Geometry driven collapses for converting a Čech complex into a triangulation of a nicely triangulable shape*

Dominique Attali† André Lieutier‡

Preprint submitted to Discrete & Computational Geometry — November 13, 2015

Abstract

Given a set of points that sample a shape, the Rips complex of the points is often used to provide an approximation of the shape easily-computed. It has been proved that the Rips complex captures the homotopy type of the shape assuming the vertices of the complex meet some mild sampling conditions. Unfortunately, the Rips complex is generally high-dimensional. To remedy this problem, it is tempting to simplify it through a sequence of collapses. Ideally, we would like to end up with a triangulation of the shape. Experiments suggest that, as we simplify the complex by iteratively collapsing faces, it should indeed be possible to avoid entering a dead end such as the famous Bing's house with two rooms. This paper provides a theoretical justification for this empirical observation.

We demonstrate that the Rips complex of a point-cloud (for a well-chosen scale parameter) can always be turned into a simplicial complex homeomorphic to the shape by a sequence of collapses, assuming the shape is nicely triangulable and well-sampled (two concepts we will explain in the paper). To establish our result, we rely on a recent work which gives conditions under which the Rips complex can be converted into a Čech complex by a sequence of collapses. We proceed in two phases. Starting from the Čech complex, we first produce a sequence of collapses that arrives to the Čech complex, restricted by the shape. We then apply a sequence of collapses that transforms the result into the nerve of some covering of the shape. Along the way, we establish results which are of independent interest. First, we show that the reach of a shape can not decrease when intersected with a (possibly infinite) collection of balls, assuming the balls are small enough. Under the same hypotheses, we show that the restriction of a shape with respect to an intersection of balls is either empty or contractible. We also provide conditions under which the nerve of a family of compact sets undergoes collapses as the compact sets evolve over time. We believe conditions are general enough to be useful in other contexts as well.

*Research partially supported by the French "Agence nationale pour la Recherche" under grant ANR-13-BS01-0008 TopData.
†Gipsa-lab – CNRS UMR 5216, Grenoble, France. Dominique.Attali@gipsa-lab.grenoble-inp.fr
‡Dassault systèmes, Aix-en-Provence, France. andre.lieutier@3ds.com

1 Introduction

This paper studies the problem of converting a Čech complex whose vertices sample a shape into a triangulation of that shape using collapses. Even if the present paper focuses exclusively on the Čech Complex, it has also implications on the simplification of Rips complexes by sequences of collapses, due to a recent result in [6].

Imagine we are given a set of points that sample a shape and we want to build an approximation of the shape from the sample points. An often used approach consists in outputting the Vietoris-Rips complex of the points (see for instance [9, 14, 21]). Formally, the *Vietoris-Rips complex* of a set of points P at scale α is the abstract simplicial complex whose simplices are subsets of points in P with diameter at most 2α. For brevity, we shall refer to it as the Rips complex. The Rips complex is an example of a flag complex — the maximal simplicial complex with a given 1-skeleton. As such, it enjoys the property to be completely determined by its 1-skeleton which therefore offers a compact form of storage easy to compute. Moreover, the Rips complex is able to reproduce the homotopy type of the shape in certain situations [17, 18, 10, 4]. Precisely, Hausmann proved in [17] that if the shape A is a compact Riemannian manifold, then the Rips complex with vertex set A is homotopy equivalent to A when the scale used to build the Rips complex is small enough. In [18], Latschev extended this result to Rips complexes with vertex set a metric space (possibly finite) whose Gromov-Hausdorff distance to the shape is small. In [6], a variant has been established in a different framework: shapes are assumed to be subsets of $\mathbb{R}^d$ with a positive μ-reach and Rips complexes are built on finite samples of the shapes using the Euclidean distance. The latter result makes the Rips complex an appealing object for reconstructing shapes living in high dimensional spaces, as for instance in machine learning.

Unfortunately, the dimension of the Rips complex can be very large, compare to the dimension of the underlying shape it is suppose to approximate. This suggests a two-phase algorithm for shape reconstruction. The first phase builds the Rips complex of the data points, thus producing an object with the right homotopy type. The second phase simplifies the Rips complex through a sequence of *collapses*. Ideally, after simplifying the Rips complex by repeatedly applying collapses, we would like to end up with a simplicial complex homeomorphic to the underlying shape.

Yet it is not at all obvious that the Rips complex whose vertices sample a shape contains a subcomplex homeomorphic to that shape. Even if such a subcomplex exists, is there a sequence of collapses that leads to it? Certainly if we want to say anything at all, the geometry of the complex will have to play a key role. As evidence for this, consider a simplicial complex whose vertex set is a noisy point-cloud that samples a 0-dimensional manifold and suppose the complex is composed of a union of Bing's houses with two rooms, one for each connected component in the manifold. Each Bing's house is a 2-dimensional simplicial complex which is contractible but not collapsible. Thus, the complex carries the homotopy type of the 0-dimensional manifold but is not collapsible. Fortunately, it seems that such bad things do not happen in practice, when we start with the Rips complex of a set of points that samples "sufficiently well" a "nice enough" space in $\mathbb{R}^d$. The primary aim of the present work is to understand why. For this, we will focus on the *Čech complex*, a closely related construction. Formally, the Čech complex of a point set P at scale α consists of all simplices spanned by points in P that fit in a ball of radius α. In [6], it was proved that the Rips complex can be reduced to the Čech complex by a sequence of collapses, assuming some sampling conditions are met. This result shows that it suffices to study Čech complexes.

In this paper, we give some mild conditions under which there is a sequence of collapses that converts the Čech complex (and therefore also the Rips complex) into a simplicial complex homeomorphic to the shape. Our result assumes the shape to be a subset of the d-dimensional Euclidean space with the property to be *nicely triangulable*, a concept we will explain later in the paper.

Perhaps unfortunately, our proof that a sequence of collapses exists is not very constructive: it starts by sweeping space with offsets of the shape — which are unknown — and builds a sequence of complexes which have no reason to remain close to flag complexes and therefore cannot benefit from the data structure developed in [5]. Nonetheless, even if results presented here do not give yet any practical algorithm, we believe that they provide a better understanding as to why the Čech complex (and therefore the Rips complex) can be simplified by collapses and how this ability is connected to the underlying metric structure of the space. In the same spirit, we should mention [1], in which the authors prove that every complex that is CAT(0) with a metric for which all vertex stars are convex, is collapsible.

We now list the principal results of the paper, materialized as brown arrows in Figure 1. We also mention some auxiliary results, which are interesting in their own rights. In Section 3, we study how the reach of a shape is modified when intersected with a (possibly infinite) collection of balls and establish the contractibility of the intersection, assuming the balls are small enough. In Section 4, we introduce the *Čech complex restricted by the shape* A and deduce conditions under which it is homotopy equivalent to A (Theorem 1). In Section 5, we provide general conditions under which the nerve of a family of compact sets undergoes collapses as the compact sets evolve over time. Applying this technical result to our context, we obtain conditions under which there is a sequence of collapses that goes from the Čech complex to the restricted Čech complex (Theorem 2). Combined with Theorem 1, this gives an alternative proof to a result [20] recalled in Section 2 (Lemma 2). In Section 6, we define *α-robust coverings* and give conditions under which the restricted Čech complex can be transformed into the nerve of an α-robust covering (Theorem 3). In Section 7, we define and study *nicely triangulable spaces*. Such spaces enjoy the property of having triangulations that can be expressed as the nerve of α-robust coverings for a large range of α. Finally, we provide examples of such spaces. Our list includes affine subspaces of the d-dimensional Euclidean space. It also contains the 2-sphere, the flat torus and all surfaces $C^{1,1}$ diffeomorphic to these two. Although our list is quite short, it is conceivable that many more spaces could be added. Actually, we conjecture that all compact smooth manifolds embedded in $\mathbb{R}^d$ are nicely triangulable and leave open this conjecture for future research. Section 8 concludes the paper.

$$
\begin{tikzcd}
\mathrm{Rips}(P,\alpha) \arrow[r, rightsquigarrow, "{[6]}"] & \mathrm{Cech}(P,\alpha) \arrow[r, rightsquigarrow, "\text{Theorem 2}"] \arrow[d, leftrightarrow, "\simeq"', "\text{Lemma 1 (Nerve Lemma)}"] & \mathrm{Cech}_A(P,\alpha) \arrow[r, rightsquigarrow, "\text{Theorem 3}"] \arrow[d, leftrightarrow, "\simeq\ \text{Theorem 1}"] & \text{Nerve of an } \alpha\text{-robust covering of } A \arrow[dl, leftrightarrow, bend left, "\approx ?"] \\
 & P^{\oplus\alpha} \arrow[r, hookrightarrow, "\text{Lemma 2 [20]}"] & A &
\end{tikzcd}
$$

Figure 1: Logical structure of our results. Brown arrows represent new results. The arrow $\hookrightarrow$ stands for "deformation retracts to". The arrow $\rightsquigarrow$ stands for "can be transformed by a sequence of collapses into". The symbol "$\simeq$" means "homotopy equivalent to" and "$\approx$" means "homeomorphic to".

2 Background

First let us explain some of our terms and introduce the necessary background. $\mathbb{R}^d$ denotes the d-dimensional Euclidean space. $\|x-y\|$ is the Euclidean distance between two points x and y of $\mathbb{R}^d$. The closed ball with center x and radius r is denoted by $\mathrm{B}(x,r)$ and its interior by $\mathrm{B}^\circ(x,r)$. Given a subset $X \subset \mathbb{R}^d$, the

α-offset of X is $X^{\oplus\alpha} = \bigcup_{x\in X} B(x,\alpha)$. The Hausdorff distance $d_H(X,Y)$ between the two compact sets X and Y of $\mathbb{R}^d$ is the smallest real number $\varepsilon \geq 0$ such that $X \subset Y^{\oplus\varepsilon}$ and $Y \subset X^{\oplus\varepsilon}$. We write $d(x,Y) = \inf_{y\in Y} \|x-y\|$ for the distance between point $x \in \mathbb{R}^d$ and the set $Y \subset \mathbb{R}^d$ and $d(X,Y) = \inf_{x\in X} \inf_{y\in Y} \|x-y\|$ for the distance between the two sets $X \subset \mathbb{R}^d$ and $Y \subset \mathbb{R}^d$.

A convenient way to build a simplicial complex is to consider the nerve of a collection of sets. Specifically, let P be a set of indices. Later on, elements of P will be points in $\mathbb{R}^d$. Let $\mathcal{C} = \{C_p \mid p \in P\}$ be a family of sets indexed by $p \in P$. The nerve of the family is the abstract simplicial complex that consists of all non-empty finite subcollections whose sets have a non-empty common intersection. Formally, writing $\operatorname{card}\sigma$ for the number of elements in σ, we have $\operatorname{Nrv}\mathcal{C} = \{\sigma \subset P \mid 0 < \operatorname{card}\sigma < +\infty \text{ and } \bigcap_{p\in\sigma} C_p \neq \emptyset\}$. In this paper, we shall consider nerves of coverings of a shape A. We recall that a *covering* of A is a collection $\mathcal{C} = \{C_p \mid p \in P\}$ of subsets of A so that $A = \bigcup_{p\in P} C_p$. It is a *closed* (*resp. compact*) covering if each set in $\mathcal{C}$ is closed (*resp.* compact). It is a *finite* covering if the set of indices P is finite. The Nerve Lemma gives a condition under which the nerve of a covering of a shape shares the topology of the shape. It has several versions [7] and we shall use the following form:

Lemma 1 (Nerve Lemma). *Consider a compact set $A \subset \mathbb{R}^d$. Let $\mathcal{C} = \{C_p \mid p \in P\}$ be a finite closed covering of A. If for every $\emptyset \neq \sigma \subset P$, the intersection $\bigcap_{z\in\sigma} C_z$ is either empty or contractible, then the underlying space of* $\operatorname{Nrv}\mathcal{C}$ *is homotopy equivalent to A.*

Hereafter, we shall omit the phrase "the underlying space of" and write $X \simeq Y$ to say that X is homotopy equivalent to Y. Given a finite set of points $P \in \mathbb{R}^d$ and a real number $\alpha \geq 0$, the Čech complex of P at scale α can be defined as $\operatorname{Cech}(P,\alpha) = \operatorname{Nrv}\{B(p,\alpha) \mid p \in P\}$. With this definition and the Nerve Lemma, it is clear that $\operatorname{Cech}(P,\alpha) \simeq P^{\oplus\alpha}$; see the black vertical arrow in Figure 1. Several recent results have expressed conditions under which $P^{\oplus\alpha}$ recovers the homotopy type of the shape A [20, 12, 11, 6]. Intuitively, the data points P must sample the shape A sufficiently densely and accurately. One of the simplest ways to measure the quality of the sample is to use the reach of the shape. Given a compact subset A of $\mathbb{R}^d$, recall that the *medial axis* $\mathcal{M}_A$ of A is the set of points in $\mathbb{R}^d$ which have at least two closest points in A. The *reach* of A is the infimum of distances between points in A and points in $\mathcal{M}_A$, $\operatorname{Reach}(A) = \inf_{a\in A, m\in\mathcal{M}_A} \|a-m\|$. It is well-known that a compact subset $C \subset \mathbb{R}^d$ has infinite reach if and only if C is convex. We have (horizontal black arrow in Figure 1):

Lemma 2 ([20]). *Let A and P be two compact subsets of $\mathbb{R}^d$. Suppose there exists a real number ε such that $d_H(A,P) \leq \varepsilon < (3-\sqrt{8})\operatorname{Reach}(A)$. Then, $P^{\oplus\alpha}$ deformation retracts to A for $\alpha = (2+\sqrt{2})\varepsilon$.*

Combining Lemma 1 and Lemma 2, we thus get conditions under which $\operatorname{Cech}(P,\alpha) \simeq A$. The next two sections will provide an alternative proof of this result along the way. We recall a result which will be useful when establishing some of the intermediate geometric lemmas. For a point $x \in A \setminus \mathcal{M}_A$, write $\pi_A(x)$ for the unique point in A closest to x. We have:

Lemma 3 ([15, p. 305]). *Let A be a compact subset of $\mathbb{R}^d$ and c a point such that $0 < d(c,A) < \operatorname{Reach}(A)$. Let Δ_c be the half-line with end point $\pi_A(c)$ and containing c. For every point $x \in \Delta_c$, if $d(x,\pi_A(c)) < \operatorname{Reach}(A)$, then $\pi_A(x) = \pi_A(c)$.*

Before starting the paper, we recall that a *collapse* of an abstract simplicial complex K is the removal of a simplex $\sigma_{\min} \in K$ together with all its cofaces assuming $\sigma_{\min}$ is non-maximal and its set of cofaces contains a unique maximal element $\sigma_{\max} \in K$. A collapse produces a simplicial complex to which K deformation retracts and thus is a simplification operation that preserves the homotopy type [13].

