$, then K is called a *triangulation* of $\mathcal{M}$ or is said to triangulate $\mathcal{M}$. Since K is canonically embedded, the link of every i-simplex of K falls into one of the following two categories: (1) it is a triangulation of the sphere of dimension $d-i-1$ or (2) it is a proper[1] subcomplex of such a triangulation. The *boundary complex* of a simplicial complex K is the subset of simplices in the second category, denoted ∂K, and it holds that $|\partial K| = \partial|K|$. Simplices in ∂K are referred to as *boundary simplices* of K. Given a set of abstract simplices Σ, if $\sigma \in \Sigma$ has no coface in Σ besides itself, then σ is said to be *inclusion-maximal* in Σ.

Suppose that $\tau \in K$ is a simplex whose star in K has a unique inclusion-maximal element $\sigma \neq \tau$. Then τ is said to be *free* in K. Equivalently, τ is free in K if and only if the link of τ in K is the closure of a simplex. Consequently, free simplices of K are always boundary simplices of K. However, not all boundary simplices of K are necessary free. There are instances where none of them are free, such as the famous example when K triangulates the 2-dimensional subspace of $\mathbb{R}^3$, known as the "house with two rooms". A *collapse* in K is the operation that removes from K a free simplex τ along with all its cofaces. This operation is known to preserve the homotopy-type of $|K|$.

Delaunay complexes, α-complexes, and α-shapes

Consider a finite collection of points $P \subseteq \mathbb{R}^d$. The Voronoi region of $q \in P$ is the collection of points $x \in \mathbb{R}^d$ that are closer to q than to any other points of P:

$$V(q, P) = \{x \in \mathbb{R}^d \mid \|x - q\| \leq \|x - p\|, \text{ for all } p \in P\}.$$

Given a subset $\sigma \subseteq P$, let $V(\sigma, P) = \bigcap_{q \in \sigma} V(q, P)$. The *Delaunay complex* is defined as

$$\operatorname{Del}(P) = \{\sigma \subseteq P \mid \sigma \neq \emptyset \text{ and } V(\sigma, P) \neq \emptyset\}.$$

A simplex $\sigma \in \operatorname{Del}(P)$ is called a *Delaunay simplex* of P and it is *dual* to its corresponding *Voronoi cell* $V(\sigma, P)$. Henceforth, we assume that the set of points P is in *general position*. This means that no $d+2$ points of P lie on the same d-dimensional sphere and no $k+2$ points of P lie on the same k-dimensional flat for $k < d$. In that case, $\operatorname{Del}(P)$ is canonically embedded [33]. For $\alpha \geq 0$, the α-complex of P is the subcomplex of $\operatorname{Del}(P)$ defined by:

$$\operatorname{Del}(P, \alpha) = \{\sigma \subseteq P \mid \sigma \neq \emptyset \text{ and } V(\sigma, P) \cap P^{\oplus\alpha} \neq \emptyset\}.$$

Its underlying space $|\operatorname{Del}(P, \alpha)| = \bigcup_{\sigma \in \operatorname{Del}(P,\alpha)} \operatorname{conv}(\sigma)$ is called the *α-shape* of P. It has the properties: (i) $|\operatorname{Del}(P, \alpha)| \subseteq P^{\oplus\alpha}$ and (ii) $|\operatorname{Del}(P, \alpha)|$ is homotopy equivalent to $P^{\oplus\alpha}$; see [28] for more details.

3 Vertically convex simplicial complexes

In this section, we define the concept of vertical convexity relative to $\mathcal{M}$ for both a set and a simplicial complex. We then study the boundary of a vertically convex simplicial complex K. Specifically, we divide the boundary of its underlying space into an upper and a lower skins,

[1] A *proper* subset A of B is such that $A \neq B$.

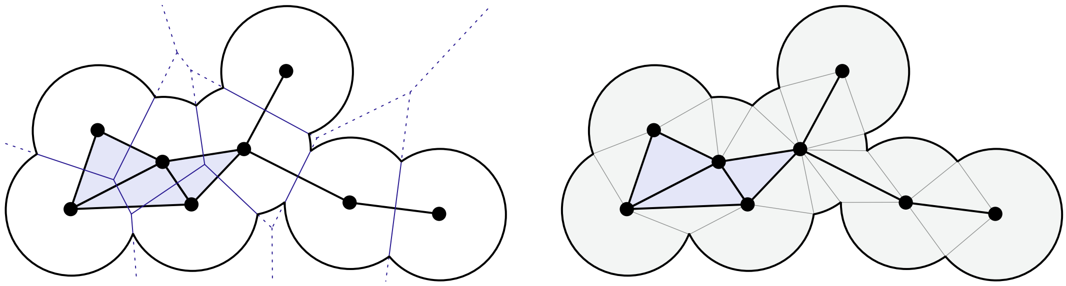

Figure 2 Left: P is such that neither $P^{\oplus\alpha}$ nor $\mathrm{Del}(P,\alpha)$ are vertically convex relative to a horizontal line. Right: Decomposition of $P^{\oplus\alpha} \setminus |\mathrm{Del}(P,\alpha)|^\circ$ in joins as described in [28].

enabling us to identify two boundary subcomplexes: an upper and a lower ones. Furthermore, we show that each of these subcomplexes triangulates the orthogonal projection of $|K|$ onto $\mathcal{M}$ (Lemma 6). We also extend the definitions for a single d-simplex.

▶ **Definition 1** (Vertical convexity). *A set $X \subseteq \mathbb{R}^d$ is* vertically convex *relative to $\mathcal{M}$ if $\exists\, r \in [0, \mathrm{Reach}(\mathcal{M}))$ such that*

1. *$X \subseteq \mathcal{M}^{\oplus r}$ and*
2. *$\forall\, m \in \mathcal{M}$, $\mathbf{N}_m\mathcal{M} \cap B(m,r) \cap X$ is convex.*

In other words, for any $m \in \mathcal{M}$, the set $\mathbf{N}_m\mathcal{M} \cap B(m,r) \cap X$ is either empty or a line segment (possibly of zero-length). A simplicial complex K is vertically convex *relative to $\mathcal{M}$ if its underlying space $|K|$ is.*

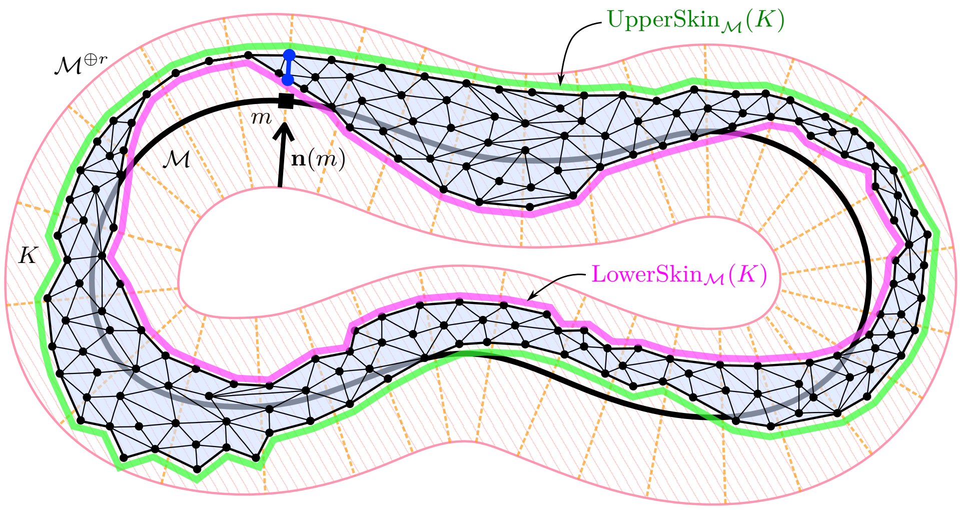

Figure 3 A simplicial complex K vertically convex relative to the curve $\mathcal{M}$. Each segment $\mathbf{N}_m\mathcal{M} \cap B(m,r)$ (in dashed orange) intersects $|K|$ in a line segment, as highlighted (blue) for the point m (represented by a black square). Lemma 4 shows that each of the two skins of $|K|$, depicted in green and pink according to the labeling arrows, is homeomorphic to $\mathcal{M}$.

