When Convexity Helps Collapsing Complexes

Dominique Attali

Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, Grenoble, France
Dominique.Attali@grenoble-inp.fr

André Lieutier

Dassault systèmes, Aix-en-Provence, France
andre.lieutier@3ds.com

David Salinas

Amazon research, Berlin, Germany
dsalina@amazon.com

Abstract

This paper illustrates how convexity hypotheses help collapsing simplicial complexes. We first consider a collection of compact convex sets and show that the nerve of the collection is collapsible whenever the union of sets in the collection is convex. We apply this result to prove that the Delaunay complex of a finite point set is collapsible. We then consider a convex domain defined as the convex hull of a finite point set. We show that if the point set samples sufficiently densely the domain, then both the Čech complex and the Rips complex of the point set are collapsible for a well-chosen scale parameter. A key ingredient in our proofs consists in building a filtration by sweeping space with a growing sphere whose center has been fixed and studying events occurring through the filtration. Since the filtration mimics the sublevel sets of a Morse function with a single critical point, we anticipate this work to lay the foundations for a non-smooth, discrete Morse Theory.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases collapsibility, convexity, collection of compact convex sets, nerve, filtration, Delaunay complex, Čech complex, Rips complex

Digital Object Identifier 10.4230/LIPIcs.SoCG.2019.11

1 Introduction

Contractibility and collapsibility. In the realm of point set topology in Euclidean spaces, a set is said to be *contractible* if it has the homotopy type of a point. Examples of contractible sets are bounded convex sets which can each be continuously retracted to a point. Consider a finite collection of convex sets $\mathcal{C}$ and the domain defined by their union $\bigcup \mathcal{C}$. We know that each convex set in the collection is contractible. Moreover, any non-empty intersection of two or more convex sets being itself convex is again contractible. In this situation, the topology of the domain $\bigcup \mathcal{C}$ is determined by the pattern in which convex sets in $\mathcal{C}$ intersect. This is asserted by the Nerve Lemma [6] also known as Leray's Theorem. Recall that the *nerve* of a collection of sets is the abstract simplicial complex whose vertices correspond to sets in the collection and whose simplices correspond to sub-collections with non-empty intersection. The nerve provides a way to record the intersection pattern of sets in a collection. The Nerve Lemma implies that the nerve inherits the homotopy type of $\bigcup \mathcal{C}$. In particular, if $\bigcup \mathcal{C}$ is convex, then it is contractible and so is the nerve.

Simplicial collapses are unitary operations on simplicial complexes that preserve the homotopy type while removing a few simplices. A simplicial complex is said to be *collapsible* if it can be reduced to a single vertex by a sequence of collapses. Collapsibility can be interpreted as a combinatorial version of contractibility. Indeed, each sequence of collapses

35th International Symposium on Computational Geometry (SoCG 2019).
Editors: Gill Barequet and Yusu Wang; Article No. 11; pp. 11:1–11:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

can be interpreted as a combinatorial analog of a deformation retract. Whereas contractibility involves continuous processes, the notion of collapsibility, thanks to its combinatorial and finitary nature, can be handled directly in a computational context [4, 2, 3].

If a simplicial complex is collapsible, then it is also contractible. However, the converse does not hold – the triangulated Bing's house is a famous example of simplicial complex which is contractible without being collapsible [9]. A natural question to ask is: Under which condition a simplicial complex is collapsible? This question has been studied in different context. For instance, it is known that all triangulation of the 2-dimensional ball are collapsible. Whether linear subdivisions of convex d-balls, or at least their subdivision, are collapsible has been a long-standing open problem [5].