3 Reach of spaces restricted by small balls

In this section, we consider a subset $A \subset \mathbb{R}^d$ such that $\mathrm{Reach}\,(A) > 0$ and prove that as we intersect A with balls of radius $\alpha < \mathrm{Reach}\,(A)$ the reach of the intersection can only get bigger; see Figure 2. More precisely, let A be a compact subset of $\mathbb{R}^d$ and $\sigma \subset \mathbb{R}^d$. Write $\mathcal{B}(\sigma, \alpha) = \bigcap_{z \in \sigma} B(z, \alpha)$ for the common intersection of balls with radius α centered at σ and assume that $A \cap \mathcal{B}(\sigma, \alpha) \neq \emptyset$. In this section, we establish that $\mathrm{Reach}\,(A) \leq \mathrm{Reach}\,(A \cap \mathcal{B}(\sigma, \alpha))$ in the following situations: first when σ is reduced to a single point z (Lemma 5), then when σ is finite (Lemma 6) and finally when σ is a compact subset of $\mathbb{R}^d$ (Lemma 8). Although the first generalization (σ finite) is all we need for establishing Theorem 1 in Section 4, the second generalization (σ compact) will turn out to be useful later on in the paper. Let us start with a preliminary geometric lemma:

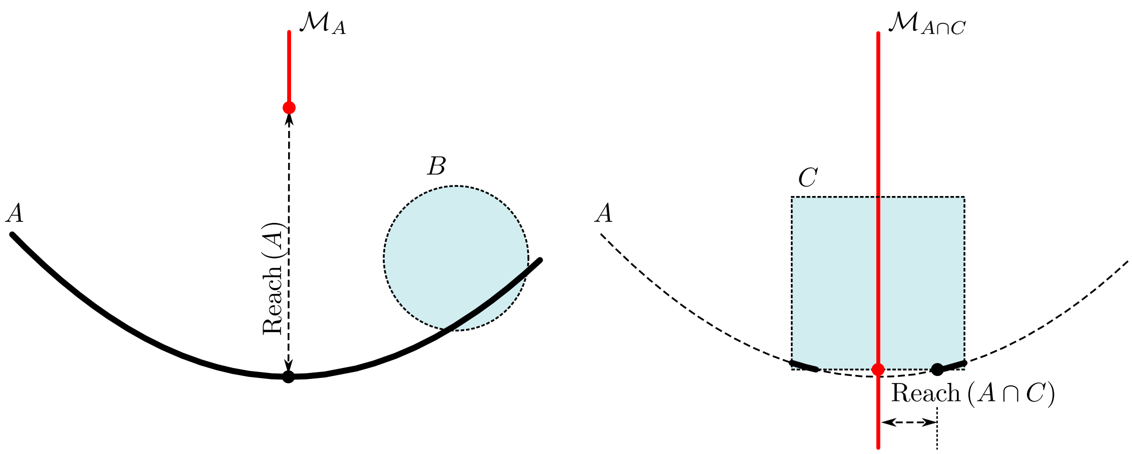

Figure 2: Left: Medial axis $\mathcal{M}_A$ of a shape A. If we intersect A with a ball B whose radius is smaller than the reach of A, the reach of the intersection $A \cap B$ can only get bigger. Right. This property does not hold if we replace the ball B by another set, even with infinite reach such as the solid cube C.

Lemma 4. *Let $X \subset \mathbb{R}^d$ be a non-empty compact set and $B(c, \rho)$ its smallest enclosing ball. For all points $z \in \mathbb{R}^d$ and all real numbers $r \geq \rho$, the following implications hold*

(i) $X \subset \mathrm{B}(z, r) \implies \mathrm{B}(c, r - \sqrt{r^2 - \rho^2}) \subset \mathrm{B}(z, r)$;

(ii) $\mathrm{B}(c, r - \sqrt{r^2 - \rho^2}) \subset \mathrm{B}^\circ(z, r) \implies X \cap \mathrm{B}^\circ(z, r) \neq \emptyset$.

Proof. To establish (i), assume for a contradiction that $\mathrm{B}(z, r)$ does not contain $\mathrm{B}(c, r - \sqrt{r^2 - \rho^2})$ or equivalently that $\|c - z\| > \sqrt{r^2 - \rho^2}$. This implies that the smallest ball enclosing $\mathrm{B}(c, \rho) \cap \mathrm{B}(z, r)$ has radius less than ρ. Since this intersection contains X, this would contradict the definition of ρ as the radius of the smallest ball enclosing X.

It is not hard to check (by contradiction) that the center of the smallest ball enclosing X lies on the convex hull of X. It follows that for any half-space H whose boundary passes through c, the intersection $X \cap H \cap B(c, \rho)$ is non-empty. To establish (ii), we may assume that $c \neq z$ for otherwise the result is clear. Let H be the half-space containing z whose boundary passes through c and is orthogonal to the segment cz; see Figure 3. If $\mathrm{B}(c, r - \sqrt{r^2 - \rho^2}) \subset \mathrm{B}^\circ(z, r)$, then $H \cap B(c, \rho) \subset \mathrm{B}^\circ(z, r)$. Since $X \cap H \cap B(c, \rho) \neq \emptyset$, it follows that $X \cap \mathrm{B}^\circ(z, r) \neq \emptyset$. $\square$

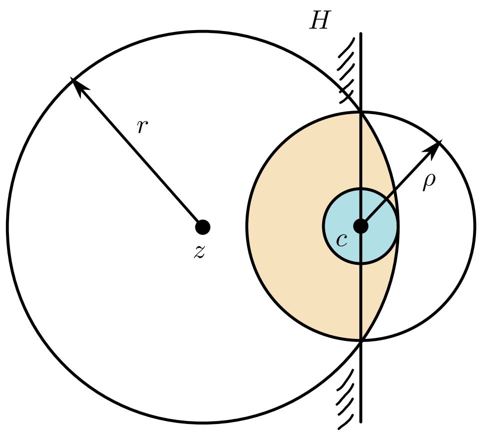

Figure 3: Notation for the proof of Lemma 4 when $\|c - z\| = \sqrt{r^2 - \rho^2}$.

Lemma 5. *Let $A \subset \mathbb{R}^d$ be a compact set and $B(z, \alpha)$ a closed ball with center z and radius α. If $0 \leq \alpha < \operatorname{Reach}(A)$ and $A \cap B(z, \alpha) \neq \emptyset$ then $\operatorname{Reach}(A) \leq \operatorname{Reach}(A \cap B(z, \alpha))$.*

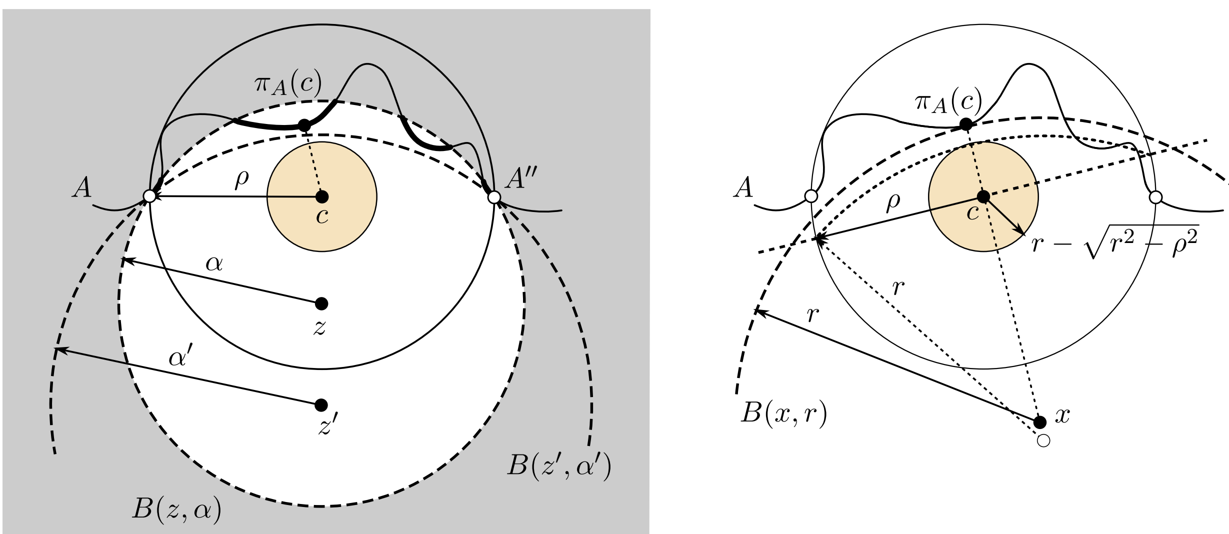

Figure 4: Notation for the proof of Lemma 5. The quantity $r - \sqrt{r^2 - \rho^2}$ represents the height of a spherical cap whose base has radius ρ and which lies on a sphere with radius r.

Proof. See Figure 4 on the left. Assume, by contradiction, that $\operatorname{Reach}(A \cap B(z, \alpha)) < \operatorname{Reach}(A)$ and consider a point z' in the medial axis of $A \cap B(z, \alpha)$ such that

$$d(z', A \cap B(z, \alpha)) = \alpha' < \operatorname{Reach}(A).$$

Introduce $A'' = A \cap B(z, \alpha) \cap B(z', \alpha')$ and denote by c and ρ the center and the radius of the smallest ball enclosing A''. Because A'' is contained in both $B(z, \alpha)$ and $B(z', \alpha')$, the radius of the smallest ball enclosing A'' satisfies $\rho \leq \min\{\alpha, \alpha'\} < \operatorname{Reach}(A)$. Because $A'' \subset A$, we get $d(c, A) \leq d(c, A'') \leq \rho < \operatorname{Reach}(A)$ and therefore c has a unique closest point $\pi_A(c)$ in A. Take r to be any real number such that $\max\{\alpha, \alpha'\} < r < \mathrm{Reach}\,(A)$. We claim that $r - \sqrt{r^2 - \rho^2} < d(c, A)$; see Figure 4 for a geometric interpretation of the quantity $r - \sqrt{r^2 - \rho^2}$. Indeed, for every $(z_0, r_0) \in \{(z, \alpha), (z', \alpha')\}$, since $A'' \subset B(z_0, r_0)$ and $r_0 \geq \rho$, Lemma 4 (i) implies that the following inclusion holds:

$$\mathrm{B}(c, r_0 - \sqrt{r_0^2 - \rho^2}) \;\subset\; \mathrm{B}(z_0, r_0).$$

Since the map $r \mapsto r - \sqrt{r^2 - \rho^2}$ is strictly decreasing in $[\rho, +\infty)$ and $\rho \leq r_0 < r$, we get that

$$\mathrm{B}(c, r - \sqrt{r^2 - \rho^2}) \;\subset\; \mathrm{B}^\circ(z, \alpha) \cap \mathrm{B}^\circ(z', \alpha').$$

By construction, $\mathrm{B}^\circ(z', \alpha')$ contains no points of $A \cap \mathrm{B}^\circ(z, \alpha)$. It follows that $\mathrm{B}^\circ(z, \alpha) \cap \mathrm{B}^\circ(z', \alpha')$ contains no points of A, and neither does $\mathrm{B}(c, r - \sqrt{r^2 - \rho^2})$. Thus, $r - \sqrt{r^2 - \rho^2} < d(c, A) = \|c - \pi_A(c)\|$ as claimed. Let us consider the point $x = \pi_A(c) + \frac{r}{d(c,A)}(c - \pi_A(c))$; see Figure 4 on the right. By construction $\|x - \pi_A(c)\| = r < \mathrm{Reach}\,(A)$ and therefore x has a unique closest point $\pi_A(x)$ in A which, by Lemma 3, satisfies $\pi_A(x) = \pi_A(c)$. Since $\mathrm{B}(c, d(c, A)) \subset \mathrm{B}(x, d(x, A))$, we deduce that

$$\mathrm{B}(c, r - \sqrt{r^2 - \rho^2}) \;\subset\; \mathrm{B}^\circ(x, r).$$

Applying Lemma 4 (ii) we get that $A'' \cap \mathrm{B}^\circ(x, r) \neq \emptyset$ and therefore $\mathrm{B}(x, r)$ contains points of A in its interior. But this contradicts $d(x, A) = r$. □

Lemma 6. *Consider a compact set $A \subset \mathbb{R}^d$ and a finite set $\sigma \subset \mathbb{R}^d$. If $0 \leq \alpha < \mathrm{Reach}\,(A)$ and $A \cap \mathcal{B}(\sigma, \alpha) \neq \emptyset$ then $\mathrm{Reach}\,(A) \leq \mathrm{Reach}\,(A \cap \mathcal{B}(\sigma, \alpha))$.*

Proof. By induction over the size of σ. □

The following lemma is a milestone for the proof of Lemma 8.

Lemma 7. *Let $(A_n)_{n\in\mathbb{N}}$ be a sequence of non-empty compact subsets of $\mathbb{R}^d$ decreasing with respect to the inclusion order. If there exists a real number r such that $0 \leq r \leq \mathrm{Reach}\,(A_n)$ for all $n \in \mathbb{N}$, then $r \leq \mathrm{Reach}\,\left(\bigcap_{n\in\mathbb{N}} A_n\right)$.*

Proof. Letting $A = \bigcap_{n\in\mathbb{N}} A_n$, we first show that the Hausdorff distance $d_H(A_n, A)$ tends to 0 as $n \to +\infty$. For $\varepsilon > 0$, introduce the set $L_\varepsilon = \{x \in \mathbb{R}^d \mid d(x, A) \geq \varepsilon\}$ and notice that $\bigcap_{n\in\mathbb{N}} (L_\varepsilon \cap A_n) = L_\varepsilon \cap A = \emptyset$. Since the sequence of compact sets $(L_\varepsilon \cap A_n)_{n\in\mathbb{N}}$ is decreasing, the only possibility is that $L_\varepsilon \cap A_i = \emptyset$ for some $i \in \mathbb{N}$. Equivalently, $d_H(A_i, A) < \varepsilon$ which proves the convergence of A_n to A under Hausdorff distance.