Examples of a non-vertically convex and a vertically convex simplicial complexes are provided in Figures 2 and 3, respectively.

3.1 Upper and lower skins

Assume that $X \subseteq \mathbb{R}^d$ is vertically convex relative to $\mathcal{M}$ and let $m \in \pi_{\mathcal{M}}(X)$. The endpoints (possibly equal) of the segment $\mathbf{N}_m\mathcal{M} \cap B(m,r) \cap X$ are denoted by $\mathrm{low}_X(m)$ and $\mathrm{up}_X(m)$, with $\mathrm{up}_X(m)$ being above $\mathrm{low}_X(m)$ along the direction of the unit normal vector $\mathbf{n}(m)$. With this notation, X can be expressed as a union of disjoint normal segments:

$$X = \dot{\bigcup}_{m \in \pi_{\mathcal{M}}(X)} [\mathrm{low}_X(m), \mathrm{up}_X(m)].$$

The *upper skin* and *lower skin* of X are, respectively:

$$\mathrm{UpperSkin}_{\mathcal{M}}(X) = \{\mathrm{up}_X(m) \mid m \in \pi_{\mathcal{M}}(X)\},$$
$$\mathrm{LowerSkin}_{\mathcal{M}}(X) = \{\mathrm{low}_X(m) \mid m \in \pi_{\mathcal{M}}(X)\}.$$

Figure 3 displays an example. Our goal is to study the skins of $|K|$, for which we need two extra definitions.

▶ **Definition 2** (Vertical simplex). *A simplex $\sigma \subseteq \mathbb{R}^d$ such that $\mathrm{conv}(\sigma) \subseteq \mathbb{R}^d \setminus \mathrm{axis}(\mathcal{M})$ is* vertical *relative to $\mathcal{M}$ if there exists a pair of distinct points in $\mathrm{conv}(\sigma)$ sharing the same projection onto $\mathcal{M}$.*

▶ **Definition 3** (Non-vertical skeleton). *Assume that $|K| \subseteq \mathbb{R}^d \setminus \mathrm{axis}(\mathcal{M})$. We say that K has a* non-vertical skeleton *relative to $\mathcal{M}$ if K contains no vertical i-simplices relative to $\mathcal{M}$ for all integers $0 < i < d$.*

The next lemma is a key property of vertically convex simplicial complexes:

▶ **Lemma 4.** *Suppose that K is vertically convex and has a non-vertical skeleton relative to $\mathcal{M}$. Then, the upper and lower skins of $|K|$ are closed sets, each homeomorphic to $\pi_{\mathcal{M}}(|K|)$. The homeomorphism is realized in both cases by $\pi_{\mathcal{M}}$. In addition,*

$$\partial|K| = \mathrm{UpperSkin}_{\mathcal{M}}(|K|) \cup \mathrm{LowerSkin}_{\mathcal{M}}(|K|). \tag{2}$$

A simple consequence follows:

▶ **Lemma 5.** *Let K be vertically convex and with non-vertical skeleton relative to $\mathcal{M}$. If $\mathrm{UpperSkin}_{\mathcal{M}}(|K|) = \mathrm{LowerSkin}_{\mathcal{M}}(|K|)$, then $K = \partial K$.*

Aiming for a simplicial version of Equation (2), we define the *upper complex* of K and the *lower complex* of K relative to $\mathcal{M}$ as follows:

$$\mathrm{UpperComplex}_{\mathcal{M}}(K) = \{\nu \subseteq \partial K \mid \mathrm{conv}(\nu) \subseteq \mathrm{UpperSkin}_{\mathcal{M}}(|K|)\},$$
$$\mathrm{LowerComplex}_{\mathcal{M}}(K) = \{\nu \subseteq \partial K \mid \mathrm{conv}(\nu) \subseteq \mathrm{LowerSkin}_{\mathcal{M}}(|K|)\}.$$

By construction, both are subcomplexes of ∂K. A combinatorial equivalent of Lemma 4 is:

▶ **Lemma 6.** *Let K be vertically convex with non-vertical skeleton relative to $\mathcal{M}$. Then,*

$$|\,\mathrm{UpperComplex}_{\mathcal{M}}(K)| = \mathrm{UpperSkin}_{\mathcal{M}}(|K|),$$
$$|\,\mathrm{LowerComplex}_{\mathcal{M}}(K)| = \mathrm{LowerSkin}_{\mathcal{M}}(|K|) \quad \textit{and}$$
$$\partial K = \mathrm{UpperComplex}_{\mathcal{M}}(K) \cup \mathrm{LowerComplex}_{\mathcal{M}}(K).$$

Moreover, if $\pi_{\mathcal{M}}(|K|) = \mathcal{M}$, both $\mathrm{UpperComplex}_{\mathcal{M}}(K)$ and $\mathrm{LowerComplex}_{\mathcal{M}}(K)$ are triangulations of $\mathcal{M}$.

3.2 Upper and lower facets of a d-simplex

Let σ be a non-degenerate d-simplex of $\mathbb{R}^d$ such that $\mathrm{conv}(\sigma) \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$. In that case, $\mathrm{Cl}\,\sigma$ is embedded and vertically convex relative to $\mathcal{M}$. The facets of σ can be partitioned into *upper facets* and *lower facets* of σ relative to $\mathcal{M}$ as follows:

$$\mathrm{UpperFacets}_{\mathcal{M}}(\sigma) = \{\nu \text{ facet of } \sigma \mid \nu \in \mathrm{UpperComplex}_{\mathcal{M}}(\mathrm{Cl}\,\sigma)\}$$
$$\mathrm{LowerFacets}_{\mathcal{M}}(\sigma) = \{\nu \text{ facet of } \sigma \mid \nu \in \mathrm{LowerComplex}_{\mathcal{M}}(\mathrm{Cl}\,\sigma)\}.$$

An example can be seen in Figure 4, where one can also observe the following property:

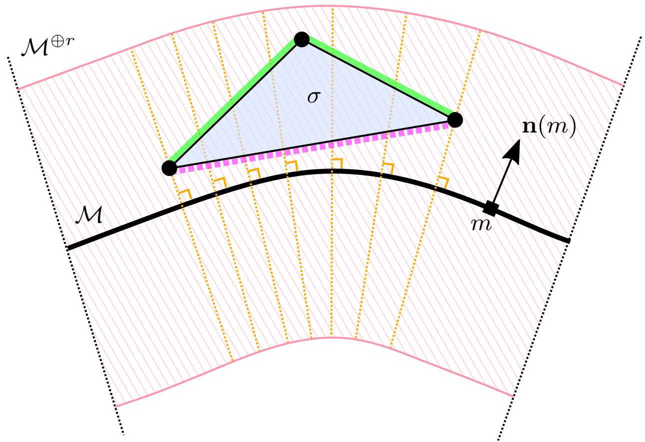

Figure 4 Upper (smooth green edges) and lower (dotted pink edge) facets of a 2-simplex $\sigma \subseteq \mathbb{R}^2$.

▶ **Lemma 7.** *Consider a non-degenerate d-simplex $\sigma \subseteq \mathbb{R}^d$ such that $\mathrm{conv}(\sigma) \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$. If σ has no vertical facets relative to $\mathcal{M}$, then $\mathrm{UpperFacets}_{\mathcal{M}}(\sigma)$ and $\mathrm{LowerFacets}_{\mathcal{M}}(\sigma)$ are non-empty sets that partition the facets of σ.*

4 Vertically collapsing simplicial complexes

In this section, assuming that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$, we introduce an algorithm for simplifying K using vertical collapses relative to $\mathcal{M}$ (Section 4.1) and establish conditions for when it outputs a triangulation of $\mathcal{M}$ (Section 4.2). We first present a naive version that requires the knowledge of $\mathcal{M}$ and then present a practical version (Section 4.3).