Contributions. In this paper, we establish collapsibility of certain simplicial complexes. First, we consider a finite collection of compact convex sets whose union is convex and prove that the nerve of the collection is collapsible. This is a stronger result than what we obtain when applying the Nerve Lemma which only entails contractibility of the nerve. Then, we focus on two simplicial complexes built upon a finite point set S. The first one is the Čech complex $\mathcal{C}(S, r)$ of S with scale parameter r defined as the nerve of the collection of balls of radius r centered at S. The second one is the Vietoris-Rips $\mathcal{R}(S, r)$ which is the largest simplicial complex sharing with $\mathcal{C}(S, r)$ the same set of vertices and edges. In other words, a simplex belongs to $\mathcal{R}(S, r)$ if and only if all its vertices and edges belong to $\mathcal{C}(S, r)$. Such a simplicial complex which enjoys the property to be completely determined by the set of its vertices and edges is called a *flag completion*. We study the situation in which the point set S samples sufficiently densely a given convex domain. As a second result, we obtain that, in this particular case, both the Čech complex and the Vietoris-Rips complex of S are collapsible for a well-chosen scale parameter. Furthermore, when reducing the Vietoris-Rips complex to a point by a sequence of collapses, we can do it in such a way that we preserve at all time the property of the complex to be a flag completion. This opens the possibility to compute the sequence of collapses by maintaining the graph formed by the vertices and edges, thus avoiding the complexity, possibly exponential in the dimension, required by an extensive complex representation; see [1] for an example of a data structure optimized for "almost-flag" completions.

2 Preliminaries

In this section we review the necessary background and explain some of our terms.

2.1 Euclidean space, distances and convex sets

$\mathbb{R}^d$ denotes the Euclidean space and $\|x - y\|$ the Euclidean distance between two points x and y of $\mathbb{R}^d$. Given a point $o \in \mathbb{R}^d$ and a subset $X \subseteq \mathbb{R}^d$, we write $d(o, X)$ for the infimum of Euclidean distances between o and points in X. By convention, we set $d(o, X) = +\infty$ whenever X is empty. We write ∂X for the boundary of X. The closed ball with center x and radius r is denoted by $B(x, r)$. The dilation of X by a ball of radius r centered at the origin is $X^{\oplus r} = \bigcup_{x \in X} B(x, r)$ and referred to as the *r-offset* of X. If X is either compact or convex, so is $X^{\oplus r}$. It is easy to check that $(X \cup Y)^{\oplus r} = X^{\oplus r} \cup Y^{\oplus r}$. In this paper, we will use many times the following fact.

▶ Remark. Let $o \in \mathbb{R}^d$ and $X \subseteq \mathbb{R}^d$. If X is a non-empty compact convex set, then there is a unique point $x \in X$ such that $d(o, X) = \|o - x\|$.

2.2 Abstract simplicial complexes

An *(abstract) simplex* is any finite non-empty set. The *dimension* of a simplex σ is one less than its cardinality. A *k-simplex* designates a simplex of dimension k. If $\tau \subseteq \sigma$ is a non-empty subset, we call τ a *face* of σ and σ a *coface* of τ. If in addition $\tau \neq \sigma$, we say that τ is a *proper face* and σ is a *proper coface*. Given a set of simplices Δ and a simplex $\sigma \in \Delta$, we say that σ is *inclusion-maximal* in Δ if it has no proper coface in Δ. Similarly, we say that σ is *inclusion-minimal* if it has no proper face in Δ. An *abstract simplicial complex* is a collection of simplices K, that contains, with every simplex, the faces of that simplex. The vertex set of the abstract simplicial complex K is the union of its elements, $\mathrm{Vert}(K) = \bigcup_{\sigma \in K} \sigma$. A *subcomplex* of K is a simplicial complex $L \subseteq K$. The *star* of σ in K, denoted $\mathrm{St}_K(\sigma)$, is the set of cofaces of σ. The *link* of σ in K, denoted $\mathrm{Lk}_K(\sigma)$, 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. Another particular subcomplex is the *i-skeleton* consisting of all simplices of dimension i or less. We call the simplicial complex formed by a simplex and all its faces the *closure* of that simplex. The closure of a simplex is an example of cone. A *cone* is a simplicial complex L which contains a vertex o such that the following holds: $\sigma \in L \implies \sigma \cup \{o\} \in L$.

2.3 Collapses

Let $\pi : \mathrm{Vert}(K) \to \mathbb{R}^n$ be an injective map that sends the n vertices of K to n affinely independent points of $\mathbb{R}^n$, such as for instance the n vectors of the standard basis of $\mathbb{R}^n$. Let $\mathrm{Hull}(X)$ denote the convex hull of $X \subseteq \mathbb{R}^n$. The *underlying space* of K is the point set $|K| = \bigcup_{\sigma \in K} \mathrm{Hull}(\pi(\sigma))$ and is defined up to a homeomorphism. We shall say that an operation preserves the homotopy-type of K if the result is a simplicial complex K' whose underlying space is homotopy equivalent to that of K. We are interested in simplifying a simplicial complex through a sequence of homotopy-preserving operations.