Let r' be a positive real number in the open interval $(0, r)$ and let z' be a point whose distance to A is r'. Let us prove that z' has a unique closest point in A. For n large enough, $d(z', A_n) < r$ and z' has a unique closest point a_n in A_n. All points a_n are contained in the closed ball $B(z', r)$ and therefore, we can extract from $(a_n)_{n\in\mathbb{N}}$ a subsequence $(a_{n_i})_{i\in\mathbb{N}}$ that converges to a point a. Since $d_H(A_{n_i}, A)$ tends to 0 as $i \to +\infty$, we deduce that the point a must belong to A. Let us define the point z_{n_i} by

$$z_{n_i} = a_{n_i} + \frac{r}{\|z' - a_{n_i}\|}(z' - a_{n_i}).$$

The sequence (z_{n_i}) converges to the point $z = a + \frac{r}{\|z'-a\|}(z' - a)$ and we get $d(z, A) \leq \|z - a\| = r$. By Lemma 3, the point z_{n_i} shares with z' the same closest point in A_{n_i}, namely a_{n_i} and by construction $\|z_{n_i} - a_{n_i}\| = r$. Thus, $\mathrm{B}^\circ(z_{n_i}, r) \cap A = \emptyset$ and by passing to the limit, we get that $\mathrm{B}^\circ(z, r) \cap A = \emptyset$, or equivalently that $d(z, A) \geq r$. Thus, $d(z, A) = r = \|z - a\|$ and since z' lies on the open line segment za, it has a unique closest point in A, namely the point a. Since this is true for all $r' \in (0, r)$, we deduce that $\mathrm{Reach}\,(A) \geq r$. □

Lemma 8. *Let A and $\sigma \neq \emptyset$ be two compact sets of $\mathbb{R}^d$. If $0 \leq \alpha < \text{Reach}(A)$ and $A \cap \mathcal{B}(\sigma, \alpha) \neq \emptyset$ then* $\text{Reach}(A) \leq \text{Reach}(A \cap \mathcal{B}(\sigma, \alpha))$.

Proof. Consider a sequence $\{z_i\}_{i \in \mathbb{N}}$ which is dense in σ. Lemma 6 implies that for all $n \geq 0$, we have $\text{Reach}(A \cap \bigcap_{i=1}^n \mathrm{B}(z_i, \alpha)) \geq \alpha$. Applying Lemma 7 we get that $\text{Reach}\left(A \cap \bigcap_{i \in \mathbb{N}} \mathrm{B}(z_i, \alpha)\right) \geq \alpha$. We claim that $A \cap \bigcap_{i \in \mathbb{N}} \mathrm{B}(z_i, \alpha) = A \cap \bigcap_{z \in \sigma} \mathrm{B}(z, \alpha)$. One direction is trivial. If $a \in A \cap \bigcap_{z \in \sigma} \mathrm{B}(z, \alpha)$, then $a \in A \cap \bigcap_{i \in \mathbb{N}} \mathrm{B}(z_i, \alpha)$. For the other direction, take $a \in A \cap \bigcap_{i \in \mathbb{N}} \mathrm{B}(z_i, \alpha)$ and let us prove that $\forall z \in \sigma, \|z - a\| \leq \alpha$. Assume, by contradiction, that for some $z \in \sigma$, one has $\|z - a\| > \alpha$. Then, there is $i \in \mathbb{N}$ such that $\|z - z_i\| < \|z - a\| - \alpha$, yielding $\|z_i - a\| \geq \|z - a\| - \|z - z_i\| > \alpha$, which is impossible. We have just shown that $a \in A \cap \bigcap_{z \in \sigma} \mathrm{B}(z, \alpha)$. $\square$

4 The restricted Čech complex

Given a subset $A \subset \mathbb{R}^d$, a finite point set P and a real number $\alpha \geq 0$, let us define the Čech complex of P at scale α, $\text{Cech}_A(P, \alpha)$, restricted by A as the set of simplices spanned by points in P that fit in a ball of radius α whose center belongs to A. Equivalently, $\text{Cech}_A(P, \alpha) = \text{Nrv}\{A \cap \mathrm{B}(p, \alpha) \mid p \in P\}$. In this section, we give conditions under which A and $\text{Cech}_A(P, \alpha)$ are homotopy equivalent (Theorem 1). Recall that $\mathcal{B}(\sigma, \alpha) = \bigcap_{z \in \sigma} B(z, \alpha)$ is the common intersection of balls with radius α centered at σ. In the proof, we will argue that $A \cap \mathcal{B}(\sigma, \alpha)$ is either empty or contractible, whenever $0 \leq \alpha < \text{Reach}(A)$. This argument is encapsulated in Lemma 10 and follows from Lemma 9. Let $\text{Radius}(X)$ designate the radius of the smallest ball enclosing the compact set X.

Lemma 9. *If $X \subset \mathbb{R}^d$ is a non-empty compact set with $\text{Radius}(X) < \text{Reach}(X)$, then X is contractible.*

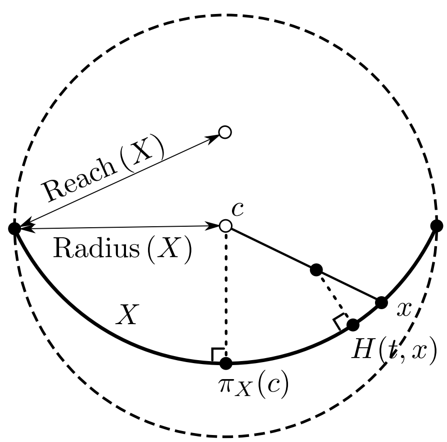

Figure 5: Notation for the proof of Lemma 9.

Proof. We recall that for every point m such that $d(m, X) < \text{Reach}(X)$ there exists a unique point of X closest to m, which we denote by $\pi_X(m)$. Furthermore, we know from [16, page 435] that for $0 < r < \text{Reach}(X)$ the projection map π_X onto X is Lipschitz for points at distance less than r from X. Denote by c the center of the smallest ball enclosing X; see Figure 5. If $x \in X$ and $t \in [0, 1]$, one has

$$d((1-t)x + tc, X) \leq \|(1-t)x + tc - x\| \leq \|c - x\| \leq \text{Radius}(X) < \text{Reach}(X).$$

Thus, the map $H : [0, 1] \times X \to X$ defined by $H(t, x) = \pi_X((1-t)x + tc)$ is Lipschitz and defines a deformation retraction of X onto $\{\pi_X(c)\}$. $\square$

We deduce immediately the following lemma. Besides being useful for proving Theorem 1, it will turn out to be a key tool in Section 6.

Lemma 10. *Let A be a compact set of $\mathbb{R}^d$ and α a real number such $0 \leq \alpha < \mathrm{Reach}\,(A)$. For all non-empty compact subsets $\sigma \subset \mathbb{R}^d$, the intersection $A \cap \mathcal{B}(\sigma, \alpha)$ is either empty or contractible.*

Proof. Suppose $A \cap \mathcal{B}(\sigma, \alpha) \neq \emptyset$. By Lemma 8,

$$\mathrm{Radius}\,(A \cap \mathcal{B}(\sigma, \alpha)) \;\leq\; \alpha \;<\; \mathrm{Reach}\,(A) \;\leq\; \mathrm{Reach}\,(A \cap \mathcal{B}(\sigma, \alpha))\,.$$

By Lemma 9, $A \cap \mathcal{B}(\sigma, \alpha)$ is contractible. □

This lemma can be seen as a variant of Lemma 7 in [2], Proposition 12 in [8] and the local reach lemma in [3] which all say that if A is a k-manifold that intersects a ball B with radius $\alpha < \mathrm{Reach}\,(A)$, then $A \cap B$ is a topological k-ball.

Theorem 1. *Let $A \subset \mathbb{R}^d$ be a compact set, $P \subset \mathbb{R}^d$ a finite point set and α a real number such that $0 \leq \alpha < \mathrm{Reach}\,(A)$ and $A \subset P^{\oplus\alpha}$. Then, $\mathrm{Cech}_A(P, \alpha)$ and A have the same homotopy type.*

Proof. Since $A \subset P^{\oplus\alpha}$, clearly $A = \bigcup_{p \in P}(A \cap B(p, \alpha))$. By Lemma 10, for all $\emptyset \neq \sigma \subset P$, the intersection $\bigcap_{z \in \sigma}(A \cap B(z, \alpha))$ is either empty or contractible. We conclude by applying the Nerve Lemma to the collection $\{A \cap B(p, \alpha) \mid p \in P\}$. □

5 Restricting the Čech complex by collapses

In this section, we state our second theorem (horizontal brown arrow in Figure 1). The theorem describes conditions under which there exists a sequence of collapses that transforms $\mathrm{Cech}(P, \alpha)$ into its restricted version $\mathrm{Cech}_A(P, \alpha)$. It can be seen as a combinatorial version of Lemma 2 which says that, under the same hypotheses, there is a deformation retraction of $P^{\oplus\alpha}$ onto A. Instrumental in proving the theorem, we need several facts about the distance between a collection of balls and a shape A. These facts are formalized in Lemma 11. As before, we let $\mathcal{B}(\sigma, \alpha)$ denote the common intersection of balls with radius α centered at σ and by convention, we set $d(A, \emptyset) = +\infty$. Hence, when we write that $d(A, \mathcal{B}(\sigma, \alpha)) = t$ for some $t \in \mathbb{R}$, this implies implicitly that $\mathcal{B}(\sigma, \alpha) \neq \emptyset$.

Lemma 11. *Let $A \subset \mathbb{R}^d$ be a compact set, $\sigma \subset \mathbb{R}^d$ a finite set and $\alpha \geq 0$ such that $d(A, \mathcal{B}(\sigma, \alpha)) = t$ for some $t \in \mathbb{R}$. If $0 < t < \mathrm{Reach}\,(A) - \alpha$, we have the following properties (see Figure 6, left):*

- *There exists a unique point $x \in \mathcal{B}(\sigma, \alpha)$ whose distance to A is t ;*
- *The set $\sigma_0 = \{p \in \sigma \mid x \in \partial B(p, \alpha)\}$ is non-empty ;*
- *$d(A, \mathcal{B}(\sigma_0, \alpha)) = t$.*

Proof. Since A and $\mathcal{B}(\sigma, \alpha)$ are both compact sets, there exists at least one pair of points $(a, x) \in A \times \mathcal{B}(\sigma, \alpha)$ such that $\|a - x\| = t$; see Figure 6. By definition, x belongs to $B(p, \alpha)$ for all $p \in \sigma$. Since x lies on the boundary of $\mathcal{B}(\sigma, \alpha)$, it lies on the boundary of $\mathrm{B}(p, \alpha)$ for at least one $p \in \sigma$, showing that $\sigma_0 \neq \emptyset$. By construction, a is the point of A closest to x and, by Lemma 3, it is also the point of A closest to point $z = x + \alpha\frac{x-a}{\|x-a\|}$. It follows that $d(A, \{z\}) = \|a - z\| = \|a - x\| + \|x - z\| = t + \alpha$ is realized by the pair of points (a, z) and the distance $d(A, B(z, \alpha)) = t$ is realized by the pair of points (a, x). To prove that x is

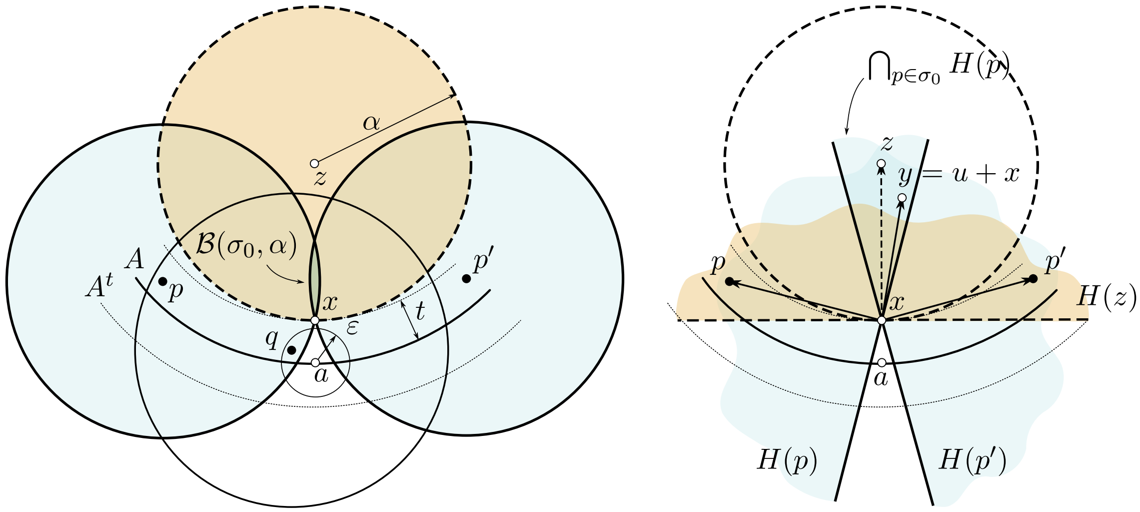

Figure 6: Notation for the proofs of Lemma 11 and Theorem 2. Black dots belong to σ and the ball $B(z,\alpha)$ is instrumental in proving Lemma 11. We show that $\mathcal{B}(\sigma_0,\alpha) \subset B(z,\alpha)$ (on the left) using the fact that $\bigcap_{p\in\sigma_0} H(p) \subset H(z)$ (on the right), where $H(m)$ designates the half-space which contains $B(m, \|x-m\|)$ and whose boundary passes through x.

the unique point of $\mathcal{B}(\sigma,\alpha)$ whose distance to A is t, it suffices to show that $\mathcal{B}(\sigma,\alpha) \subset \mathrm{B}(z,\alpha)$. Actually, we will prove a stronger result, namely that $\mathcal{B}(\sigma_0,\alpha) \subset \mathrm{B}(z,\alpha)$ which will also imply the third item of the lemma, that is, $d(A,\mathcal{B}(\sigma_0,\alpha)) = t$.