4.1 Naive algorithm

▶ **Definition 8** (Vertically free simplices)**.** *A simplex τ is said to be* free from above (resp., free from below) *in K relative to $\mathcal{M}$ if*

- *τ is a free simplex of K;*
- *the unique inclusion-maximal simplex σ in $\mathrm{St}(\tau, K)$ has dimension d;*
- *the set of $(d-1)$-simplices in $\mathrm{St}(\tau, K)$ is exactly the set of upper (resp., lower) facets of σ relative to $\mathcal{M}$.*

We say that τ is vertically free *in K relative to $\mathcal{M}$ if τ is either free from above or from below in K relative to $\mathcal{M}$. See Figures 5 and 7 for a depiction.*

▶ Remark 9. Definition 8 can be naturally extended to non-compact submanifolds $\mathcal{M}$. In particular, it holds for hyperplanes, a fact that we use in Algorithm 2.

▶ **Definition 10** (Vertical collapse). *A* vertical collapse *of K relative to $\mathcal{M}$ is the operation of removing the star of a simplex $\tau \in K$ that is vertically free relative to $\mathcal{M}$.*

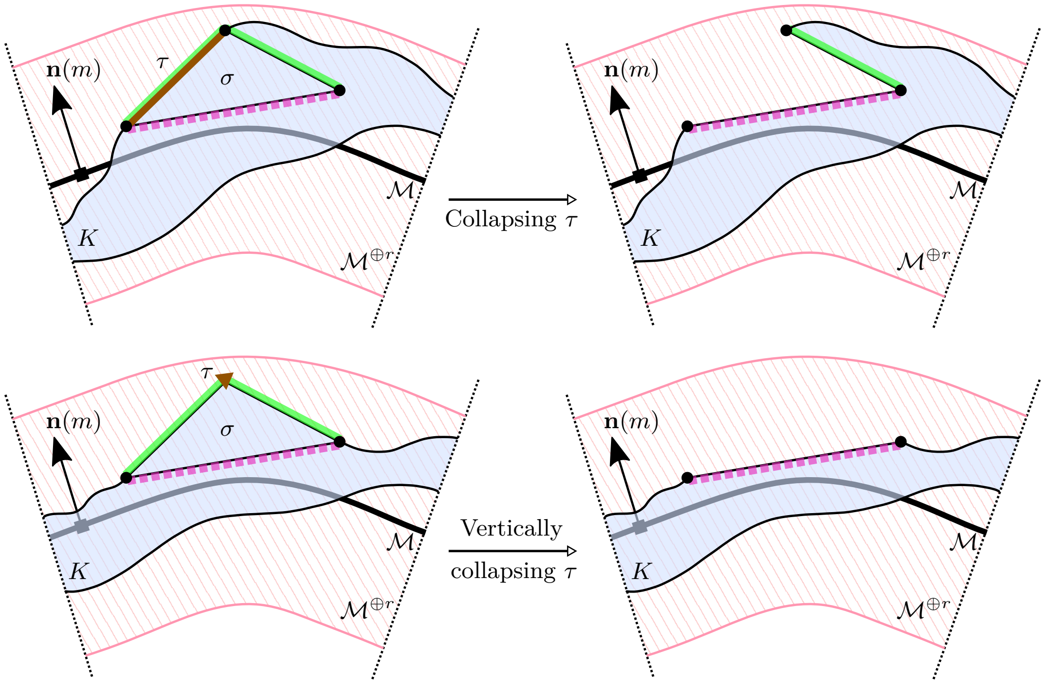

Figure 5 Schematic drawings of K in blue (smooth filled areas). Top row: the edge τ is free but not vertically free relative to $\mathcal{M}$ and collapsing τ does not preserve the vertical convexity of K. Bottom row: the vertex τ is free from above relative to $\mathcal{M}$, so that collapsing τ preserves the vertical convexity of K (Lemma 14). The $(d-1)$-simplices of K that disappear with τ are precisely the upper facets of σ (smooth edges, in green).

A vertical collapse of K can be seen as compressing the underlying space of K by shifting its upper or lower skin along directions normal to $\mathcal{M}$; see Figure 5. Our first algorithm, outlined in Algorithm 1, simplifies K by iteratively applying vertical collapses relative to $\mathcal{M}$. It is worth noting that the algorithm operates on any simplicial complex K with $|K| \subseteq \mathcal{M}^{\oplus r}$ for $r < \mathrm{Reach}(\mathcal{M})$, irrespective of whether K is vertically convex relative to $\mathcal{M}$ or not.

Algorithm 1 `NaiveVerticalSimplification`(K).

while there is a simplex τ vertically free in K relative to $\mathcal{M}$ **do**
  Collapse τ in K;
end while

4.2 Correctness

We now establish conditions under which `NaiveVerticalSimplification`(K) transforms K into a triangulation of $\mathcal{M}$. For that, we introduce a binary relation over d-simplices:

▶ **Definition 11** (Below relation $\prec_{\mathcal{M}}$). *Let $\sigma_0, \sigma_1 \subseteq \mathbb{R}^d$ be two d-simplices sharing a common facet $\nu = \sigma_0 \cap \sigma_1$ and let $\mathrm{conv}(\sigma_0) \cup \mathrm{conv}(\sigma_1) \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$. We say that σ_0 is* below *σ_1 (or that σ_1 is* above *σ_0) relative to $\mathcal{M}$, denoted $\sigma_0 \prec_{\mathcal{M}} \sigma_1$, if ν is an upper facet of σ_0 and a lower facet of σ_1 relative to $\mathcal{M}$.*

Note that the relation $\prec_{\mathcal{M}}$ is not acyclic in general, see Figure 6.

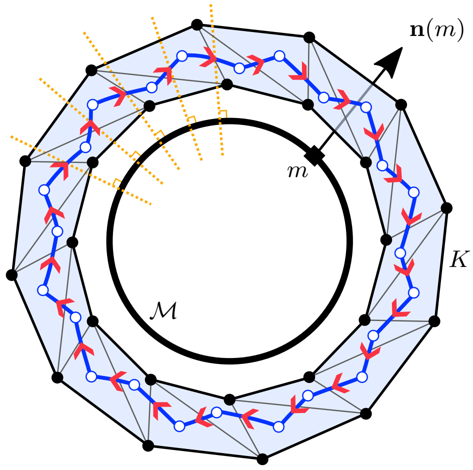

Figure 6 Non-Delaunay triangles that form a cycle in the $\prec_{\mathcal{M}}$ relation and their dual graph.

▶ **Theorem 12** (Correctness). *Consider K such that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$ and assume the following:*
Injective projection: *K has a non-vertical skeleton relative to $\mathcal{M}$.*
Covering projection: *$\pi_{\mathcal{M}}(|K|) = \mathcal{M}$.*
Vertical convexity: *K is vertically convex relative to $\mathcal{M}$.*
Acyclicity: *$\prec_{\mathcal{M}}$ is acyclic over d-simplices of K.*
Then, `NaiveVerticalSimplification`*(K) transforms K into a triangulation of $\mathcal{M}$.*

The remaining of this section aims to prove Theorem 12 and we consider K such that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$. Using the relation $\prec_{\mathcal{M}}$, associate to K its dual graph $G_{\mathcal{M}}(K)$ that has one node for each d-simplex of K and one arc for each pair of d-simplices $\sigma_0, \sigma_1 \in K$ that share a common facet $\sigma_0 \cap \sigma_1$. Direct an arc from σ_0 to σ_1 if $\sigma_0 \prec_{\mathcal{M}} \sigma_1$, and from σ_1 to σ_0 otherwise. Since either $\sigma_0 \prec_{\mathcal{M}} \sigma_1$ or $\sigma_1 \prec_{\mathcal{M}} \sigma_0$, this yields a well-defined orientation for each arc in the dual graph. Figures 6 and 7 show examples.

In a directed graph, a *source* is a node with only outgoing arcs, while a *sink* is a node with only incoming arcs. The next lemma states that a vertical collapse in K corresponds to the removal of either a sink or a source in $G_{\mathcal{M}}(K)$ and conversely. For that, given a finite set of abstract simplices $\Sigma = \{\sigma_1, \sigma_2, \dots, \sigma_k\}$, let $\bigcap \Sigma = \bigcap_{i=1}^{k} \sigma_i$ denote the set of vertices that belong to all simplices in Σ. If $\bigcap \Sigma \neq \emptyset$, it forms an abstract simplex.