Consider the operation that removes from K the set of simplices $\Delta = \mathrm{St}_K(\sigma)$. This operation is known to preserve the homotopy-type in the following three cases:

1. $\Delta = \{\sigma, \tau\}$ with $\sigma \neq \tau$. This case can also be characterized by the fact that the link of σ is reduced to a singleton. The operation is called an *(elementary) collapse*.
2. $\Delta = \{\eta \mid \sigma \subseteq \eta \subseteq \tau\}$ with $\sigma \neq \tau$. This case can also be characterized by the fact that the link of σ is the *closure* of a simplex. The operation is called a *(classical) collapse*.
3. The link of σ is a cone. The operation is called an *(extended) collapse*.

Both classical and extended collapses can be expressed as compositions of elementary collapses. A simplicial complex is said to be *collapsible* if it can be reduced to a single vertex by a finite sequence of collapses (either elementary, classical or extended).

2.4 Filtrations

In the next two sections, we establish collapsibility of certain simplicial complexes using the following strategy. We associate to simplicial complex K a filtration $\{K(t)\}_{t \in \mathbb{R}}$ which is a nested one-parameter family of simplicial complexes such that $K(-\infty) = \emptyset$ and $K(+\infty) = K$. The filtration induces a strict order on simplices of K defined by $\eta \prec_K \nu \Leftrightarrow \{\eta \in K(t_i)$ and $\nu \in K(t_j) \setminus K(t_i)$ for some $t_i < t_j\}$. Simply put, as time t goes by, if a simplex η shows up in $K(t)$ strictly before another simplex ν, then $\eta \prec_K \nu$. As we continuously increases the parameter t from $-\infty$ to $+\infty$, the simplicial complex $K(t)$ changes only at finitely many values of t for which the set of simplices $\Delta(t) = K(t) \setminus \lim_{u \to t^-} K(u)$ is non-empty, where $\lim_{u \to t^-} K(u)$ designates the limit of $K(u)$ as u approaches t from below. Let t_0 be the first time at which a non-empty simplicial complex appears in the filtration. In other words, $K(t_0)$

is the smallest non-empty simplicial complex in the filtration. To prove that K is collapsible, it suffices to show that: (1) $K(t_0)$ is collapsible; (2) the operation that removes $\Delta(t)$ from $K(t)$ is a collapse for all $t > t_0$. For this purpose, it will be crucial to build filtrations that are simple (see Figure 1):

▶ **Definition 1.** *The filtration $\{K(t)\}_{t\in\mathbb{R}}$ is* simple *if for all $t > t_0$, then $\Delta(t)$ has a unique inclusion-minimal element σ_t.*

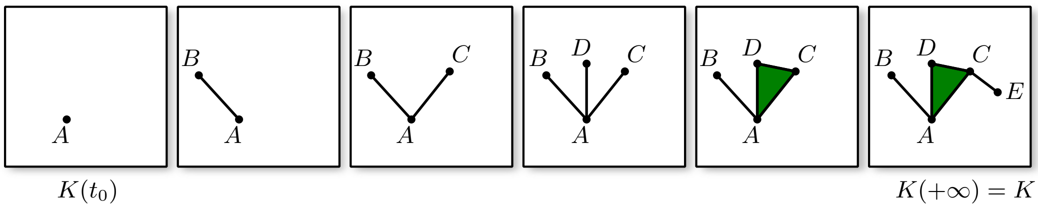

Figure 1 The 6 simplicial complexes form a filtration of K which is simple.

When the filtration is simple, $\Delta(t)$ is precisely the star of σ_t in $K(t)$ and removing $\Delta(t)$ from $K(t)$ is a collapse if one of the three conditions listed in Section 2.3 is satisfied.