Let us associate to every point $m \in \mathbb{R}^d$ the closed half-space $H(m)$ whose boundary passes through x and which contains the ball $B(m, \|m-x\|)$:

$$H(m) = \{y \in \mathbb{R}^d,\ \langle m-x, y-x\rangle \ \geq\ 0\}.$$

We establish the following four statements:

(A) $\bigcap_{p\in\sigma} B(p,\alpha) \subset H(z)$;

(B) $\bigcap_{p\in\sigma} H(p) \subset H(z)$;

(C) $\bigcap_{p\in\sigma_0} H(p) \subset H(z)$;

(D) $\bigcap_{p\in\sigma_0} B(p,\alpha) \subset B(z,\alpha)$;

To establish (A), we note that, by construction, x is the point of $\mathcal{B}(\sigma,\alpha)$ whose distance to a is smallest and therefore $\mathcal{B}(\sigma,\alpha) \cap B(a,t) = \{x\}$ from which statement (A) follows by convexity of $\mathcal{B}(\sigma,\alpha)$. Indeed, if there were $y \in \mathcal{B}(\sigma,\alpha) \setminus H(z)$, then $y \neq x$ and the segment xy would intersect the interior of $B(a,t)$. But, this is impossible since xy is contained in $\mathcal{B}(\sigma,\alpha)$ and $\mathcal{B}(\sigma,\alpha)$ does not intersect the interior of $B(a,t)$. To prove (A) $\Rightarrow$ (B), we apply to the two sets on both sides of (A) an homothety with center x and ratio s. Consider the half-line with origin at x and passing through p and let p_s be the point on this half-line whose distance to x is $s \geq 0$. Clearly, the image of the left side is $\bigcap_{p\in\sigma} B(p_s, s\alpha)$ and the image of the right side is $H(z)$. We thus get that $\bigcap_{p\in\sigma} B(p_s, s\alpha) \subset H(z)$ for all $s \geq 0$. Taking the limit as s tends to infinity (or equivalently, taking the union of left sides for all values of s), we get (B). Statement (B) means that for all $y \in \mathbb{R}^d$, the following implication holds:

$$\min_{p\in\sigma} \langle p-x, y-x\rangle \geq 0 \implies \langle z-x, y-x\rangle \geq 0.$$

Noting that if the above implication holds for all y in a small neighborhood of x, then it holds for all $y \in \mathbb{R}^d$, we deduce that (B) $\Rightarrow$ (C). To prove (C) $\Rightarrow$ (D), we observe that (C) implies that for all $y \in \bigcap_{p\in\sigma_0} H(p)$, the distance between y and the boundary of $H(z)$ is always larger than or equal to the distance between y and the boundary of $H(p)$ for some $p \in \sigma_0$. Formally, this means that for all $u \in \mathbb{R}^d$ and all $\delta \geq 0$, the following implication holds:

$$\min_{p\in\sigma_0} \langle p - x, u\rangle \geq \delta \implies \langle z - x, u\rangle \geq \delta$$

Plugging $\delta = \frac{\|u\|^2}{2}$ in the above implication and noting that for all $m \in \sigma_0 \cup \{z\}$, the following inequality $2\langle m - x, u\rangle \geq \|u\|^2$ can be rewritten as $\|(m - x) - u\|^2 \leq \alpha^2$, we get that for all $u \in \mathbb{R}^d$, the following implication holds:

$$\max_{p\in\sigma_0} \|p - u - x\|^2 \leq \alpha^2 \implies \|z - u - x\|^2 \leq \alpha^2.$$

Equivalently, (D) holds and $\mathcal{B}(\sigma_0, \alpha) \subset \mathrm{B}(z, \alpha)$, as required. □

Theorem 2. *Let $\varepsilon \geq 0$, $\alpha \geq 0$, and $r \geq 0$. Consider a compact set $A \subset \mathbb{R}^d$ with $\mathrm{Reach}\,(A) \geq r$. Let $P \subset \mathbb{R}^d$ be a finite set such that $d_H(A, P) < \varepsilon$. There exists a sequence of collapses from $\mathrm{Cech}(P, \alpha)$ to $\mathrm{Cech}_A(P, \alpha)$ whenever ε, α and r satisfy the following two conditions:*

(i) $\sqrt{2}\alpha < r - \varepsilon$;

(ii) $r - \sqrt{(r - \varepsilon)^2 - \alpha^2} < \alpha - \varepsilon$.

In particular, for $\varepsilon < (3 - \sqrt{8})r$ and $\alpha = (2 + \sqrt{2})\varepsilon$, conditions (i) *and* (ii) *are fulfilled.*

Proof. Letting $\beta = r - \sqrt{(r - \varepsilon)^2 - \alpha^2}$, we observe that condition (i) implies $\beta < r - \alpha$ and condition (ii) is equivalent to $\beta < \alpha - \varepsilon$. For $t \geq 0$, we define the simplicial complex $K_t = \mathrm{Nrv}\{A^{\oplus t} \cap \mathrm{B}(p, \alpha) \mid p \in P\}$. Notice that $K_0 = \mathrm{Cech}_A(P, \alpha)$ and $K_{+\infty} = \mathrm{Cech}(P, \alpha)$. Using the fact that K_t can equivalently be defined as $K_t = \{\sigma \subset P \mid d(A, \mathcal{B}(\sigma, \alpha)) \leq t\}$, we deduce that, as t continuously decreases from $+\infty$ to 0, the complex K_t can only loose simplices and the set of simplices that disappear at time t is:

$$\Delta_t = \{\sigma \subset P \mid d(A, \mathcal{B}(\sigma, \alpha)) = t\}.$$

Generic case. We first establish the theorem under the following generic condition:

($\star$) For all $s \in \mathbb{R}^+$, the set of simplices Δ_s is either empty or has a unique inclusion-minimal element.

At the end of the proof, we will explain what to do if the above condition is not satisfied. Assuming we are in the generic case, we proceed in two stages:

(a) First, we prove that K_t does not change at all as t decreases continuously from $+\infty$ to β. In other words, $K_t = \mathrm{Cech}(P, t)$ for all $t \geq \beta$. Note that this is equivalent to proving that for all non-empty subsets $\sigma \subset P$ and all $t \geq \beta$,

$$\bigcap_{p\in\sigma} \mathrm{B}(p, \alpha) \neq \emptyset \iff A^{\oplus t} \cap \bigcap_{p\in\sigma} \mathrm{B}(p, \alpha) \neq \emptyset.$$

One direction is trivial: if a point belongs to the intersection on the right, then it belongs to the intersection on the left. If the intersection on the left is non-empty, then it contains the center z of the smallest ball enclosing σ and $\mathrm{Radius}\,(\sigma) \leq \alpha < \sqrt{2}\alpha < r - \varepsilon$. Lemma 14 in [4] states that if a subset σ satisfies the following two conditions: (1) $\sigma \subset A^{\oplus\varepsilon}$ and (2) $\mathrm{Radius}\,(\sigma) < r - \varepsilon$, then $\mathrm{Conv}(\sigma) \subset A^{\oplus t}$ for all $t \geq \beta$.

Since z belongs to $\mathrm{Conv}(\sigma)$, it follows that z also belongs to the t-offset $A^{\oplus t}$ and therefore the intersection on the right is non-empty.

(b) Second, we prove that as t decreases continuously from β to 0, the deletion of simplices Δ_t from K_t is a collapse for all $t \in (0, \beta]$. Suppose $\Delta_t \neq \emptyset$ for some $t \in (0, \beta]$ and let $\sigma_{\min}$ be the unique inclusion-minimal element of Δ_t. Since $\sigma_{\min}$ disappears at time t, so does all its cofaces and it follows that Δ_t is the set of cofaces of $\sigma_{\min}$. Since $0 < t \leq \beta < r - \alpha \leq \mathrm{Reach}\,(A) - \alpha$, Lemma 11 implies that there exists a unique point $x \in \mathcal{B}(\sigma_{\min}, \alpha)$ whose distance to A is t; see Figure 6, left. It is easy to see that Δ_t has a unique inclusion-maximal element $\sigma_{\max} = \{p \in P \mid x \in \mathrm{B}(p, \alpha)\}$. Thus, Δ_t consists of all cofaces of $\sigma_{\min}$ and these cofaces are all faces of $\sigma_{\max}$. To prove that removing Δ_t from K_t is a collapse, it suffices to establish that $\sigma_{\min} \neq \sigma_{\max}$. By Lemma 11, we know that $\sigma_0 = \{p \in \sigma_{\min} \mid x \in \partial B(p, \alpha)\}$ is non-empty and belongs to Δ_t. By the choice of $\sigma_{\min}$ as the minimal element of Δ_t, we have $\sigma_{\min} \subset \sigma_0$ and therefore x lies on the boundary of $B(p, \alpha)$ for all $p \in \sigma_{\min}$. Since $d(x, A) = t \leq \beta < \mathrm{Reach}\,(A)$, there exists a unique point $a \in A$ such that $\|a - x\| = t$. Because $d_H(A, P) < \varepsilon$, we know that there exists a point $q \in P$ such that $\|q - a\| \leq \varepsilon$. Since $\|q - x\| \leq \|q - a\| + \|a - x\| \leq \varepsilon + t \leq \varepsilon + \beta < \alpha$, we get that x lies in the interior of $B(q, \alpha)$. Therefore, q belongs to $\sigma_{\max}$ but not to $\sigma_{\min}$. Hence, $\sigma_{\min} \neq \sigma_{\max}$.

Getting rid of the genericity assumption. We need first some definitions and notations. Given a collection of maps $\xi_p : \mathbb{R}^+ \to \mathbb{R}^+$, one for each $p \in P$, we define the simplicial complex

$$K_t^\xi = \mathrm{Nrv}\{A^{\oplus \xi_p(t)} \cap \mathrm{B}(p, \alpha) \mid p \in P\}.$$

If each ξ_p is an increasing continuous bijection, the simplicial complex K_t^ξ can only loose simplices as t continuously decreases from $+\infty$ to 0. Precisely, the set of simplices that disappear at time t is:

$$\Delta_t^\xi = \{\sigma \subset P \mid d(A, \mathcal{B}(\sigma, \alpha)) = \min_{p \in \sigma} \xi_p(t)\}.$$

Given $\eta > 0$, we say that the map $\alpha : \mathbb{R}^+ \to \mathbb{R}^+$ is a *standard η-perturbation* of the identity map if (1) α is a continuous bijection; (2) $\alpha(0) = 0$; (3) $\lim_{p \to +\infty} \alpha(t) = +\infty$; (4) $t \leq \alpha(t) \leq t + \eta$ for all $t \in \mathbb{R}^+$. One can easily check that the composition of two standard η-perturbations is a standard (2η)-perturbation. Suppose now that each ξ_p is a standard η-perturbation and notice that $K_0^\xi = \mathrm{Cech}_A(P, \alpha)$ and $K_{+\infty}^\xi = \mathrm{Cech}(P, \alpha)$. By slightly adapting the first part of the proof above, it is not difficult to establish that for $\eta > 0$ small enough, the simplicial complex K_t^ξ only undergoes collapses as t continuously decreases from $+\infty$ to 0 under the following generic condition:

($\star^\xi$) For all $s \in \mathbb{R}^+$, the set of simplices Δ_s^ξ is either empty or has a unique inclusion-minimal element.

We start by setting ξ_p to the identity map for all $p \in P$. If the generic condition ($\star^\xi$) is not satisfied, we apply a small perturbation to the maps ξ_p so that after perturbation the generic condition ($\star^\xi$) is satisfied and each ξ_p is a standard η-perturbation. The construction can be made so that $\eta > 0$ is as small as desired and we can apply our previous findings. For this, we proceed as follows. We say that two simplices σ_1 and σ_2 are *in conjunction at time t* if they are both inclusion-minimal elements of Δ_t^ξ for some $t \in \mathbb{R}^+$. We say that t is an *event time* if $\Delta_t^\xi \neq \emptyset$. Consider two simplices that are in conjunction at time t, say σ_1 and σ_2. Suppose $q \in \sigma_1$ and $q \notin \sigma_2$. Consider an increasing continuous bijection $\psi : [0, 1] \to [0, 1]$ that differs from identity only in a small neighborhood of t that does not include any other event times. Furthermore, we choose ψ such that $\psi(t) \geq t$. Replacing ξ_q by $\xi_q \circ \psi$ and leaving unchanged ξ_p for all $p \in P \setminus \{q\}$, we change the time at which σ_1 disappears while keeping unchanged the time at which σ_2 disappears. After this operation, σ_1 and σ_2 are not in conjunction anymore. Furthermore, the operation does not create any new pair of simplices in conjunction. By repeating this operation a finite number of times, we thus get a new collection of maps ξ_p as required. □

6 Collapsing the restricted Čech complex

In this section, we find conditions under which there is a sequence of collapses transforming the restricted Čech complex $\mathrm{Cech}_A(P, \alpha)$ into the nerve of an α-robust covering of A. We define α-robust coverings and state our result in Section 6.2. Our proof technique consists in introducing a family of compact sets $\mathcal{D} = \{D_p(t) \mid (p, t) \in P \times [0, 1]\}$ and monitoring the evolution of its nerve as the parameter t increases continuously from 0 to 1. In Section 6.1, we give some general conditions on $\mathcal{D}$ that guarantee that the simplicial complex $K(t) = \mathrm{Nrv}\{D_p(t) \mid p \in P\}$ only undergoes collapses as t increases from 0 to 1. We believe that these conditions are sufficiently general to be applied to other situations and therefore are interesting in their own right. Armed with this tool, we establish our third result in Section 6.2, that is, we find a family of compact sets $\mathcal{D}$ which enjoys the properties required in Section 6.1 and such that $K(0) = \mathrm{Cech}_A(P, \alpha)$ and $K(1)$ is isomorphic to the nerve of an α-robust covering of A.

6.1 Evolving families of compact sets

In this section, we present a tool that will be useful in the next section for establishing Theorem 3. Consider a covering of a topological space and suppose this covering evolves over time. We state conditions under which the evolution of the nerve of this covering only undergoes collapses. Conditions are formulated in a very general setting. We do not even need to endow the topological space with a metric structure. We only require the topological space to be compact and T_1 separable. Recall that a topological space X is said to be T_1 *separable* if for every pair of distinct points $(a, b) \in X^2$, there exist two open sets U_a and U_b such that $a \in U_a \setminus U_b$ and $b \in U_b \setminus U_a$. For instance, metric spaces are T_1 separable.

In Lemma 12, we will use the notion of connectedness as defined in general topology: a topological space (*resp.* subspace) is *connected* if it cannot be represented as the union of two disjoint non-empty open subsets (*resp.* relatively open subsets). Observe that if X is a topological T_1 space, then for any point $a \in X$, the subspace $X \setminus \{a\}$ is open. It follows that if X is connected and $X \setminus \{a\}$ is non-empty, then $\{a\}$ cannot be open and $\{a\}^\circ = \emptyset$. Indeed, if $\{a\}$ were open, then $X = (X \setminus \{a\}) \cup \{a\}$ would be expressed as the union of two disjoint non-empty open subsets, a contradiction.