▶ **Lemma 13** (Sinks and sources). *Consider K such that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$ and assume that K satisfies the injective projection, covering projection and vertical convexity assumptions of Theorem 12. Consider a d-simplex $\sigma \in K$ and let $\tau = \bigcap \mathrm{UpperFacets}_{\mathcal{M}}(\sigma)$ and $\tau' = \bigcap \mathrm{LowerFacets}_{\mathcal{M}}(\sigma)$. Then,*
- *τ is a free simplex of K from above relative to $\mathcal{M}$ $\iff$ σ is a sink of $G_{\mathcal{M}}(K)$.*
- *τ' is a free simplex of K from below relative to $\mathcal{M}$ $\iff$ σ is a source of $G_{\mathcal{M}}(K)$.*

The next lemma provides an invariant for the while-loop of Algorithm 1.

▶ **Lemma 14** (Loop invariant). *Consider K such that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$. Let τ be a vertically free simplex in K relative to $\mathcal{M}$. Let K' be obtained from K by collapsing τ in K. If K satisfies the assumptions of Theorem 12, then so does K'.*

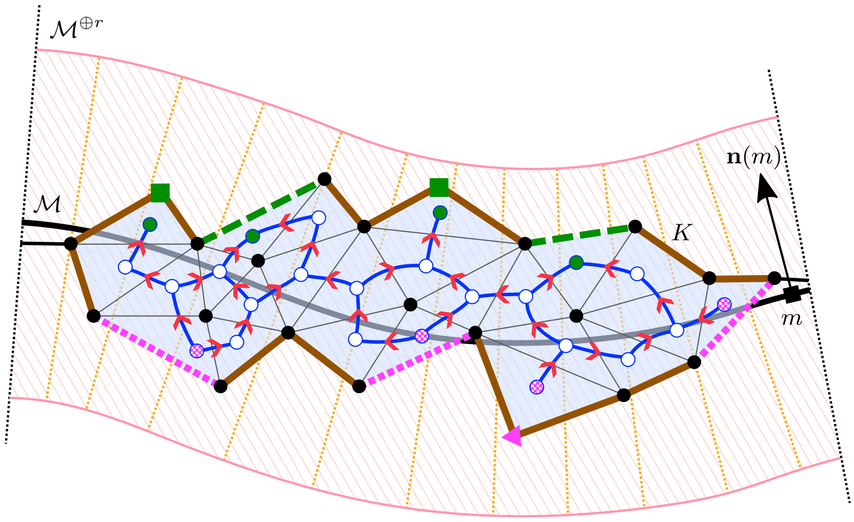

Figure 7 A vertically convex simplicial complex K relative to $\mathcal{M}$. All free simplices are highlighted by thickness. There are four simplices free from below (represented by three dotted edges and one triangular vertex, in pink) and four simplices free from above (represented by two dashed edges and two square vertices, in green). The dual graph (oriented edges, in blue) has four sources (vertices filled by dots, in pink) and four sinks (vertices with a smooth filling, in green), in one-to-one correspondence with the free simplices from below and above.

▶ **Lemma 15** (Upon termination). *Consider K such that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $r < \mathrm{Reach}(\mathcal{M})$ and assume that K satisfies the injective projection and vertical convexity assumptions of Theorem 12. If $G_{\mathcal{M}}(K) = \emptyset$, then $\mathrm{LowerSkin}_{\mathcal{M}}(|K|) = \mathrm{UpperSkin}_{\mathcal{M}}(|K|)$.*

Proof. We establish the contrapositive:

$$\mathrm{LowerSkin}_{\mathcal{M}}(|K|) \neq \mathrm{UpperSkin}_{\mathcal{M}}(|K|) \quad\Longrightarrow\quad G_{\mathcal{M}}(K) \neq \emptyset$$

Suppose that the two skins are distinct, in other words, that there exists $m \in \mathcal{M}$ such that $\mathrm{low}_{|K|}(m) \neq \mathrm{up}_{|K|}(m)$ and let us show that the segment $[\mathrm{low}_{|K|}(m), \mathrm{up}_K(m)]$ intersects the support of at least one d-simplex of K, implying $G_{\mathcal{M}}(K) \neq \emptyset$. Suppose, for a contradiction, that $[\mathrm{low}_{|K|}(m), \mathrm{up}_{|K|}(m)]$ only intersects the support of i-simplices of K for $i < d$. As $\mathrm{low}_{|K|}(m) \neq \mathrm{up}_{|K|}(m)$ and K has a finite number of simplices, at least one of these i-simplices, say ν, intersects $[\mathrm{low}_{|K|}(m), \mathrm{up}_{|K|}(m)]$ in a non-zero length segment containing distinct points $x, y \in \mathrm{conv}(\nu) \cap [\mathrm{low}_{|K|}(m), \mathrm{up}_{|K|}(m)]$. Hence, x and y share the same orthogonal projection m onto $\mathcal{M}$, implying that ν is vertical relative to $\mathcal{M}$. This contradicts the injective projection assumption on K and therefore establishes the contrapositive. ◀

We now prove the correctness of `NaiveVerticalSimplification`(K).

Proof of Theorem 12. The algorithm starts with K that satisfies the theorem assumptions. By Lemma 14, after each iteration of the while-loop we obtain a new K that continues to satisfy those assumptions. Since each iteration involves a vertical collapse of K relative to $\mathcal{M}$, the number of d-simplices of K is reduced. Thus, the algorithm must terminate. Upon termination, there are no vertically free simplices in K relative to $\mathcal{M}$. By Lemma 13, this implies that, when Algorithm 1 terminates, $G_{\mathcal{M}}(K)$ has no terminal node (neither a source nor a sink) and is therefore empty. By Lemma 15, it follows that $\mathrm{LowerSkin}_{\mathcal{M}}(|K|) = \mathrm{UpperSkin}_{\mathcal{M}}(|K|)$. By Lemma 5, we have $K = \partial K$ and Lemma 6 implies

$$K = \partial K = \mathrm{UpperComplex}_{\mathcal{M}}(K) = \mathrm{LowerComplex}_{\mathcal{M}}(K),$$

with K being a triangulation of $\mathcal{M}$. ◀

4.3 Practical version

Algorithm 1 relies on knowledge of $\mathcal{M}$, which renders it impractical for implementation since $\mathcal{M}$ is typically unknown. In this section, we introduce Algorithm 2, a feasible variant that is correct if the $(d-1)$-simplices of K form a sufficiently small angle with $\mathcal{M}$; see [7, App. A] for a definition of the angle between affine spaces. In this variant, we assign an affine space $\mathcal{H}_\tau$ to each $\tau \in K$: for a free simplex τ with a d-dimensional coface σ, $\mathcal{H}_\tau$ is defined as the hyperplane spanned by any facet of σ. Otherwise, set $\mathcal{H}_\tau = \emptyset$. We also use the notion of vertically free simplices relative to $\mathcal{H}_\tau$, extending Definition 8 as indicated in Remark 9.

Algorithm 2 `PracticalVerticalSimplification`(K).

---

while there is a simplex τ vertically free in K relative to $\mathcal{H}_\tau$ **do**
  Collapse τ in K;
end while

▶ **Theorem 16** (Correctness). *Suppose that K satisfies the assumptions of Theorem 12 and, in addition, for all $(d-1)$-simplices ν of K*

$$\max_{a\in\nu} \angle(\mathrm{Aff}(\nu), \mathbf{T}_{\pi_{\mathcal{M}}(a)}\mathcal{M}) < \frac{\pi}{4}. \tag{3}$$

Then, `PracticalVerticalSimplification`(K) *transforms K into a triangulation of $\mathcal{M}$.*

5 Correct reconstructions from α-complexes

In this section, we assume that $\mathcal{M}$ is sampled by a finite point set P and consider Algorithms 3 and 4, which apply vertical collapses to $\mathrm{Del}(P, \alpha)$ either straightforwardly or practically. We introduce two parameters, $\varepsilon \geq 0$ and $\delta \geq 0$, to control the sample density and noise of P, respectively, and a scale parameter $\alpha \geq 0$. Section 5.1 establishes conditions ensuring the correctness of Algorithms 3 and 4, with the proof outlined in Section 5.2. Section 5.3 shows how these conditions hold for a wide range of $\frac{\varepsilon}{\mathcal{R}}$ and $\frac{\alpha}{\mathcal{R}}$ when $\mathcal{M}$ is a surface in $\mathbb{R}^3$ and P is noiseless. This result is extended to the restricted Delaunay complex of P in Section 5.4.