▶ **Definition 2.** *Two simplices σ^1 and σ^2 are in* conjunction *in filtration $\{K(t)\}_{t\in\mathbb{R}}$ if they are both inclusion-minimal elements of $\Delta(t)$ for some $t > t_0$ (see Figure 2).*

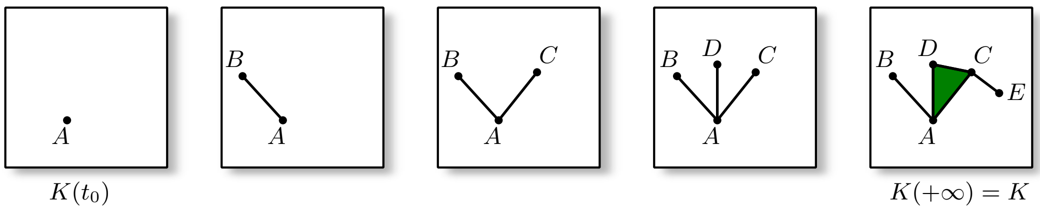

Figure 2 The 5 simplicial complexes form a filtration of K which is not simple as simplices CD and E are in conjunction.

Clearly, a filtration is simple if and only if it contains no pair of simplices in conjunction. The following definition will be useful when, in the next section, we perturb a filtration to make it simple.

▶ **Definition 3.** *Consider a filtration $\{K(t)\}_{t\in\mathbb{R}}$ of K and a filtration $\{L(t)\}_{t\in\mathbb{R}}$ of L. We say that the filtration $\{L(t)\}_{t\in\mathbb{R}}$ is* finer *than the filtration $\{K(t)\}_{t\in\mathbb{R}}$ if for all $t \in \mathbb{R}$, there exists $t' \in \mathbb{R}$ such that $K(t) = L(t')$.*

▶ Remark. If $\{L(t)\}_{t\in\mathbb{R}}$ finer than $\{K(t)\}_{t\in\mathbb{R}}$, then $\eta \prec_K \nu \implies \eta \prec_L \nu$.

3 Collapsing nerves of compact convex sets

3.1 Statement of results

Given a finite collection of sets $\mathcal{C}$, we write $\bigcup \mathcal{C}$ for the union of sets in $\mathcal{C}$ and $\bigcap \mathcal{C}$ for the common intersection of sets in $\mathcal{C}$. The nerve of $\mathcal{C}$ is the abstract simplicial complex that consists of all non-empty subcollections whose sets have a non-empty common intersection,

$$\mathrm{Nrv}\,\mathcal{C} = \{\eta \subseteq \mathcal{C} \mid \bigcap \eta \neq \emptyset\}.$$

The nerve theorem implies that if all sets in $\mathcal{C}$ are compact and convex, then the nerve of $\mathcal{C}$ is homotopy equivalent to the union of sets in $\mathcal{C}$, that is $\operatorname{Nrv}\mathcal{C} \simeq \bigcup \mathcal{C}$. We get immediately that if furthermore $\bigcup \mathcal{C}$ is a non-empty convex set, then $\operatorname{Nrv}\mathcal{C}$ is contractible. In this paper, we prove that under the same hypotheses on $\mathcal{C}$, we have a stronger result, namely, that $\operatorname{Nrv}\mathcal{C}$ is collapsible. Formally:

▶ **Theorem 4.** *If $\mathcal{C}$ is a finite collection of compact convex sets whose union $\bigcup \mathcal{C}$ is a non-empty convex set, then $\operatorname{Nrv}\mathcal{C}$ is collapsible.*

Two corollaries follow immediately. First, we obtain collapsibility of the Delaunay complex. Recall that given a finite point set $S \subseteq \mathbb{R}^d$, the *Voronoi region* of $a \in S$ is $V_a = \{x \in \mathbb{R}^d \mid \|x - a\| \leq \|x - s\|, \forall s \in S\}$ and the *Delaunay complex* of S is the nerve of Voronoi regions, $\operatorname{Del}(S) = \operatorname{Nrv}\{V_s \mid s \in S\}$.

▶ **Corollary 5.** *The Delaunay complex of any finite point set $S \subseteq \mathbb{R}^d$ is collapsible.*

Proof. Apply Theorem 4 to the collection $\mathcal{C} = \{V_s \cap B \mid s \in S\}$, where B is any ball large enough to intersect all common intersection $\bigcap_{s\in\sigma} V_s$ for σ ranging over $\operatorname{Nrv} S$. The result follows because $\operatorname{Del}(S)$ is isomorphic to $\operatorname{Nrv}\mathcal{C}$ which is collapsible by Theorem 4. ◀

To state the second corollary, recall that the Čech complex of a finite point set $S \subseteq \mathbb{R}^d$ with parameter $r \in \mathbb{R}$ is the nerve of the collection of balls, $\mathcal{C}(S, r) = \operatorname{Nrv}\{B(s, r) \mid s \in S\}$ and denote by $\operatorname{Hull}(S)$ the convex hull of S.