Before stating our lemma, let us introduce one additional piece of notation. Given a finite set σ and a map $\phi : \sigma \to [0, 1)$, we write $\phi' \succ \phi$ to designate a map $\phi' : \sigma \to [0, 1]$ such that $\phi'(p) > \phi(p)$ for all $p \in \sigma$. We will say that the map ϕ is constant if $\phi(p) = \phi(q)$ for all $(p, q) \in \sigma^2$.

Lemma 12. *Let A be a compact topological T_1 space and P a finite set. Consider a family of compact subsets of A, $\mathcal{D} = \{D_p(t) \mid (p, t) \in P \times [0, 1]\}$ which satisfies the following five properties:*

(a) *For all $0 \leq t < t' \leq 1$ and all $p \in P$, we have $D_p(t') \subset D_p(t)^\circ$;*

(b) $\bigcup_{p \in P} D_p(1) = A$;

(c) *For all $\emptyset \neq \sigma \subset P$ and all maps $\phi : \sigma \to [0, 1]$, the intersection $\mathcal{D}(\sigma, \phi) = \bigcap_{p \in \sigma} D_p \circ \phi(p)$ is either empty or connected;*

(d) *For all $\emptyset \neq \sigma \subset P$ and all maps $\phi : \sigma \to [0, 1)$, the following implication holds: $\mathcal{D}(\sigma, \phi) \neq \emptyset$ and $\mathcal{D}(\sigma, \phi') = \emptyset$ for all $\phi' \succ \phi$ implies that $\mathcal{D}(\sigma, \phi)$ is reduced to a single point.*

(e) *For all* $0 < \tau \leq 1$ *and all* $p \in P$*, one has* $D_p(\tau) = \bigcap_{t\in[0,\tau)} D_p(t)$

Then, as t *increases continuously from* 0 *to* 1*, the simplicial complex* $K_t = \mathrm{Nrv}\{D_p(t) \mid p \in P\}$ *only undergoes collapses.*

Proof. To prove the lemma, we may assume that A is neither disconnected nor reduced to a single point. Indeed, if A is not connected then condition (c) implies that for each $p \in P$, the subset $D_p(0)$ is contained entirely within one connected component of A and the connected components of A can be considered separately. If A is reduced to a single point, then the result is clear.

Assuming A is neither disconnected nor reduced to a single point, we study the changes that occur in K_t as t increases continuously from 0 to 1. Because of condition (a), some simplices may disappear from K_t but no simplices can ever appear in K_t. Given a simplex $\sigma \in K_0 \setminus K_1$, we call $\tau_\sigma = \sup\{t \in [0,1] \mid \sigma \in K_t\}$ the *death time* of σ and claim that $\sigma \in K_{\tau_\sigma}$. Indeed, if $\tau_\sigma = 0$ then $\sigma \in K_0 = K_{\tau_\sigma}$. Now if $0 < \tau_\sigma \leq 1$, condition (e) gives

$$\bigcap_{p\in\sigma} D_p(\tau_\sigma) \;=\; \bigcap_{p\in\sigma}\;\bigcap_{t\in[0,\tau_\sigma)} D_p(t) \;=\; \bigcap_{t\in[0,\tau_\sigma)}\;\bigcap_{p\in\sigma} D_p(t) \;=\; \bigcap_{n\in\mathbb{N}^*}\;\bigcap_{p\in\sigma} D_p(\tau_\sigma - \tau_\sigma/n).$$

Since the intersection of a sequence of decreasing non-empty compact sets is non-empty, the right-hand side above is non-empty and so is the left-hand side. Hence $\sigma \in K_{\tau_\sigma}$ and since $\sigma \in K_0 \setminus K_1$, one has $0 \leq \tau_\sigma < 1$. In other words, the simplex σ belongs to the complex till its death time and disappears from the complex right after. For $t \in [0,1)$, let Δ_t be the set of simplices with death time t.

Generic case. We first establish the lemma under the following generic condition:

($\star$) For all $s \in [0,1)$, the set of simplices Δ_s is either empty or has a unique inclusion-minimal element.

At the end of the proof, we will explain how to get rid of this genericity assumption. Consider $t \in [0,1)$ and suppose $\Delta_t \neq \emptyset$. We prove that the deletion of simplices Δ_t from K_t is a collapse. Let $\sigma_{\min}$ be the unique inclusion-minimal element of Δ_t. Assuming we are in the generic situation, we do not need anymore conditions (c) and (d) but can replace them with the weaker conditions (c') and (d') obtained by considering constant maps for ϕ and ϕ'. Since $\bigcap_{p\in\sigma_{\min}} D_p(t) \neq \emptyset$ and $\bigcap_{p\in\sigma_{\min}} D_p(t+\eta) = \emptyset$ for all $0 < \eta \leq 1-t$, condition (d') implies that $\bigcap_{p\in\sigma_{\min}} D_p(t) = \{a\}$ for some $a \in A$. It is easy to see that Δ_t has a unique inclusion-maximal element $\sigma_{\max} = \{p \in P \mid a \in D_p(t)\}$. Hence, Δ_t consists of all cofaces of $\sigma_{\min}$ and these cofaces are faces of $\sigma_{\max}$. To prove that removing Δ_t from K_t is a collapse, it suffices to establish that $\sigma_{\min} \neq \sigma_{\max}$. We proceed in two steps:

<u>Step 1:</u> Let us prove that a lies on the boundary of $D_p(t)$ for all $p \in \sigma_{\min}$. For this, we start by proving that a lies on the boundary of at least one $D_p(t)$ for some $p \in \sigma_{\min}$. Suppose for a contradiction that a belongs to the interior of $D_p(t)$ for all $p \in \sigma_{\min}$. This implies that, for all $p \in \sigma_{\min}$, there exists an open neighborhood U_p of a such that $a \in U_p \subset D_p(t)$ and $a \in U = \bigcap_{p\in\sigma_{\min}} U_p \subset \bigcap_{p\in\sigma_{\min}} D_p(t) = \{a\}$. It follows that $U = \{a\}$ and therefore a is an isolated point of A. Since A is assumed to be connected, it entails that $A = \{a\}$. We thus reach a contradiction since we obtain a case we have excluded. Defining $\sigma_0 = \{p \in \sigma_{\min} \mid a \in \partial D_p(t)\}$, we have just proved that $\sigma_0 \neq \emptyset$.

Let us now prove that $\sigma_0 = \sigma_{\min}$. Suppose for a contradiction that σ_0 is a proper subset of $\sigma_{\min}$. As before, we can define an open set U such that $a \in U \subset D_p(t)$ for all $p \in \sigma_{\min} \setminus \sigma_0$. We have

$$a \;\in\; \bigcap_{p\in\sigma_0} D_p(t) \cap U \;\subset\; \bigcap_{p\in\sigma_{\min}} D_p(t) = \{a\}.$$

Setting $X = \bigcap_{p\in\sigma_0} D_p(t)$, we thus have $X \cap U = \{a\}$ which is open in the subspace topology on X. Since A is T_1 separable, the subset $X \setminus \{a\}$ is also open in the subspace topology on X. It follows that $X = \{a\} \cup (X \setminus \{a\})$ is the union of two disjoint relatively open subsets. Since X is connected by condition (c'), one of the two subsets must be empty. The only possibility is that $X \setminus \{a\} = \emptyset$ and $X = \bigcap_{p\in\sigma_0} D_p(t) = \{a\}$.

Because of (a), for all $t < t' \leq 1$, we get that $\bigcap_{p\in\sigma_0} D_p(t') \subset \bigcap_{p\in\sigma_0} D_p(t)^\circ = \{a\}^\circ = \emptyset$. It follows that the death time of σ_0 is t and the minimality of $\sigma_{\min}$ implies that $\sigma_0 = \sigma_{\min}$, yielding a contradiction. Thus, a lies on the boundary of $D_p(t)$ for all $p \in \sigma_{\min}$.

Step 2: Let us prove that $\sigma_{\max} \neq \sigma_{\min}$. By condition (b), we have $a \in A = \bigcup_{p\in P} D_p(1)$ and therefore a belongs to $D_q(1)$ for some $q \in P$. Since $t < 1$, condition (a) implies that $a \in D_q(1) \subset D_q(t)^\circ$ and therefore $q \in \sigma_{\max}$. On the other hand, $a \notin \partial D_q(t)$ and therefore $q \notin \sigma_{\min}$. It follows that $\sigma_{\max} \neq \sigma_{\min}$.

Getting rid of the genericity assumption. If we are not in the generic case, the idea is to apply a small perturbation to the family $\mathcal{D}$ which will leave unchanged K_0 and K_1 and such that after perturbation (1) $\mathcal{D}$ will still satisfy the hypotheses of the lemma; (2) the generic condition ($\star$) will hold. We say that two simplices σ_1 and σ_2 are *in conjunction at time t* if they are both inclusion-minimal elements of Δ_t for some $t \in [a, b)$. We say that t is an *event time* if $\Delta_t \neq \emptyset$. Consider two simplices that are in conjunction at time t, say σ_1 and σ_2. Suppose $q \in \sigma_1$ and $q \notin \sigma_2$. Consider an increasing continuous bijection $\psi : [0, 1] \to [0, 1]$ that differs from identity only in a small neighborhood of t that does not include any other event times. Replacing $D_q(t)$ by $D_q \circ \psi(t)$ and leaving unchanged $D_p(t)$ for all $p \in P \setminus \{q\}$, we change the time at which σ_1 disappears while keeping unchanged the time at which σ_2 disappears. After this operation, σ_1 and σ_2 are not in conjunction anymore. Furthermore, the operation does not create any new pair of simplices in conjunction. By repeating this operation a finite number of times, we thus get a new collection as required. $\square$

Remark. Somewhat surprisingly, condition (c) of Lemma 12 is weaker than the condition required by the Nerve Lemma for guaranteeing that the simplicial complex $K_t = \mathrm{Nrv}\{D_p(t) \mid p \in P\}$ is homotopy equivalent to A at some particular value of $t \in [0, 1]$. In particular, if the Nerve Lemma applies at time $t = 0$, that is, if $\bigcap_{p\in\sigma} D_p(0)$ is either empty or contractible for all $\emptyset \neq \sigma \subset P$ and if furthermore the five conditions of Lemma 12 hold, then K_t will have the right homotopy type for all $t \in [0, 1]$.

6.2 Towards the nerve of α-robust coverings

To state and prove our third theorem, we need some definitions. Given a subset $X \subset \mathbb{R}^d$, we call the intersection of all balls of radius α containing X the *α-hull of X* and denote it by $\mathrm{Hull}_\alpha(X)$. By construction, $\mathrm{Hull}_\alpha(X)$ is convex and $\mathrm{Hull}_{+\infty}(X)$ is the convex hull of X. Setting $\mathrm{Clenchers}_\alpha(X) = \{z \in \mathbb{R}^d \mid X \subset B(z, \alpha)\}$, we have

$$\mathrm{Hull}_\alpha(X) = \bigcap_{z\in\mathrm{Clenchers}_\alpha(X)} B(z, \alpha).$$

Notice that $\mathrm{Clenchers}_\alpha(X)$ is also convex; see Figure 7, left. Indeed, if two balls $B(z_1, \alpha)$ and $B(z_2, \alpha)$ contain X, then any ball $B(\lambda_1 z_1 + \lambda_2 z_2, \alpha)$ with $\lambda_1 + \lambda_2 = 1$, $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$ also contains X. Furthermore, if X is compact, so is $\mathrm{Clenchers}_\alpha(X)$.

Definition 1 (α-robust coverings)**.** *A covering $\mathcal{C} = \{C_v \mid v \in V\}$ of A is* α-robust *if (1) each set in $\mathcal{C}$ can be enclosed in an open ball with radius α; (2) $\mathrm{Nrv}\,\mathcal{C} = \mathrm{Nrv}\{A \cap \mathrm{Hull}_\alpha(C_v) \mid v \in V\}$.*

Of course, one may wonder if α-robust coverings of a shape A often arise in practice. Section 7 will address this issue. For now we focus on establishing properties of α-robust coverings.

Lemma 13. *If $\mathcal{C}$ is a finite compact α-robust covering of A and $0 \leq \alpha < \mathrm{Reach}\,(A)$, then $\mathrm{Nrv}\,\mathcal{C} \simeq A$.*

Proof. We apply the Nerve Lemma to the collection $\{A \cap \mathrm{Hull}_\alpha(C_v) \mid v \in V\}$. Clearly, $A = \bigcup_{v \in V}(A \cap \mathrm{Hull}_\alpha(C_v))$. By Lemma 10, for all $\emptyset \neq \sigma \subset V$, the intersection $A \cap \bigcap_{v\in\sigma} \mathrm{Hull}_\alpha(C_v)$ is either empty or contractible. □

Combining the above lemma and Theorem 1 we thus get that $\mathrm{Nrv}\,\mathcal{C} \simeq \mathrm{Cech}_A(P,\alpha)$ for all finite compact α-robust coverings $\mathcal{C} = \{C_v \mid v \in V\}$ of A with $0 \leq \alpha < \mathrm{Reach}\,(A)$. Next theorem strengthens this result and states mild conditions on P and V under which there exists a sequence of collapses transforming $\mathrm{Cech}_A(P,\alpha)$ into a simplicial complex isomorphic to $\mathrm{Nrv}\,\mathcal{C}$.