Algorithm 3 `NaiveSquash`(P, α).

$K \leftarrow \mathrm{Del}(P, \alpha)$; `NaiveVerticalSimplification`$(K)$; **return** K;

Algorithm 4 `PracticalSquash`(P, α).

$K \leftarrow \mathrm{Del}(P, \alpha)$; `PracticalVerticalSimplification`$(K)$; **return** K;

5.1 Sampling and angular conditions in $\mathbb{R}^d$

The next definition enables us to express our results in $\mathbb{R}^d$ more concisely.

▶ **Definition 17** (Strict homotopy condition). *We say that $\varepsilon, \delta \geq 0$ satisfy the* strict homotopy condition *if $(\mathcal{R}-\delta)^2 - \varepsilon^2 > (4\sqrt{2}-5)\mathcal{R}^2$ for $\delta \leq \varepsilon$ and $\varepsilon + \sqrt{2}\delta < (\sqrt{2}-1)\mathcal{R}$ for $\delta \geq \varepsilon$.*

Let $I(\varepsilon, \delta)$ be an interval of α values so that $P^{\oplus\alpha}$ is vertically convex with relation to $\mathcal{M}$. The exact definition can be found in [7, App. D.1]. The fact that this interval guarantees vertical convexity follows from the specialization of Propositions 5 and 7 in [8] to the case where $\mathcal{M}$ has codimension-one:

▶ **Theorem 18** (Specialization of [8]). *Suppose that $\mathcal{M} \subseteq P^{\oplus\varepsilon}$ and $P \subseteq \mathcal{M}^{\oplus\delta}$ with $\varepsilon, \delta \geq 0$ that satisfy the strict homotopy condition. Then, for all $\alpha \in I(\varepsilon, \delta)$, $\pi_{\mathcal{M}}(P^{\oplus\alpha}) = \mathcal{M}$ and $P^{\oplus\alpha}$ is vertically convex relative to $\mathcal{M}$. Thus, $P^{\oplus\alpha}$ has associated upper and lower skins and deformation-retracts onto $\mathcal{M}$ along $\pi_{\mathcal{M}}$. In addition, the two skins partition $\partial P^{\oplus\alpha}$.*

The above concepts can be put together to state our main theorem:

▶ **Theorem 19.** *Assume $\mathcal{M} \subseteq P^{\oplus\varepsilon}$ and $P \subseteq \mathcal{M}^{\oplus\delta}$ for $\varepsilon, \delta \geq 0$ that satisfy the strict homotopy condition. Let $\alpha \in \left[\delta, \frac{2(\mathcal{R}-\delta)}{3}\right) \cap I(\varepsilon, \delta)$ and $\beta > 0$ such that $\mathcal{M}^{\oplus\beta} \subseteq P^{\oplus\alpha}$. Suppose that for all i-simplices $\tau \in \mathrm{Del}(P, \alpha)$, $0 < i < d$ and all $(d-1)$-simplices $\nu \in \mathrm{Del}(P, \alpha)$ it holds:*

$$\max_{x \in \mathrm{conv}(\tau)} \angle(\mathrm{Aff}(\tau), \mathbf{T}_{\pi_{\mathcal{M}}(x)}\mathcal{M}) \;<\; \frac{\pi}{2}, \tag{4}$$

$$\min_{a \in \tau} \angle(\mathrm{Aff}(\tau), \mathbf{T}_{\pi_{\mathcal{M}}(a)}\mathcal{M}) \;<\; \arcsin\left(\frac{(\mathcal{R}+\beta)^2 - (\mathcal{R}+\delta)^2 - \alpha^2}{2(\mathcal{R}+\delta)\alpha}\right) \text{ and} \tag{5}$$

$$\min_{x \in \mathrm{conv}(\nu)} \angle(\mathrm{Aff}(\nu), \mathbf{T}_{\pi_{\mathcal{M}}(x)}\mathcal{M}) \;<\; \frac{\pi}{2} - 2\arcsin\left(\frac{\alpha}{2(\mathcal{R}-\delta-\alpha)}\right). \tag{6}$$

Then, $\mathrm{Del}(P, \alpha)$ satisfies the injective projection, covering projection, vertical convexity and acyclicity assumptions of Theorem 12. Furthermore, both the upper and lower complexes of $\mathrm{Del}(P, \alpha)$ relative to $\mathcal{M}$ are triangulations of $\mathcal{M}$ and `NaiveSquash`(P, α) *returns a triangulation of $\mathcal{M}$.*

One can check that Conditions (4), (5) and (6) are well-defined; see Remark E1 in [7].

▶ **Corollary 20.** *Suppose the assumptions of Theorem 19 are satisfied and furthermore that for all $(d-1)$-simplices $\nu \in \mathrm{Del}(P, \alpha)$,*

$$\max_{a \in \nu} \angle(\mathrm{Aff}(\nu), \mathbf{T}_{\pi_{\mathcal{M}}(a)}\mathcal{M}) \;<\; \frac{\pi}{4}. \tag{7}$$

Then, `PracticalSquash`(P, α) *returns a triangulation of $\mathcal{M}$.*

5.2 Partial proof technique for Theorem 19

In this section, we establish the covering projection and vertical convexity of $\mathrm{Del}(P, \alpha)$, as guaranteed by Theorem 19. The complete proof of that theorem, which is unfortunately too lengthy to include here, can be found in [7, App. E].

▶ **Lemma 21.** *Assume $\mathcal{M} \subseteq P^{\oplus\varepsilon}$ and $P \subseteq \mathcal{M}^{\oplus\delta}$ for $\varepsilon, \delta \geq 0$ that satisfy the strict homotopy condition. Let $\alpha \in [\delta, \mathcal{R} - \delta) \cap I(\varepsilon, \delta)$ and $\beta > 0$ be such that $\mathcal{M}^{\oplus\beta} \subseteq P^{\oplus\alpha}$. Suppose that for all i-simplices $\tau \in \mathrm{Del}(P, \alpha)$ with $0 < i < d$, Conditions (4) and (5) hold. Then, $\pi_{\mathcal{M}}(|\mathrm{Del}(P, \alpha)|) = \mathcal{M}$ and $\mathrm{Del}(P, \alpha)$ is vertically convex relative to $\mathcal{M}$.*

Upper and lower joins. For proving this lemma, we introduce upper and lower joins. Consider $\mathcal{M}$, P, ε, δ, α and β that satisfy the assumptions of Lemma 21 and notice that they also meet the conditions of Theorem 18. Therefore, $P^{\oplus\alpha}$ has associated upper and lower skins and the two skins form a partition of $\partial P^{\oplus\alpha}$. Using that partition, we decompose the set difference $P^{\oplus\alpha} \setminus |\mathrm{Del}(P, \alpha)|^\circ$ into upper and lower joins, slightly adapting what is done in [28]; see Figures 2 and 8. For that, notice that $\partial P^{\oplus\alpha}$ can be decomposed into faces, each face being the restriction of $\partial P^{\oplus\alpha}$ to a Voronoi cell of P. There is a one-to-one correspondence

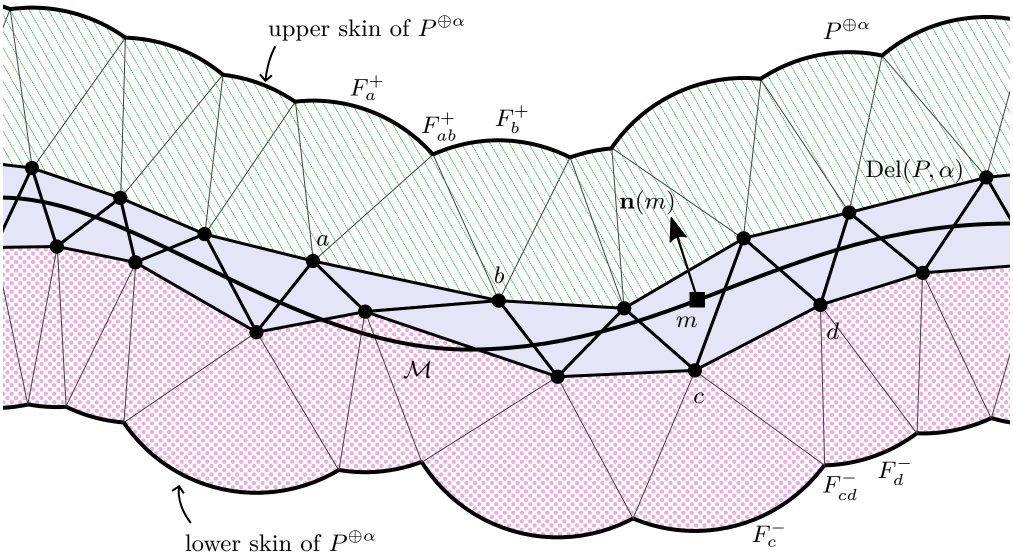

Figure 8 Decomposing $P^{\oplus\alpha} \setminus |\mathrm{Del}(P,\alpha)|^\circ$ into upper joins (hashed, in green) and lower joins (dotted, in pink).