▶ **Corollary 6.** *The Čech complex of a finite point set $S \subseteq \mathbb{R}^d$ with parameter $r \in \mathbb{R}$ is collapsible whenever $\operatorname{Hull}(S) \subseteq S^{\oplus r}$.*

Before proving the corollary, notice that the condition on S is tight as for any $\delta > 0$, if $\operatorname{Hull}(S) \subseteq S^{\oplus(r+\delta)}$, then $\mathcal{C}(S, r)$ is not necessarily collapsible. Take for instance the Cech complex of two points at distance $2(r + \delta)$.

Proof. Apply Theorem 4 to the collection $\mathcal{C} = \{B(s, r) \cap \operatorname{Hull}(S) \mid s \in S\}$ and notice that $\operatorname{Nrv}\mathcal{C}$ is isomorphic to $\mathcal{C}(S, r)$. ◀

We prove Theorem 4 by adopting the following strategy. Consider a finite collection $\mathcal{C}$ of compact convex sets whose union $\bigcup \mathcal{C}$ is non-empty and convex. We first build a filtration of $\operatorname{Nrv}\mathcal{C}$ from which we derive a sequence of collapses reducing $\operatorname{Nrv}\mathcal{C}$ to a vertex. The rest of the section is devoted to the proof of Theorem 4. In Section 3.2, we build the filtration and show that the smallest non-empty simplicial complex in the filtration is collapsible. In Section 3.3, we show that we can always perturb the collection $\mathcal{C}$ so that the filtration associated to $\operatorname{Nrv}\mathcal{C}$ enjoys nice properties (in a sense to be made precise). In Section 3.4, we study the filtration of the perturbed collection and show that events through the filtration are collapses.

3.2 Building a filtration

Let $\mathcal{C}$ be a family of subsets of $\mathbb{R}^d$ whose union $\bigcup \mathcal{C}$ is non-empty. Let o be a fix point in $\mathbb{R}^d$ that belongs to $\bigcup \mathcal{C}$. We build a filtration of $\operatorname{Nrv}\mathcal{C}$ by sweeping the space $\mathbb{R}^d$ with a sphere centered at o and whose radius $t \geq 0$ continuously increases from 0 to $+\infty$. Simplicial complexes $K_{o,\mathcal{C}}(t)$ in the filtration are obtained by keeping simplices in $\operatorname{Nrv}\mathcal{C}$ which are subcollections of $\mathcal{C}$ whose common intersections have a distance to o equal to or less than t:

$$K_{o,\mathcal{C}}(t) = \{\eta \subseteq \mathcal{C} \mid d(o, \bigcap \eta) \leq t\}.$$

Clearly, $K_{o,\mathcal{C}}(t) = \emptyset$ for all $t < 0$ and $K_{o,\mathcal{C}}(+\infty) = \operatorname{Nrv}\mathcal{C}$. Notice that $K_{o,\mathcal{C}}(0)$ is non-empty because point o belongs to at least one set C in the collection $\mathcal{C}$ and thus $K_{o,\mathcal{C}}(0)$ contains at least vertex C. It follows that $K_{o,\mathcal{C}}(0)$ is the smallest non-empty simplicial complex in the filtration. As we continuously increases the parameter t from $-\infty$ to $+\infty$, the simplicial complex $K_{o,\mathcal{C}}(t)$ changes only at finitely many values of t for which the set of simplices

$$\Delta_{o,\mathcal{C}}(t) = \{\eta \subseteq \mathcal{C} \mid d(o, \bigcap \eta) = t\}$$

is non-empty. When these events happen, the sphere centered at o with radius t passes through particular points of $\bigcup \mathcal{C}$ that we call *trigger points* and defined below; see Figure 3.