Theorem 3. *Let A be a compact set of $\mathbb{R}^d$ and α a real number such that $0 \leq \alpha < \mathrm{Reach}\,(A)$. Let $\mathcal{C} = \{C_v \mid v \in V\}$ be a compact α-robust covering of A. Let P be a finite point set and suppose there exists an injective map $f : V \to P$ such that $C_v \subset \mathrm{B}^\circ(f(v),\alpha)$ for all $v \in V$. Then, there exists a sequence of collapses from $\mathrm{Cech}_A(P,\alpha)$ to $f(\mathrm{Nrv}\,\mathcal{C}) = \{f(\sigma) \mid \sigma \in \mathrm{Nrv}\,\mathcal{C}\}$.*

Proof. We build a family of compact sets $\mathcal{D} = \{D_p(t) \mid (p,t) \in P \times [0,1]\}$ in such a way that if we let $K(t) = \mathrm{Nrv}\{D_p(t) \mid p \in P\}$, then $\mathrm{Cech}_A(P,\alpha) = K(0)$ and $f(\mathrm{Nrv}\,\mathcal{C}) = K(1)$. We then prove that this family meets the hypotheses of Lemma 12, implying that $\mathrm{Cech}_A(P,\alpha)$ can be transformed into $f(\mathrm{Nrv}\,\mathcal{C})$ by a sequence of collapses obtained by increasing continuously t from 0 to 1. To define the family $\mathcal{D}$, let us first associate to every point $p \in P$ the set

$$\mathrm{Split}(p) = \begin{cases} \mathrm{Clenchers}_\alpha(C_v) & \text{if } f^{-1}(p) = \{v\}, \\ \{p^+, p^-\} & \text{if } f^{-1}(p) = \emptyset, \end{cases}$$

where p^+ and p^- are two points which are symmetric with respect to p and chosen such that $B(p^+,\alpha) \cap B(p^-,\alpha) = \emptyset$. We then set

$$D_p(t) = A \cap \bigcap_{s \in \mathrm{Split}(p)} \mathrm{B}((1-t)p + ts, \alpha).$$

Let us check that $\mathrm{Cech}_A(P,\alpha) = K(0)$ and $f(\mathrm{Nrv}\,\mathcal{C}) = K(1)$. We claim that $\mathrm{Split}(p) \neq \emptyset$ for all $p \in P$. Let us consider two cases. First, if $f^{-1}(p) = \emptyset$, then by definition $\mathrm{Split}(p) = \{p^+, p^-\} \neq \emptyset$. Second, if $f^{-1}(p) = \{v\}$, then $\mathrm{Split}(p)$ contains at least p since $C_v \subset \mathrm{B}^\circ(p,\alpha)$ by hypothesis. Thus, $D_p(0) = A \cap \bigcap_{s\in\mathrm{Split}(p)} B(p,\alpha) = A \cap B(p,\alpha)$ and $K(0) = \mathrm{Cech}_A(P,\alpha)$. On the other hand, we have $D_p(1) = A \cap \bigcap_{s\in\mathrm{Split}(p)} B(s,\alpha)$ which we can rewrite as

$$D_p(1) = \begin{cases} A \cap \mathrm{Hull}_\alpha(C_v) & \text{if } f^{-1}(p) = \{v\}, \\ \emptyset & \text{if } f^{-1}(p) = \emptyset. \end{cases}$$

Thus, $K(1) = f(\mathrm{Nrv}\,\mathcal{C})$. Let us make some more remarks. Writing $Z(p,t) = \{(1-t)p + ts \mid s \in \mathrm{Split}(p)\}$, we can express $D_p(t)$ as $A \cap \mathcal{B}(Z(p,t),\alpha)$. Since $\mathrm{Split}(p)$ is compact, so is $Z(p,t)$ and by Lemma 10, $D_p(t)$ is either empty or contractible. Furthermore, $C_v \subset D_p(t)$ for all $p \in f(V)$, showing that the collection of cells $D_p(t)$ cover the shape. Applying the Nerve Lemma, we thus get that $K(t) \simeq A$ for all $t \in [0,1]$. We are now ready to prove a stronger result, namely that as t increases continuously from 0 to 1, the only

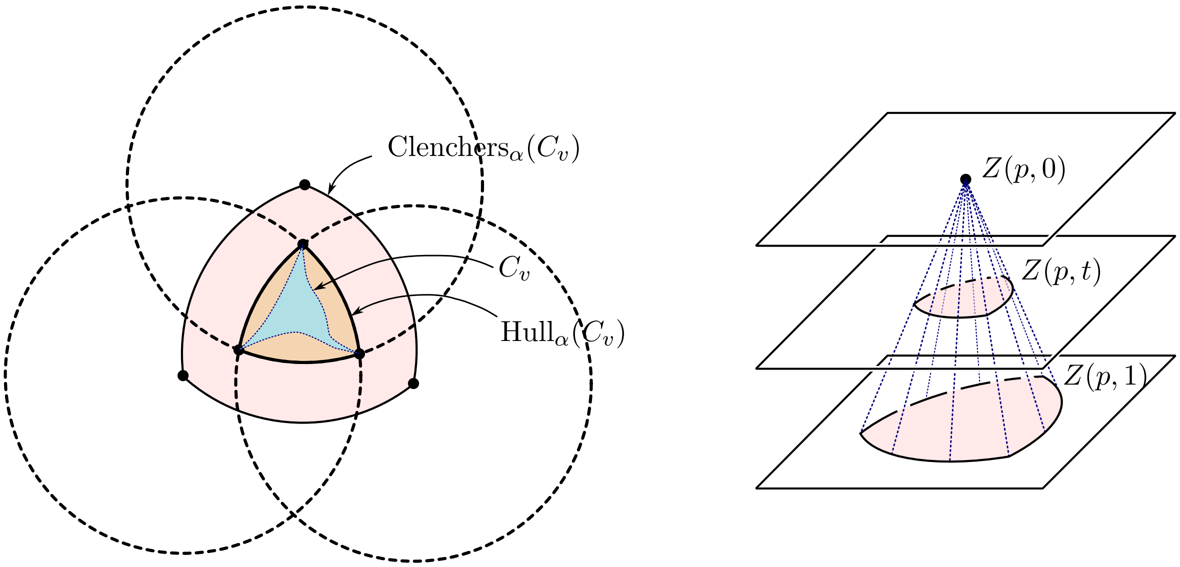

Figure 7: Left: α-hull and α-clenchers of a planar cell C_v. Right: $Z(p,t)$ is the image of $Z(p,1)$ by an homothety centered at p with scale factor t.

changes that may occur in $K(t)$ are collapses. For this, it suffices to establish that the family $\mathcal{D}$ defined above satisfies conditions (a), (b), (c), (d) and (e) of Lemma 12.

(a) Let us prove that for all $0 \leq t < t' \leq 1$ and all $p \in P$, we have $D_p(t') \subset D_p(t)^\circ$. If $f^{-1}(p) = \emptyset$, this is easy to see. Suppose $f^{-1}(p) = \{v\}$; see Figure 7. We note that $Z(p,1) = \text{Split}(p) = \text{Clenchers}_\alpha(C_v)$ is convex and by construction, so are all $Z(p,t)$ for all $t \in [0,1]$. Since $C_v \subset \mathrm{B}^\circ(p,\alpha)$, it follows that p belongs to the interior of $Z(p,1)$ and $Z(p,t) \subset Z^\circ(p,t')$ for all $0 \leq t < t' \leq 1$. This implies that $D_p(t') \subset D_p(t)$. It remains to show that no point of $D_p(t')$ is in $\partial D_p(t)$.

Suppose for a contradiction that $d \in \partial D_p(t) \cap D_p(t')$ and let $r = \max\{\|d - z\| \mid z \in Z(p,t)\}$. The real number r is well-defined since $Z(p,t)$ is compact. From $d \in \partial D_p(t)$ we can easily deduce that $r \geq \alpha$. Indeed, otherwise some neighborhood of d would belong to $D_p(t)$ which is impossible. Let $z \in Z(p,t)$ be such that $\|z - d\| = r$. Since $Z(p,t) \subseteq Z^\circ(p,t')$, there is $z' \in Z(p,t')$ such that $\|z' - d\| > r \geq \alpha$. But, this contradicts $d \in D_p(t')$.

(b) Clearly, $\bigcup_{p \in P} D_p(1) = A$.

(c) Given $\sigma \subset P$ and a map $\phi : \sigma \to [0,1]$, we introduce the set

$$\mathcal{D}(\sigma,\phi) = \bigcap_{p\in\sigma} D_p \circ \phi(p) = A \cap \bigcap_{p\in\sigma} \bigcap_{z \in Z(p)} B((1-\phi(p))p + \phi(p)z, \alpha).$$

By Lemma 10, the intersection $\mathcal{D}(\sigma,\phi)$ is either empty or connected.

(d) Consider $\sigma \subset P$ and a map $\phi : \sigma \to [0,1)$ such that $\mathcal{D}(\sigma,\phi) \neq \emptyset$. Let us prove that if $\mathcal{D}(\sigma,\phi') = \emptyset$ for all maps $\phi' : \sigma \to [0,1]$ with $\phi' \succ \phi$, then $\mathcal{D}(\sigma,\phi)$ is a singleton. Assume, by contradiction, that $\mathcal{D}(\sigma,\phi)$ contains two points x_1 and x_2 and let us prove that we can find a map $\phi' : \sigma \to [0,1]$ such that $\phi' \succ \phi$ and $\mathcal{D}(\sigma,\phi') \neq \emptyset$. Take α' such that $\alpha < \alpha' < \text{Reach}\,(A)$. Since $A \cap \text{Hull}_{\alpha'}(\{x_1,x_2\})$ contains both x_1 and x_2, it is non-empty and therefore contractible by Lemma 10; see Figure 8. In particular, there is a path connecting the points x_1 and x_2 in $A \cap \text{Hull}_{\alpha'}(\{x_1,x_2\})$. This path has to intersect the largest ball B

contained in $\mathrm{Hull}_{\alpha'}(\{x_1, x_2\})$ and therefore $A \cap B \neq \emptyset$. For $\xi > 0$ sufficiently small we have

$$A \cap B^{\oplus\xi} \subset A \cap \mathrm{Hull}_\alpha(\{x_1, x_2\}) \subset \mathcal{D}(\sigma, \phi).$$

By moving slightly the centers of the balls defining $\mathcal{D}(\sigma, \phi)$, that is, by replacing the map ϕ by a map $\phi' \succ \phi$ such that $\phi'(p) - \phi(p)$ is small enough for all $p \in \sigma$, we get a new set $\mathcal{D}(\sigma, \phi')$ that still contains B. Since $\emptyset \neq A \cap B$, we thus get that $\mathcal{D}(\sigma, \phi') \neq \emptyset$, reaching a contradiction.

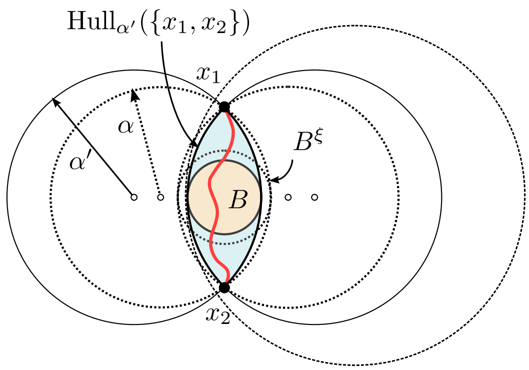

Figure 8: Notation for the proof of Theorem 3.

(e) It is not difficult to see that $\bigcap_{t\in[0,\tau)} Z(p,t) = Z(p,\tau)$ and $D_p(\tau) = \bigcap_{t\in[0,\tau)} D_p(t)$ for $0 < \tau \leq 1$. $\square$

7 Nicely triangulable spaces

Given a space A and a finite sample P of A, we are seeking a sequence of collapses that transform the Čech complex of P with scale parameter α into a *triangulation* of A. We recall that a *triangulation* of A is a simplicial complex whose underlying space is homeomorphic to A. If A has a triangulation, then A is said to be *triangulable*. In particular, we know that compact smooth manifolds are triangulable [22]. Unfortunately, the proof involves barycentric subdivisions whose dual meshes are not likely to have convex cells and therefore have little chance of being α-robust coverings. And yet, we know that if a triangulation T of a space A is the nerve of an α-robust covering of A, then the previous section provides conditions under which $\mathrm{Cech}_A(P, \alpha)$ can be transformed into T by a sequence of collapses. This raises the question of whether, given a space A and a scale parameter α, it is possible to find a triangulation T of A which is the nerve of some α-robust covering of A. In this section, we focus on the question and present examples of spaces enjoying this property.

As a warm-up, we study the easy case $A = \mathbb{R}^2$; see Figure 9. Consider a Delaunay triangulation T of $\mathbb{R}^2$ with vertex set V and write $C_v = \{x \in \mathbb{R}^2 \mid \|x - v\| \leq \|x - u\| \text{ for all } u \in V\}$ for the Voronoi cell of $v \in V$. Setting $\mathcal{C} = \{C_v \mid v \in V\}$ for the collection of Voronoi cells, we have that $T = \mathrm{Nrv}\,\mathcal{C}$. If all angles in T are acute, then the Voronoi cell C_v is contained in the star of v and so is $\mathrm{Hull}_\alpha(C_v)$ for α large enough. In particular, by choosing carefully V and α, we can ensure that $\mathcal{C}$ is an α-robust covering of the plane.

To facilitate our discussion for more general spaces A, we first introduce some more notations and definitions. Given an abstract simplicial complex T, we let $g : V \to \mathbb{R}^{|V|-1}$ be an injective map that sends the vertex set V of T to affinely independent points of $\mathbb{R}^{|V|-1}$. The *underlying space* of T is the point set $|T| = \bigcup_{\sigma \in T} |\sigma|$, where $|\sigma|$ is the geometric simplex obtained by taking the convex hull of $g(\sigma)$. If v is a vertex of T, the *open star* of v in T, denoted by $\mathrm{St}_T(v)$, is the union of the relative interiors of $|\sigma|$ for all σ

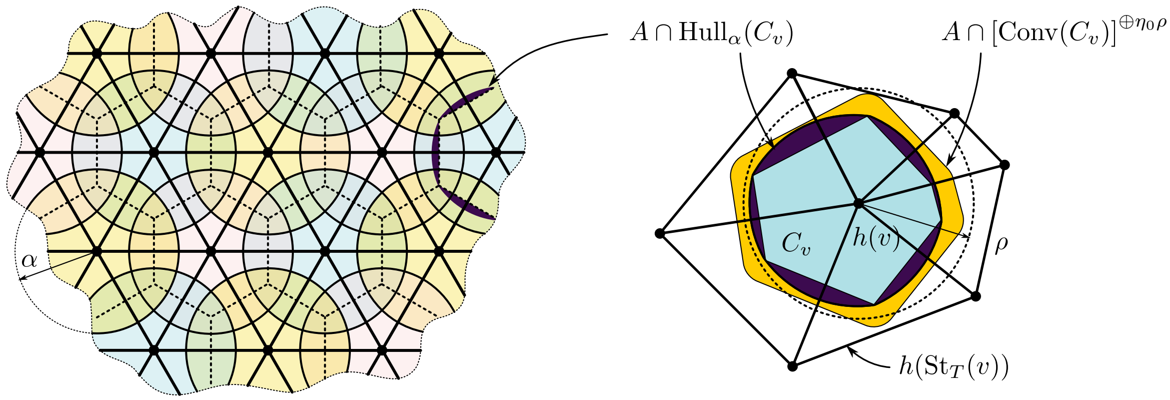

Figure 9: Left: the collection of disks and the collection of Voronoi regions both form an α-robust covering of the plane. Right: A triangulation is nice in our context when, among other things, it is the nerve of a collection of cells C_v with size ρ such that $A \cap [\mathrm{Conv}(C_v)]^{\oplus \eta_0 \rho} \subset h(\mathrm{St}_T(v))$ for some $\eta_0 > 0$. This property will be preserved by $C^{1,1}$ diffeomorphisms for ρ small enough.

of T that contain v [19]. By definition, the set $\mathrm{St}_T(v)$ is thus an open subset of $|T|$. For brevity, we shall write $h(v)$ instead of $h(|v|)$ and $h(\sigma)$ instead of $h(|\sigma|)$. Writing $\mathrm{Conv}(X) = \mathrm{Hull}_{+\infty}(X)$ for the convex hull of $X \subset \mathbb{R}^d$, we introduce the following definition:

Definition 2 (nice triangulation). *Let ρ and δ be two positive real numbers. A triangulation T of $A \subset \mathbb{R}^d$ is said to be (ρ, δ)-*nice *with respect to $(h, \mathcal{C})$ if h is a homeomorphism from $|T|$ to A, $\mathcal{C} = \{C_v \mid v \in V\}$ is a finite compact covering of A such that $\mathrm{Nrv}\, \mathcal{C} = T$ and the following conditions hold:*

(i) $h(\sigma) \subset \bigcup_{v \in \sigma} C_v$ *for all simplices* $\sigma \in T$*;*

(ii) $C_v \subset \mathrm{B}^\circ(h(v), \rho)$ *for all* $v \in V$*;*

(iii) $A \cap [\mathrm{Conv}(C_v)]^{\oplus \delta} \subset h(\mathrm{St}_T(v))$ *for all* $v \in V$.