between simplices of $\partial\,\mathrm{Del}(P,\alpha)$ and faces of $\partial P^{\oplus\alpha}$: the simplex $\tau \in \partial\,\mathrm{Del}(P,\alpha)$ corresponds to the face $F_\tau = V(\tau, P) \cap \partial P^{\oplus\alpha}$ and conversely. We can further partition each face F_τ of $\partial P^{\oplus\alpha}$ into a portion F_τ^+ that lies on the upper skin of $\partial P^{\oplus\alpha}$ and a portion F_τ^- that lies on the lower skin of $\partial P^{\oplus\alpha}$. Note that F_τ^+ or F_τ^- can be empty. We refer to F_τ^+ as an upper face and F_τ^- as a lower face. The set of upper faces decompose the upper skin, while the set of lower faces decompose the lower skin. A *join* $X * Y$ is defined as the set of segments $[x, y]$ where $x \in X$ and $y \in Y$ [28]. We call $F_\tau^+ * \mathrm{conv}(\tau)$ an *upper join* and $F_\tau^- * \mathrm{conv}(\tau)$ a *lower join*. The next lemma, proved in [7, App. E.2], identifies points in $\partial|\,\mathrm{Del}(P,\alpha)|$ that are connected to upper or lower joins.

▶ Remark 22. The collection of upper and lower joins cover the set $P^{\oplus\alpha} \setminus |\,\mathrm{Del}(P,\alpha)|^\circ$.

▶ Remark 23. If an upper join and a lower join have a non-empty intersection, the common intersection belongs to $|\,\mathrm{Del}(P,\alpha)|$.

▶ **Lemma 24.** *Under the assumptions of Lemma 21, let $\gamma \in \mathrm{Del}(P,\alpha)$ and $x \in \mathrm{relint}(\mathrm{conv}(\gamma))$. If for some $\lambda > 0$ (*resp. *$\lambda < 0$), the segment $(x, x + \lambda\mathbf{n}(\pi_\mathcal{M}(x))]$ lies outside $|\,\mathrm{Del}(P,\alpha)|$, then it intersects an upper (*resp. *lower join).*

Proof of Lemma 21. Consider $\mathcal{M}$, P, ε, δ, α and β that satisfy the assumptions of Lemma 21. As noted before, they also meet the conditions of Theorem 18. Therefore, $\pi_\mathcal{M}(P^{\oplus\alpha}) = \mathcal{M}$ and $P^{\oplus\alpha}$ is vertically convex relative to $\mathcal{M}$. Hence, there exists $r < \mathrm{Reach}(\mathcal{M})$ such that $P^{\oplus\alpha} \subseteq \mathcal{M}^{\oplus r}$ and $P^{\oplus\alpha} \cap \mathbf{N}_m\mathcal{M} \cap B(m,r)$ is a line segment for any $m \in \mathcal{M}$. Fix $m \in \mathcal{M}$ arbitrarily. We show that $|\,\mathrm{Del}(P,\alpha)| \cap \mathbf{N}_m\mathcal{M} \cap B(m,r)$ is also a line segment.

First, we show by contradiction that it is non-empty. Let u^+ (*resp.* u^-) be the endpoint of the segment $P^{\oplus\alpha} \cap \mathbf{N}_m\mathcal{M} \cap B(m,r)$ that lies on the upper (*resp.* lower) skin of $P^{\oplus\alpha}$ and hence is contained in an upper (*resp.* lower) join. By Remark 22, the entire segment $[u^+, u^-]$ is covered by upper and lower joins. Thus, at some point c of $[u^+, u^-]$, an upper join and a lower join intersect. By Remark 23, such an intersection places c on $|\,\mathrm{Del}(P,\alpha)|$ as well.

Second, we show by contradiction that $|\,\mathrm{Del}(P,\alpha)| \cap \mathbf{N}_m\mathcal{M} \cap B(m,r)$ is connected. Suppose that $a, b \in |\,\mathrm{Del}(P,\alpha)| \cap \mathbf{N}_m\mathcal{M} \cap B(m,r)$ are such that $[a,b] \cap |\,\mathrm{Del}(P,\alpha)| = \{a, b\}$ with a being above b along the direction of $\mathbf{n}(m)$. Since $a, b \in |\,\mathrm{Del}(P,\alpha)| \subseteq P^{\oplus\alpha}$ and $P^{\oplus\alpha}$ is

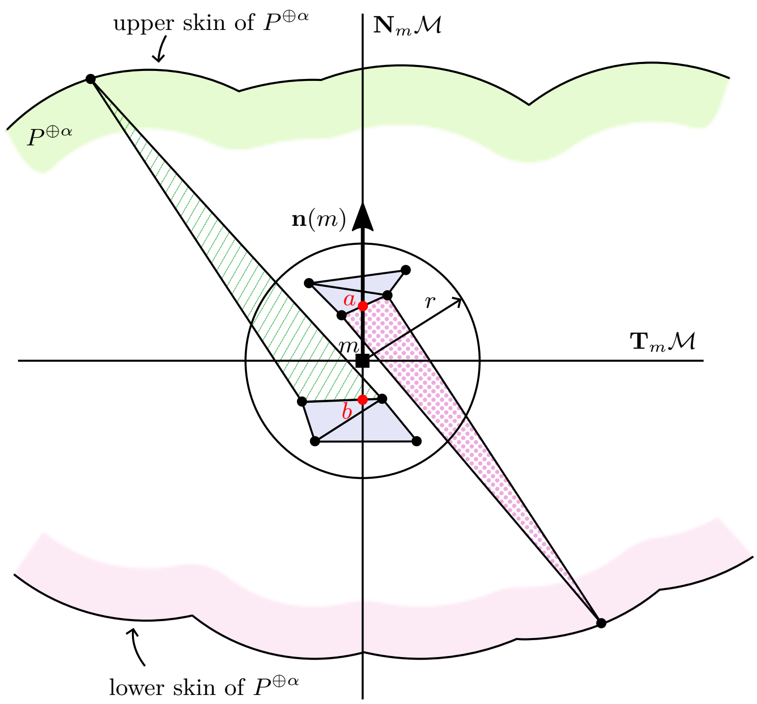

Figure 9 Reaching a contradiction in the proof of Lemma 21. We see the simplices of $\mathrm{Del}(P,\alpha)$ whose support intersects $\mathbf{N}_m\mathcal{M} \cap B(m,r)$ (smooth filling, in pale blue) and one upper (hashed, in green) and one lower (dotted, in pink) joins that intersect $[a,b]$.

vertically convex, the segment $[a,b]$ is contained in $P^{\oplus\alpha}$. By Remark 22, the entire segment $[a,b]$ is covered by upper and lower joins. Let γ_a and γ_b be the simplices of $\mathrm{Del}(P,\alpha)$ that contain a and b, respectively, in their relative interior. Letting $\lambda = \frac{\|a-b\|}{2}$, then both segments $(a, a-\lambda\mathbf{n}(m)]$ and $(b, b+\lambda\mathbf{n}(m)]$ lie outside $|\,\mathrm{Del}(P,\alpha)|$. It follows, by Lemma 24, that there are at least one lower and one upper joins among the joins that cover (a,b); see Figure 9. Hence, an upper and a lower joins intersect at a point c of the segment (a,b). By Remark 23, c lies in $|\,\mathrm{Del}(P,\alpha)|$, a contradiction. Therefore, for all $m \in \mathcal{M}$, $|\,\mathrm{Del}(P,\alpha)| \cap \mathbf{N}_m\mathcal{M} \cap B(m,r)$ is non-empty and connected, thus forming a line segment. ◀