▶ **Definition 7.** *Consider $\eta \subseteq \mathcal{C}$ such that $\bigcap \eta \neq \emptyset$. The* trigger point *of η (with respect to o) is the point of $\bigcap \eta$ at which the smallest distance to o is achieved. There is thus one trigger point per simplex in* $\operatorname{Nrv}\mathcal{C}$ *and the set of all these points is referred to as the trigger points of* $\operatorname{Nrv}\mathcal{C}$ *(with respect to o).*

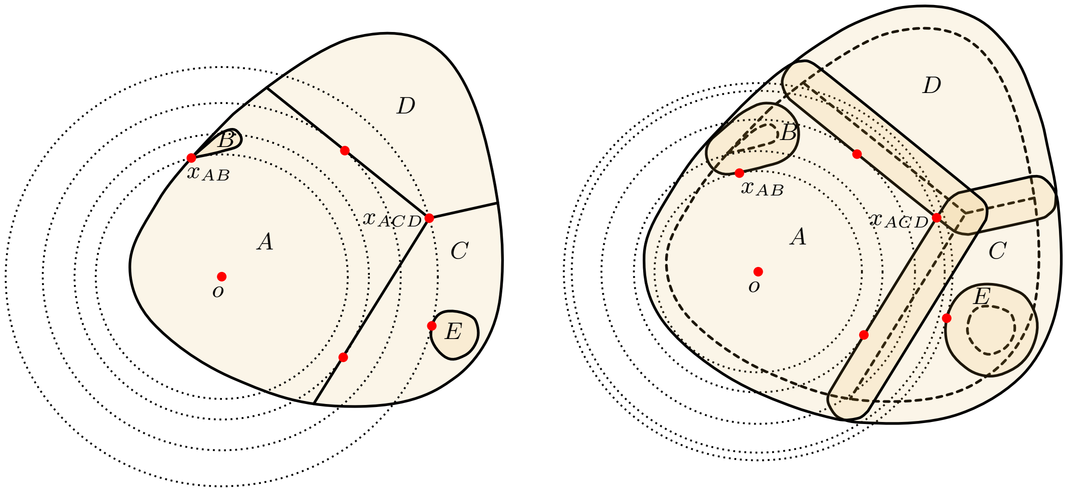

Figure 3 Left: Collection of five convex sets $\{A, B, C, D, E\}$ whose union is convex. The nerve possesses six trigger points (red dots) among which one of them x_{AB} lies on the boundary of the union and another one x_{ACD} does not lie in the interior of any convex set in the collection. The filtration associated to the nerve is depicted in Figure 2 and is not simple. Right: as we offset sets in the collection while keeping the union convex and the nerve unchanged, the trigger point x_{AB} moves in the interior of the union and the trigger point x_{ACD} moves in the interior of at least one convex set (namely A). The associated filtration is depicted in Figure 1 and is simple.

▶ **Lemma 8.** *Let $\mathcal{C}$ be a family of subsets of $\mathbb{R}^d$ whose union $\bigcup \mathcal{C}$ is non-empty and convex and let $o \in \bigcup \mathcal{C}$. Then, $K_{o,\mathcal{C}}(0)$ is collapsible.*

Proof. Let $\tau_0 = \{C \in \mathcal{C} \mid o \in C\}$. The simplicial complex $K_{o,\mathcal{C}}(0) = \{\eta \subseteq \mathcal{C} \mid d(o, \bigcap \eta) = 0\}$ consists of τ_0 together with all its faces and is clearly collapsible. ◀

3.3 Perturbing filtrations

Let $\mathcal{C}$ be a finite family of compact convex sets whose union $\bigcup \mathcal{C}$ is convex and non-empty and pick a point o in $\bigcup \mathcal{C}$. In this section, we establish that it is always possible to perturb the family $\mathcal{C}$ into a family of compact convex sets so as to make the filtration $K_{o,\mathcal{C}}(t)$ simple while leaving $\bigcup \mathcal{C}$ convex and $\operatorname{Nrv}\mathcal{C}$ unchanged. Roughly speaking, we shall perturb sets in the collection by thickening them.