The use of $\mathrm{Conv}(C_v)$ in the last item of the definition is motivated by the following geometric lemma:

Lemma 14. *Let $X \subset \mathbb{R}^d$ be a non-empty compact set and $B(c, \rho)$ its smallest enclosing ball. For all α and δ such that $\alpha \geq \rho$ and $\alpha - \sqrt{\alpha^2 - \rho^2} \leq \delta$, the following inclusion holds:* $\mathrm{Hull}_\alpha(X) \subset [\mathrm{Conv}(X)]^{\oplus \delta}$.

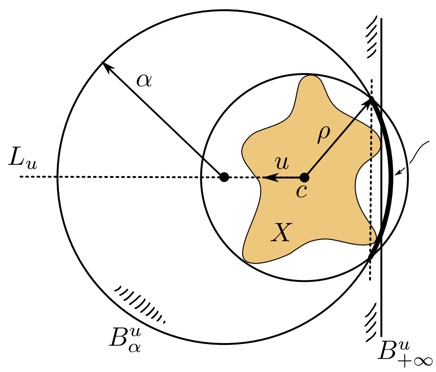

Figure 10: Notation for the proof of Lemma 14.

Proof. See Figure 10. Consider a unit vector $u \in \mathbb{S}^{d-1}$ and let L_u be the half-line emanating from c in direction u. Let B^u_α denote the ball with radius α centered on L_u containing X and whose center is furthest away from c. By construction, X is contained in the intersection of the two balls B^u_α and $B(c, \rho)$. The boundary of $B^u_\alpha \cap B(c, \rho)$ consists of two spherical caps and we let C^u_α be the one lying on the sphere bounding B^u_α. Observe that X has a non-empty intersection with C^u_α and for all $\beta \geq \alpha$, the ball B^u_β intersects C^u_α. The largest distance between a point of C^u_α and B^u_β is upper bounded by the height of C^u_α which is less than or equal to $\alpha - \sqrt{\alpha^2 - \rho^2} \leq \delta$. We thus get that $B^u_\alpha \subset [B^u_\beta]^{\oplus\delta}$. Considering this inclusion over all directions u for $\beta = +\infty$ yields the result. □

It follows that if T is a (ρ, δ)-nice triangulation of A with respect to $(h, \mathcal{C})$, we are able to derive conditions on α, ρ and δ which guarantee that $\mathcal{C}$ is an α-robust covering of A.

Lemma 15. *Let A be a compact set of $\mathbb{R}^d$ and suppose T is a (ρ, δ)-nice triangulation of A with respect to $(h, \mathcal{C})$. Then $\mathcal{C}$ is an α-robust covering of A whenever the following two conditions are fulfilled: (1) $\rho \leq \alpha$ and (2) $\alpha - \sqrt{\alpha^2 - \rho^2} \leq \delta$.*

Proof. Suppose $\mathcal{C} = \{C_v \mid v \in V\}$ and let $v \in V$. By Lemma 14, $\mathrm{Hull}_\alpha(C_v) \subset [\mathrm{Conv}(C_v)]^{\oplus\delta}$; see Figure 9, right. It follows that $C_v \subset A \cap \mathrm{Hull}_\alpha(C_v) \subset h(\mathrm{St}_T(v))$ from which we deduce the sequence of inclusions

$$T = \mathrm{Nrv}\,\mathcal{C} \subset \mathrm{Nrv}\{A \cap \mathrm{Hull}_\alpha(C_v) \mid v \in V\} \subset \mathrm{Nrv}\{h(\mathrm{St}_T(v)) \mid v \in V\} = T.$$

The nerves on the left and on the right are equal, showing that $\mathrm{Nrv}\,\mathcal{C} = \mathrm{Nrv}\{A \cap \mathrm{Hull}_\alpha(C_v) \mid v \in V\}$. □

Observe that if T is a $(\rho, \eta_0\rho)$-nice triangulation of A for some $\eta_0 > 0$, then conditions (1) and (2) of Lemma 15 are satisfied for $\delta = \eta_0\rho$ as soon as ρ is small enough. Of course, the difficult question is whether such a triangulation T can always be found for arbitrarily small ρ.

Definition 3 (nicely triangulable)**.** *We say that $A \subset \mathbb{R}^d$ is nicely triangulable if we can find $\rho_0 > 0$ and $\eta_0 > 0$ such that for all $0 < \rho < \rho_0$, there is a $(\rho, \eta_0\rho)$-nice triangulation of A.*

Theorem 4. *Suppose $A \subset \mathbb{R}^d$ is nicely triangulable. For every $0 < \alpha < \mathrm{Reach}\,(A)$, there exists $\varepsilon_0 > 0$ such that for all finite point set $P \subset \mathbb{R}^d$ and all $0 < \varepsilon < \varepsilon_0$ satisfying $A \subset P^{\oplus\varepsilon}$, the complex $\mathrm{Cech}_A(P, \alpha)$ can be transformed into a triangulation of A by a sequence of collapses.*

Proof. By definition, we can find $\rho_0 > 0$ and $\eta_0 > 0$ such that for all $0 < \rho < \rho_0$, there is a $(\rho, \eta_0\rho)$-nice triangulation T of A with respect to $(h, \mathcal{C})$. Let us choose ρ small enough so that $\rho < \alpha$ and $\alpha - \sqrt{\alpha^2 - \rho^2} \leq \eta_0\rho$. Lemma 15 then implies that $\mathcal{C}$ is a compact α-robust covering of A. Set $e(T, h) = \frac{1}{2} \inf \|h(v_1) - h(v_2)\|$ where the infimum is over all pairs of vertices $v_1 \neq v_2$ of T and let ε_0 be the minimum of $e(T, h)$ and $\alpha - \rho$. Consider a function $f : \mathrm{Vert}(T) \to P$ that maps each vertex v to a point of P closest to $h(v)$. Note that f is injective, $\|h(v) - f(v)\| \leq \varepsilon$ and $C_v \subset \mathrm{B}^\circ(f(v), \alpha)$ for all $v \in V$. Applying Theorem 3 yields the existence of a sequence of collapses from $\mathrm{Cech}_A(P, \alpha)$ to $f(T)$. □

The next theorem provides a few examples of nicely triangulable manifolds.

Theorem 5. *The following embedded manifolds are nicely triangulable:*

1. *The unit 2-sphere $\mathbb{S}^2 = \{x = (x_1, x_2, x_3) \in \mathbb{R}^3 \mid \sum_{i=1}^3 x_i^2 = 1\}$;*
2. *The flat torus $\mathbb{T}^2 = \{x = (x_1, x_2, x_3, x_4) \in \mathbb{R}^4 \mid x_1^2 + x_2^2 = 1$ and $x_3^2 + x_4^2 = 1\}$;*

3. *The m-dimensional Euclidean space $\mathbb{R}^m$, embedded in $\mathbb{R}^d$ for some $m \leq d$.*

Proof. For $A \in \{\mathbb{S}^2, \mathbb{T}^2, \mathbb{R}^m\}$, we proceed as follows. We build a triangulation T parameterized by some integer n and consider a map $h : |T| \to A$. The integer n will control the size of elements in $h(T)$: the larger n the smaller the image of simplices under h. We then consider the barycentric subdivision K of T and associate to each vertex v of T the cell $C_v = \bigcup_{\sigma \ni v} h(\sigma)$. The collection of cells C_v forms a covering $\mathcal{C}$ of A. In the three cases, it is not difficult to see that we can find $\eta_0 > 0$ such that T is $(\rho, \eta_0 \rho)$-nice with respect to $(h, \mathcal{C})$ for some $\rho > 0$. Furthermore, the value of ρ can be made as small as desired by increasing n. We thus conclude that A is nicely triangulable. Below, we just describe how T and h are chosen in each case.

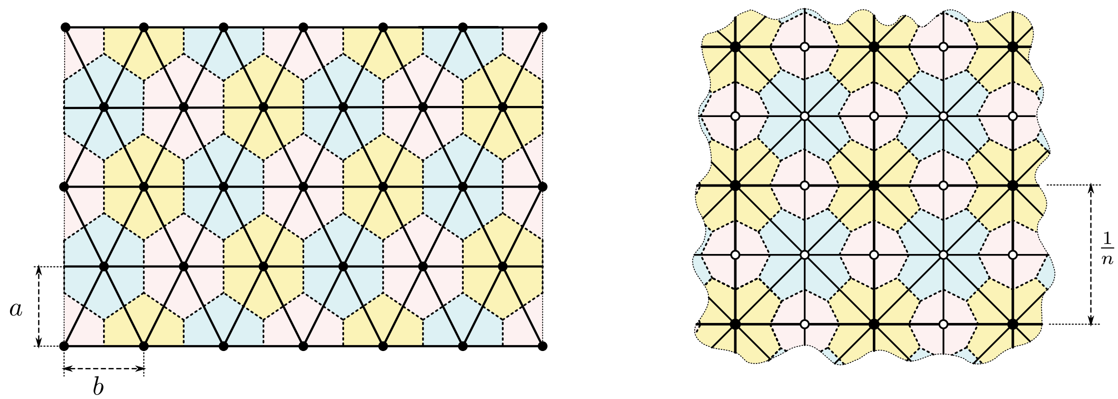

Figure 11: Triangulating $\mathbb{T}^2$ (left) and $\mathbb{R}^2$ (right).

1. $\mathbb{S}^2$ is nicely triangulable. We start with an icosahedron centered at the origin and subdivide each triangular face into 4^n equilateral triangles. Notice that all vertices of the resulting triangulation T have degree 6 but the 12 vertices in the original icosahedron which have degree 5. The triangulation T is then projected onto the sphere, using the projection map $h : |T| \to \mathbb{S}^2$ defined by $h(x) = \frac{x}{\|x\|}$.

2. $\mathbb{T}^2$ is nicely triangulable. The map $H : \mathbb{R}^2 \to \mathbb{T}^2$ defined by $H(s,t) = (\cos s, \sin s, \cos t, \sin t)$ is locally isometric and its restriction $h : [0, 2\pi)^2 \to \mathbb{T}^2$ is an homeomorphism. The idea is to build a periodic tiling of $\mathbb{R}^2$ made up of identical isosceles triangles as in Figure 11, left. Let a and b be the respective height and basis of the triangles. Consider two integers n and k such that $na = kb = 2\pi$. Taking $k = \lfloor \frac{\sqrt{3}}{2} n \rfloor$ we get that the ratio $\frac{a}{b}$ tends to $\frac{\sqrt{3}}{2}$ as $n \to +\infty$. Thus, the map H turns the periodic tiling of $\mathbb{R}^2$ into a triangulation of $\mathbb{T}^2$ whose triangles become arbitrarily close to equilateral triangles with edge length b as $n \to +\infty$.

3. $\mathbb{R}^m$ is nicely triangulable. We start with a cubical regular grid and define T as the barycentric subdivision of that grid; see Figure 11, right. Precisely, for each cell in the grid, we insert one vertex at its centroid. So each edge is subdivided into 2 edges sharing the inserted vertex. We then recursively subdivide the cells by ascending dimension. Each cubical k-cell has $2k$ cubical $(k-1)$-cells on its boundary. We subdivide each k-cell as a cone whose apex is the inserted vertex and whose basis is the subdivided boundary of that cell. We claim that all stars in T are convex. Indeed each vertex in T is the centroid of an initial cubical cell of dimension between 0 and m. Consider the vertex v that was inserted at the center of the k-dimensional cubical cell D_v and let us describe the set of vertices V_v in the link of v in T. The vertices of V_v can be partitioned in two subsets. The first subset contains vertices in the k-flat that supports D_v while the second subset contains vertices in the $(d-k)$-flat passing through v and orthogonal to the k-flat supporting D_v. The vertices in the first subset lie on the boundary of D_v and the vertices in the second subset lie on the boundary of a $(m-k)$-cube. Since in both flats of respective dimension k and $m-k$ the vertices in V_v are in convex position, it results that vertices in V_v are in convex position in $\mathbb{R}^m$. As a result, it can be proved (details are skipped) that the star of v is the convex hull of V_v. Finally, we let h be the identity map and n the inverse of the size of the grid. □

We now establish that the property of being nicely triangulable is preserved by $C^{1,1}$ diffeomorphisms between manifolds. Let us make precise what we mean in Theorem 6 by embedded $C^{1,1}$ k-manifolds. A $C^{1,1}$ function is a differentiable function with a Lipschitz derivative. A $C^{1,1}$ structure on a manifold is an equivalence class of atlases whose transition functions are $C^{1,1}$. Finally, $C^{1,1}$ diffeormorphisms between $C^{1,1}$ manifolds are defined accordingly. In Theorem 6, we restrict our attention to shapes which are $C^{1,1}$ compact manifolds without boundary embedded in $\mathbb{R}^d$ and whose embeddings are themselves regular and $C^{1,1}$ (for the differential structure induced by $\mathbb{R}^d$) , where "regular" means that the derivative of the embedding has full rank everywhere. We will say that such shapes are compact $C^{1,1}$ manifolds embedded in $\mathbb{R}^d$ for short. The assumption of regular embeddings entails the existence of well-defined tangent affine spaces. A compact manifold embedded in $\mathbb{R}^d$ is $C^{1,1}$ if and only if it has a positive reach [16].