5.3 Sampling conditions for surfaces in $\mathbb{R}^3$

Theorem 19 requires that the i-simplices of $\mathrm{Del}(P,\alpha)$ form a small angle with the manifold $\mathcal{M}$, for $0 < i < d$. Ensuring this can be challenging in practice, especially for $i \geq 3$. However, in the specific case of noiseless edges ($i = 1$) or triangles ($i = 2$), it is possible to upper bound the angle these simplices form with $\mathcal{M}$. For edges, it is known that:

▶ **Lemma 25** ([15, Lemma 7.8])**.** *If ab is a non-degenerate edge with $a, b \in \mathcal{M}$, then*

$$\sin \angle\, \mathrm{Aff}(ab), \mathbf{T}_a\mathcal{M} \leq \frac{\|b-a\|}{2\,\mathcal{R}}.$$

For triangles, let $\rho(\tau)$ be the radius of the smallest $(d-1)$-sphere circumscribing τ. We establish a simple bound that is tighter than the previous one (see [23, Lemma 3.5]):

▶ **Lemma 26.** *If abc is a non-degenerate triangle with longest edge bc, for $a, b, c \in \mathcal{M}$, then*

$$\sin \angle \operatorname{Aff}(abc), \mathbf{T}_a\mathcal{M} \leq \frac{\rho(abc)}{\mathcal{R}} \qquad \textit{if } abc \textit{ is an obtuse triangle and}$$
$$\sin \angle \operatorname{Aff}(abc), \mathbf{T}_a\mathcal{M} \leq \frac{\sqrt{3}\,\rho(abc)}{\mathcal{R}} \qquad \textit{if } abc \textit{ is an acute triangle.}$$

If abc is obtuse, the bound is tight and happens when $\mathcal{M}$ is a sphere of radius $\mathcal{R}$.

The proof is technical and is therefore provided in [7, App. A]. For the same reason, the proof of the following result, where we use the bounds for edges and triangles to establish sampling conditions for surfaces in $\mathbb{R}^3$, is in [7, App. F]. Let us define

$$\beta_{\varepsilon,\alpha} = -\frac{\varepsilon^2}{2\mathcal{R}} + \sqrt{\alpha^2 + \frac{\varepsilon^4}{4\mathcal{R}^2} - \varepsilon^2},$$

which is one particular value of β that guarantees $\mathcal{M}^{\oplus\beta} \subseteq P^{\oplus\alpha}$; see [7, App. D.2].

▶ **Theorem 27.** *Let $\mathcal{M}$ be a C^2 surface in $\mathbb{R}^3$ whose reach is at least $\mathcal{R} > 0$. Let P be a finite point set such that $P \subseteq \mathcal{M} \subseteq P^{\oplus\varepsilon}$.*

1. *For all $\varepsilon, \alpha \geq 0$ that satisfy $\frac{\sqrt{3}\alpha}{\mathcal{R}} < \min\left\{\frac{(\mathcal{R}+\beta_{\varepsilon,\alpha})^2-\mathcal{R}^2-\alpha^2}{2\mathcal{R}\alpha}, \cos\left(2\arcsin\left(\frac{\alpha}{\mathcal{R}}\right)\right)\right\}$,*
   - *the upper and lower complexes of* $\operatorname{Del}(P, \alpha)$ *relative to $\mathcal{M}$ are triangulations of $\mathcal{M}$;*
   - `NaiveSquash`(P, α) *returns a triangulation of $\mathcal{M}$.*
2. *For all $\varepsilon, \alpha \geq 0$ that satisfy in addition $\frac{\sqrt{3}\alpha}{\mathcal{R}} < \sin\left(\frac{\pi}{4} - 2\arcsin\left(\frac{\alpha}{\mathcal{R}}\right)\right)$,*
   - `PracticalSquash`(P, α) *returns a triangulation of $\mathcal{M}$.*

The pairs of $(\frac{\varepsilon}{\mathcal{R}}, \frac{\alpha}{\mathcal{R}})$ that satisfy **1** and **2** are depicted in Figures 10a and 10b, respectively.

▶ Remark 28. In particular, **1** holds for $\frac{\varepsilon}{\mathcal{R}} \leq 0.225$ and $\frac{\alpha}{\mathcal{R}} = 0.359$; and **2** for $\frac{\varepsilon}{\mathcal{R}} \leq 0.178$ and $\frac{\alpha}{\mathcal{R}} = 0.207$, better bounds than the previous existing ones [22, Theorem 13.16][24].

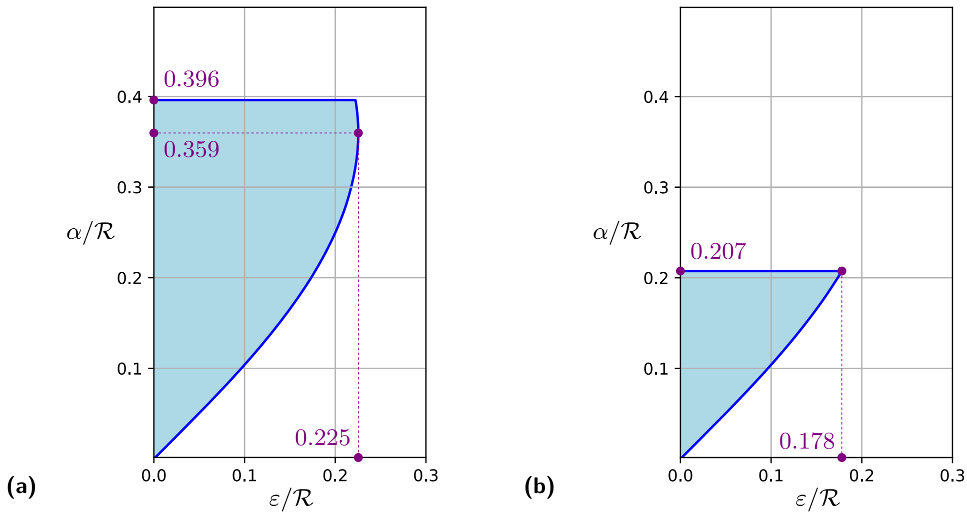

Figure 10 Pairs of $(\frac{\varepsilon}{\mathcal{R}}, \frac{\alpha}{\mathcal{R}})$ for which `NaiveSquash`(P, α) (a) and `PracticalSquash`(P, α) (b) are correct, for $P \subseteq \mathcal{M} \subseteq P^{\oplus\varepsilon}$ and $d = 3$.

5.4 The restricted Delaunay complex

We recall from [32] that the *restricted Delaunay complex* is

$$\mathrm{Del}_{\mathcal{M}}(P) = \{\sigma \subseteq P \mid \sigma \neq \emptyset \text{ and } V(\sigma, P) \cap \mathcal{M} \neq \emptyset\}.$$

▶ **Theorem 29.** *Let $\mathcal{M}$ be a C^2 surface in $\mathbb{R}^3$ whose reach is at least $\mathcal{R} > 0$. Let P be a finite set such that $P \subseteq \mathcal{M} \subseteq P^{\oplus\varepsilon}$ for $0 \leq \frac{\varepsilon}{\mathcal{R}} \leq 0.225$. Under the additional generic assumption that all Voronoi cells of P intersect $\mathcal{M}$ transversally, $\mathrm{Del}_{\mathcal{M}}(P)$ is a triangulation of $\mathcal{M}$.*

References

1 Karim Adiprasito and Bruno Benedetti. Barycentric subdivisions of convex complexes are collapsible. *Discrete & Computational Geometry*, 64(3):608–626, 2020. doi:10.1007/S00454-019-00137-3.

2 Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T Silva. Point set surfaces. In *Proceedings Visualization, 2001. VIS'01.*, pages 21–29. IEEE, 2001. doi:10.1109/VISUAL.2001.964489.

3 N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22(4):481–504, 1999. doi:10.1007/PL00009475.

4 Nina Amenta, Sunghee Choi, Tamal K. Dey, and Naveen Leekha. A simple algorithm for homeomorphic surface reconstruction. *Int. J. Comput. Geom. Appl.*, 12(1-2):125–141, 2002. doi:10.1142/S0218195902000773.

5 Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In David C. Anderson and Kunwoo Lee, editors, *Sixth ACM Symposium on Solid Modeling and Applications, Sheraton Inn, Ann Arbor, Michigan, USA, June 4-8, 2001*, pages 249–266. ACM, 2001. doi:10.1145/376957.376986.

6 D. Attali and A. Lieutier. Flat delaunay complexes for homeomorphic manifold reconstruction, 2022. arXiv:2203.05943.

7 Dominique Attali, Mattéo Clémot, Bianca B. Dornelas, and André Lieutier. When alpha-complexes collapse onto codimension-1 submanifolds, 2024. doi:10.48550/arXiv.2411.10388.

8 Dominique Attali, Hana Dal Poz Kouřimská, Christopher Fillmore, Ishika Ghosh, André Lieutier, Elizabeth Stephenson, and Mathijs Wintraecken. Tight bounds for the learning of homotopy à la Niyogi, Smale, and Weinberger for subsets of Euclidean spaces and of Riemannian manifolds. In *Proc. 26th Ann. Sympos. Comput. Geom.*, Athens, Greece, june 11-14 2024.

9 Dominique Attali and André Lieutier. Delaunay-like triangulation of smooth orientable submanifolds by ℓ_1-norm minimization. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry, SoCG 2022, June 7-10, 2022, Berlin, Germany*, volume 224 of *LIPIcs*, pages 8:1–8:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.SoCG.2022.8.

10 Dominique Attali, André Lieutier, and David Salinas. When convexity helps collapsing complexes. In *SoCG 2019-35th International Symposium on Computational Geometry*, page 15, 2019.

11 Håvard Bakke Bjerkevik. Tighter bounds for reconstruction from ε-samples. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry (SoCG 2022)*, volume 224 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:17, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2022.9.

12 Ulrich Bauer and Herbert Edelsbrunner. The morse theory of čech and delaunay complexes. *Transactions of the American Mathematical Society*, 369(5):3741–3762, 2017.

13 Ulrich Bauer and Fabian Roll. Wrapping cycles in delaunay complexes: Bridging persistent homology and discrete morse theory. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:16, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPICS.SOCG.2024.15.

14 Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics*, 5(4):349–359, 1999. doi:10.1109/2945.817351.

15 Jean-Daniel Boissonnat, Frédéric Chazal, and Mariette Yvinec. *Geometric and topological inference*, volume 57. Cambridge University Press, 2018.

16 Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh, and Steve Oudot. Equating the witness and restricted Delaunay complexes. Research Report CGL-TR-24, CGL, November 2011. URL: https://inria.hal.science/hal-00772486.

17 Jonathan C Carr, Richard K Beatson, Jon B Cherrie, Tim J Mitchell, W Richard Fright, Bruce C McCallum, and Tim R Evans. Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques (SIGGRAPH 2001)*, pages 67–76, 2001. doi:10.1145/383259.383266.

18 F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in Euclidean space. *Discrete and Computational Geometry*, 41(3):461–479, 2009. doi:10.1007/S00454-009-9144-8.

19 F. Chazal and A. Lieutier. Smooth Manifold Reconstruction from Noisy and Non Uniform Approximation with Guarantees. *Computational Geometry: Theory and Applications*, 40:156–170, 2008. doi:10.1016/J.COMGEO.2007.07.001.

20 Aaron Chen, Florian Frick, and Anne Shiu. Neural codes, decidability, and a new local obstruction to convexity. *SIAM Journal on Applied Algebra and Geometry*, 3(1):44–66, 2019. doi:10.1137/18M1186563.

21 Siu-Wing Cheng, Tamal K Dey, and Edgar A Ramos. Manifold reconstruction from point samples. In *SODA*, volume 5, pages 1018–1027, 2005. URL: http://dl.acm.org/citation.cfm?id=1070432.1070579.

22 Siu-Wing Cheng, Tamal Krishna Dey, Jonathan Shewchuk, and Sartaj Sahni. *Delaunay mesh generation*. CRC Press Boca Raton, 2013.

23 Tamal K Dey. *Curve and surface reconstruction: algorithms with mathematical analysis*, volume 23. Cambridge University Press, 2006.

24 Tamal K Dey. Curve and surface reconstruction. In *Handbook of Discrete and Computational Geometry*, pages 915–936. Chapman and Hall/CRC, 2017.

25 Tamal K Dey, Joachim Giesen, Edgar A Ramos, and Bardia Sadri. Critical Points of the Distance to an ε-Sampling of a Surface and Flow-Ccomplex-Based Surface Reconstruction. *International Journal of Computational Geometry & Applications*, 18, 2008.

26 M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, Upper Saddle River, New Jersey, 1976.

27 Herbert Edelsbrunner. Surface reconstruction by wrapping finite sets in space. *Discrete and computational geometry: the Goodman-Pollack Festschrift*, pages 379–404, 2003.

28 Herbert Edelsbrunner. Alpha shapes-a survey. In R. van de Weygaert, G. Vegter, J. Ritzerveld, and V. Icke, editors, *Tessellations in the Sciences: Virtues, Techniques and Applications of Geometric Tilings*. Springer, 2011.

29 Herbert Edelsbrunner and John L. Harer. *Computational topology*. American Mathematical Society, Providence, RI, 2010. An introduction. doi:10.1090/mbk/069.

30 Herbert Edelsbrunner, David G. Kirkpatrick, and Raimund Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inf. Theory*, 29(4):551–558, 1983. doi:10.1109/TIT.1983.1056714.

31 Herbert Edelsbrunner and Ernst P Mücke. Three-dimensional alpha shapes. *ACM Transactions On Graphics (TOG)*, 13(1):43–72, 1994. doi:10.1145/174462.156635.

32 Herbert Edelsbrunner and Nimish R Shah. Triangulating topological spaces. *International Journal of Computational Geometry & Applications*, 7(04):365–378, 1997. doi:10.1142/S0218195997000223.

33 Steven Fortune. Voronoi diagrams and delaunay triangulations. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry, Second Edition*, pages 513–528. Chapman and Hall/CRC, 2004. doi:10.1201/9781420035315.ch23.

34 Joachim Giesen and Matthias John. Surface reconstruction based on a dynamical system. *Computer graphics forum*, 21(3):363–371, 2002. doi:10.1111/1467-8659.00596.

35 Joachim Giesen and Matthias John. The flow complex: a data structure for geometric modeling. In *SODA*, volume 3, pages 285–294, 2003. URL: http://dl.acm.org/citation.cfm?id=644108.644157.

36 Simon Giraudot, David Cohen-Steiner, and Pierre Alliez. Noise-adaptive shape reconstruction from raw point sets. *Computer Graphics Forum*, 32(5):229–238, 2013. doi:10.1111/CGF.12189.

37 Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In *Proceedings of the 19th annual conference on computer graphics and interactive techniques (SIGGRAPH)*, pages 71–78, 1992. URL: https://dl.acm.org/citation.cfm?id=134011.

38 Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. *Computer Graphics Forum*, 7(4), 2006.

39 J.R. Munkres. *Elements of algebraic topology.* Perseus Books, 1993.

40 P. Niyogi, S. Smale, and S. Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete Computational Geometry*, 39(1-3):419–441, 2008. doi:10.1007/S00454-008-9053-2.

41 Stefan Ohrhallinger, Jiju Peethambaran, Amal Dev Parakkat, Tamal K Dey, and Ramanathan Muthuganapathy. 2d points curve reconstruction survey and benchmark. *Computer Graphics Forum*, 40(2):611–632, 2021. doi:10.1111/CGF.142659.

42 Hans Samelson. Orientability of hypersurfaces in r^n. *Proceedings of the American Mathematical Society*, 22(1):301–302, 1969.

43 Sebastian Scholtes. On hypersurfaces of positive reach, alternating Steiner formulae and Hadwiger's problem. *arXiv preprint*, April 2013. arXiv:1304.4179.