Theorem 6. *Let M and M' be two compact $C^{1,1}$ k-manifolds without boundary embedded respectively in $\mathbb{R}^d$ and $\mathbb{R}^{d'}$ and $\Phi : M \to M'$ a $C^{1,1}$ diffeomorphism. M is nicely triangulable if and only if M' is nicely triangulable.*

The proof is given in the Appendix.

8 Discussion

The paper leaves unanswered a few questions that we discuss now:

(1) Our result assumes the shape to be nicely triangulated. In the paper, we list a few simple spaces which enjoy this property. Is it possible to extend the list to a larger class of spaces? We conjecture that compact smooth k-manifolds embedded in $\mathbb{R}^d$ are nicely triangulable. Indeed, for $k=2$, it is known that any compact connected surface (without boundary) embedded in $\mathbb{R}^3$ is homeomorphic to either the 2-sphere or a connected sum of g tori for $g \geq 1$. Hence, thanks to Theorem 6, it would suffice to provide a template of nicely triangulable surface of genus g for each $g \geq 2$, in a way similar to what we did for $g=0$ and $g=1$. Unfortunately, for higher dimensional manifolds, one cannot rely anymore on an existing classification. Another approach has to be considered.

(2) Our proof is not constructive. Indeed, the order in which to collapse faces in the Čech complex is determined by sweeping space with a t-offset of the shape for decreasing values of t. Since the common setting consists in describing the shape through a finite sample, the knowledge of the t-offsets of the shape is lost. Nonetheless, is it possible to turn our proof into an algorithm? Can we do the same for Rips complexes? A positive answer is even more desirable for the second class of complexes due to their computational tractability. We leave those questions open for future work.

9 Acknowledgments

We thank the anonymous reviewers for many helpful suggestions.

References

[1] K. Adiprasito and B. Benedetti. Metric geometry and collapsibility. *arXiv preprint arXiv:1107.5789*, 2011.

[2] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22(4):481–504, 1999.

[3] D. Attali, H. Edelsbrunner, and Y. Mileyko. Weak witnesses for Delaunay triangulations of submanifold. In *ACM Sympos. Solid and Physical Modeling*, pages 143–150, Beijing, China, June 4-6 2007. [download].

[4] D. Attali and A. Lieutier. Reconstructing shapes with guarantees by unions of convex sets. In *Proc. 26th Ann. Sympos. Comput. Geom.*, pages 344–353, Snowbird, Utah, June 13-16 2010.

[5] D. Attali, A. Lieutier, and D. Salinas. Efficient data structure for representing and simplifying simplicial complexes in high dimensions. *International Journal of Computational Geometry and Applications (IJCGA)*, 22(4):279–303, 2012. [download].

[6] D. Attali, A. Lieutier, and D. Salinas. Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes. *Computational Geometry: Theory and Applications (CGTA)*, 2012. [download].

[7] A. Bjorner. Topological methods. In *Handbook of combinatorics (vol. 2)*, page 1850. MIT Press, 1996.

[8] J.-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Computational Geometry: Theory and Applications*, 19(2):155–173, 2001.

[9] E. Carlsson, G. Carlsson, and V. De Silva. An algebraic topological method for feature identification. *International Journal of Computational Geometry and Applications*, 16(4):291–314, 2006.

[10] E. Chambers, V. De Silva, J. Erickson, and R. Ghrist. Vietoris–rips complexes of planar point sets. *Discrete & Computational Geometry*, 44(1):75–90, 2010.

[11] 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.

[12] 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.

[13] M. M. Cohen. *A course in Simple-Homotopy theory*. Springer-Verlag New York, 1973.

[14] V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology*, 7:339–358, 2007.

[15] M. C. Delfour and J.-P. Zolésio. *Shapes and geometries: Metrics, Analysis, Differential Calculus, and Optimization*. Siam, 2011.

[16] H. Federer. Curvature measures. *Trans. Amer. Math. Soc*, 93:418–491, 1959.

[17] J. Hausmann. On the Vietoris-Rips Complexes and a Cohomology Theory for Metric Spaces. *Ann. Math. Studies*, 138:175–188, 1995.

[18] J. Latschev. Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold. *Archiv der Mathematik*, 77(6):522–528, 2001.

[19] J. Munkres. *Elements of algebraic topology*. Perseus Books, 1993.

[20] 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.

[21] A. Tahbaz-Salehi and A. Jadbabaie. Distributed coverage verification in sensor networks without location information. *Automatic Control, IEEE Transactions on*, 55(8):1837–1849, 2010.

[22] H. Whitney. *Geometric integration theory*. Dover Publications, 2005.

A $C^{1,1}$ diffeomorphisms preserve nicely triangulable manifolds

The goal of this section is to prove Theorem 6.

Proof of Theorem 6. Let $x \in M$. Since M is a compact $C^{1,1}$ k-manifold embedded in $\mathbb{R}^d$, there exists a k-dimensional affine space $T_M(x) \subset \mathbb{R}^d$ tangent to M at x. Let $\pi_x : M \to T_M(x)$ be the orthogonal projection onto the tangent space $T_M(x)$ and let $\pi'_{\Phi(x)} : M' \to T_{M'}(\Phi(x))$ the orthogonal projection onto $T_{M'}(\Phi(x))$. Since M and M' are compact, we can find two constants K and K' independent of x such that:

$$\forall y \in M, \quad \|y - \pi_x(y)\| < K\|y - x\|^2 \tag{1}$$

$$\forall y \in M', \quad \|y - \pi'_{\Phi(x)}(y)\| < K'\|y - \Phi(x)\|^2 \tag{2}$$

Given $t_0 > 0$, we consider the open set $U_x = M \cap \mathrm{B}^\circ(x, t_0)$ and adjust t_0 in such a way that

1. The restriction $\pi_x : U_x \to \pi_x(U_x)$ is an homeomorphism for all $x \in M$;
2. The restriction $\pi'_{\Phi(x)} : \Phi(U_x) \to \pi'_{\Phi(x)}(\Phi(U_x))$ is also an homeomorphism for all $x \in M$.

For sake of conciseness, we only sketch a justification for the existence of such a $t_0 > 0$. The local property (i.e. the existence of $t_0 > 0$ for a given $x \in M$) follows easily from the definition of embedded $C^{1,1}$ k-manifolds. Indeed, the assumption of a regular embedding entails that π_x has full rank derivative at x and the inverse function theorem can be applied to get the local property. In order to get a uniform $t_0 > 0$ (the requested global property) one can establish first the following strengthening of the local property: For any $x \in M$, there is $t_x > 0$ such that for any $y \in M \cap B^\circ(x, t_x)$, the restriction of π_y to $M \cap B^\circ(x, t_x)$ is a C^1 homeomorphism. Compactness of M can then be used in the usual manner to get a uniform t_0.

The collection of pairs $\{(U_x, \pi_x)\}_{x \in M}$ forms an atlas in the $C^{1,1}$ structure of M. Similarly, the collection of pairs $\{(\Phi(U_x), \pi'_{\Phi(x)})\}_{x \in M}$ forms an atlas in the $C^{1,1}$ structure of M'. Let $\mathcal{T}_M(x)$ be the linear space associated to $T_M(x)$ and denote by $D\Phi_x$ the derivative of Φ at x, seen as a linear map between $\mathcal{T}_M(x)$ and $\mathcal{T}_{M'}(\Phi(x))$. Since M is compact, there is a constant K_Φ independent of x such that:

$$\forall y \in M, \quad \|\Phi(y) - \Phi(x) - D\Phi_x(\pi_x(y) - x)| < K_\Phi \|y - x\|^2, \tag{3}$$

and two positive numbers $\kappa_2 \geq \kappa_1 > 0$, again independent of x by compactness of M, such that

$$\forall u \in \mathcal{T}_M(x), \quad \kappa_1\|u\| \leq \|D\Phi_x(u)\| \leq \kappa_2\|u\|. \tag{4}$$

Consider the affine function $\hat{\Phi}_x : T_M(x) \to T_{M'}(\Phi(x))$ defined by $\hat{\Phi}_x(y) = \Phi(x) + D\Phi_x(y - x)$. Combining Equations (1) (2) and (3), we can find a constant L_Φ independent of x such that for all $t < t_0$ and all compact sets $A \subset M \cap B(x, t)$:

$$d_H(\hat{\Phi}_x \circ \pi_x(A), \pi'_{\Phi(x)} \circ \Phi(A)) < L_\Phi t^2 \tag{5}$$

Now, assume that M is nicely triangulable and let us prove that M' is also nicely triangulable. By definition, we can find $\rho_0 > 0$ and $\eta_0 > 0$ such that, for all $0 < \rho < \rho_0$, there is a $(\rho, \eta_0\rho)$-nice triangulation T of M with respect to some $(h, \mathcal{C})$. Suppose $\mathcal{C} = \{C_v \mid v \in V\}$ and consider the covering $\mathcal{C}' = \{\Phi(C_v) \mid v \in V\}$, the homeomorphism $h' = \Phi \circ h : |T| \to M'$, the real numbers $\rho' = 2\kappa_2\rho$ and $\eta'_0 = \frac{\kappa_1\eta_0 - 5L_\Phi\rho}{2\kappa_2}$. Let us prove that by choosing ρ small enough, T is a $(\rho', \eta'_0\rho')$-nice triangulation of M' with respect to $(\mathcal{C}', h')$. In other words, we need to check that conditions (ii) and (iii) of Definition 2 are satisfied for $\mathcal{C} = \mathcal{C}'$, $h = h'$, $\rho = \rho'$ and $\delta = \eta'_0\rho'$. Take $v \in V$ and set $x = h(v)$, $C = C_v$, $S = h(\mathrm{St}_T(v))$.

(ii) By definition of T, we have $x \in C \subset \mathrm{B}^\circ(x,\rho)$. Taking the image of this relation under Φ and choosing $\rho > 0$ small enough, we get that $\Phi(x) \in \Phi(C) \subset \Phi(\mathrm{B}^\circ(x,\rho)) \subset \mathrm{B}^\circ(\Phi(x),\rho')$. The last inclusion is obtained by combining Equations (1), (3) and (4).

(iii) Let us choose a positive real number $\rho < \min\{\rho_0, \frac{t_0}{2}\}$ small enough to ensure that $\eta_0' > 0$ and let us prove that $M' \cap [\mathrm{Conv}(C)]^{\oplus \eta_0' \rho'} \subset S$. By choice of T as a $(\rho, \eta_0 \rho)$-nice triangulation of M with respect to $(h, \mathcal{C})$, we have that $M \cap \mathrm{Conv}(C)^{\oplus \eta_0 \rho} \subset S$. Furthermore, $C \subset B(x,\rho)$ and $S \subset B(x, 2\rho)$. Thus, by choosing $\rho < \frac{t_0}{2}$, we have $S \subset U_x$ and

$$U_x \cap \mathrm{Conv}(C)^{\oplus \eta_0 \rho} \subset S.$$

Taking the image by the homeomorphism $\pi_x : U_x \to \pi_x(U_x)$ on both sides and using $\pi_x(A \cap B) = \pi_x(A) \cap \pi_x(B)$ we get

$$\pi_x(\mathrm{Conv}(C)^{\oplus \eta_0 \rho}) \subset \pi_x(S).$$

Let $B_k(0,r)$ denote the k-dimensional ball of $\mathcal{T}_M(x)$ centered at the origin with radius r. Writing $A \oplus B = \{a + b \mid a \in A, b \in B\}$ for the Minkowski sum of A and B, it is not too difficult to prove that $\pi_x(A^{\oplus \delta}) = \pi_x(A) \oplus B_k(0,\delta)$. It follows that

$$\pi_x(\mathrm{Conv}(C)) \oplus B_k(0, \eta_0 \rho) \subset \pi_x(S).$$

Taking the image under $\hat{\Phi}_x$ on both sides we get

$$\hat{\Phi}_x \circ \pi_x(\mathrm{Conv}(C)) \oplus D\Phi_x B_k(0, \eta_0 \rho) \subset \hat{\Phi}_x \circ \pi_x(S)).$$

Let $B_k'(0,r)$ denote the k-dimensional ball of $\mathcal{T}_{M'}(\Phi(x))$ centered at the origin with radius r. Using Equation (4) we get that $B_k'(0, \kappa_1 \eta_0 \rho) \subset D\Phi_x B_k(0, \eta_0 \rho)$. Since $\hat{\Phi}_x$ and π_x are both affine, so is the composition and therefore $\hat{\Phi}_x \circ \pi_x(\mathrm{Conv}(C)) = \mathrm{Conv}(\hat{\Phi}_x \circ \pi_x(C))$. It follows that

$$\mathrm{Conv}(\hat{\Phi}_x \circ \pi_x(C)) \oplus B_k'(0, \kappa_1 \eta_0 \rho) \subset \hat{\Phi}_x \circ \pi_x(S).$$

Recalling that $C \subset B(x,\rho)$ and $S \subset B(x, 2\rho)$ and combining the above inclusion with Equation (5) we obtain

$$\mathrm{Conv}(\pi_x' \circ \Phi(C)) \oplus B_k(0, \kappa_1 \eta_0 \rho - 5 L_\Phi \rho^2) \subset \pi_x' \circ \Phi(S).$$

Interchanging Conv and π_x', noting that $\eta_0' \rho' = \kappa_1 \eta_0 \rho - 5 L_\Phi \rho^2$ and using $\pi_x'(A^{\oplus \delta}) = \pi_x'(A) \oplus B_k(0,\delta)$ we get

$$\pi_x'(\mathrm{Conv}(\Phi(C))^{\oplus \eta_0' \rho'}) \subset \pi_x' \circ \Phi(S).$$

Since $\pi_x' : \Phi(U_x) \to \pi_{\Phi(x)}'(\Phi(U_x))$ is homeomorphic, we thus obtain $M' \cap \mathrm{Conv}(\Phi(C))^{\oplus \eta_0' \rho} \subset \Phi(S)$ as desired. $\square$