▶ **Lemma 9.** *For all $\delta > 0$, there exists a non-negative map $\alpha : \mathcal{C} \to \mathbb{R}$ bounded above by δ such that if we replace each set $C \in \mathcal{C}$ by set $C^{\oplus\alpha(C)} \cap \bigcup \mathcal{C}$, we perturb $\mathcal{C}$ in such a way that (1) the nerve of $\mathcal{C}$ is left unchanged; (2) the filtration $K_{o,\mathcal{C}}(t)$ becomes simple.*

The proof of Lemma 9 relies on three technical lemmas. To state them, we need some notation and definitions. A *perturbation* of $\mathcal{C}$ is a map $f : \mathcal{C} \to \mathcal{P}(\mathbb{R}^d)$, where $\mathcal{P}(\mathbb{R}^d)$ designates the power set of $\mathbb{R}^d$. For a subcollection $\eta \subseteq \mathcal{C}$ and a simplicial complex L over $\mathcal{C}$, we write $f(\eta) = \{f(C) \mid C \in \eta\}$ and $f(L) = \{f(\eta) \mid \eta \in L\}$. For any non-negative map $\alpha : \mathcal{C} \to \mathbb{R}$, we write $\eta^{\oplus\alpha} = \{C^{\oplus\alpha(C)} \mid C \in \eta\}$ and $L^{\oplus\alpha} = \{\eta^{\oplus\alpha} \mid \eta \in L\}$. A simple compactness argument implies the first lemma of which the proof is omitted.

▶ **Lemma 10.** *Let $\mathcal{X}$ be a finite collection of compact sets. There exists $\varepsilon > 0$ such that for all non-negative maps $\xi : \mathcal{X} \to \mathbb{R}$ bounded above by ε, we have $(\operatorname{Nrv} \mathcal{X})^{\oplus\xi} = \operatorname{Nrv}(\mathcal{X}^{\oplus\xi})$.*

▶ **Lemma 11.** *There exists $\varepsilon > 0$ such that for all non-negative maps $\alpha : \mathcal{C} \to \mathbb{R}$ bounded above by ε and all $0 \leq \beta \leq \varepsilon$, if we replace each set $C \in \mathcal{C}$ by set $C^{\oplus\alpha(C)} \cap (\bigcup \mathcal{C})^{\oplus\beta}$, we perturb $\mathcal{C}$ in such a way that (1) the nerve of $\mathcal{C}$ is left unchanged; (2) the filtration $K_{o,\mathcal{C}}(t)$ after perturbation is finer than what it was before perturbation.*

Proof. Consider a non-negative map α bounded above by ε and $0 \leq \beta \leq \varepsilon$. Write $f_{\alpha,\beta}(C) = C^{\oplus\alpha(C)} \cap (\bigcup \mathcal{C})^{\oplus\beta}$. To prove (1), we apply Lemma 10 to the set $\mathcal{X} = \mathcal{C} \cup \{\bigcup \mathcal{C}\}$ and the map ξ defined by $\xi(C) = \alpha(C)$ for all $C \in \mathcal{C}$ and $\xi(\bigcup \mathcal{C}) = \beta$. We know that for $\varepsilon > 0$ small enough, $\operatorname{Nrv} \mathcal{X} = \operatorname{Nrv}(\mathcal{X}^{\oplus\xi})$ or equivalently $\bigcap \eta \neq \emptyset \Leftrightarrow \bigcap f_{\alpha,\beta}(\eta) \neq \emptyset$ for all $\eta \subseteq \mathcal{C}$, showing that (1) holds. To prove (2), write

$$K(t) = \{\eta \subseteq \mathcal{C} \mid d(o, \bigcap \eta) \leq t\}$$
$$L_{\alpha,\beta}(t) = \{\eta \subseteq \mathcal{C} \mid d(o, \bigcap f_{\alpha,\beta}(\eta)) \leq t\}$$

and let us establish that $L_{\alpha,\beta}(t)$ is finer than $K(t)$. As we continuously increases the parameter t from $-\infty$ to $+\infty$, the simplicial complex $K(t)$ changes only at finitely many values $0 = t_0 < \ldots < t_m$ of t. Let $B_i = B(o, t_i)$ for $i \in \{0, \ldots, m\}$. Let us apply Lemma 10 to the set $\mathcal{X} = \mathcal{C} \cup \{\bigcup \mathcal{C}\} \cup \{B_0, \ldots, B_m\}$ and the map ξ defined by $\xi(C) = \alpha(C)$ for all $C \in \mathcal{C}$, $\xi(\bigcup \mathcal{C}) = \beta$ and $\xi(B_i) = 0$. Lemma 10 implies that for $\varepsilon > 0$ small enough, we have the following five equivalences: $\eta \in K(t_i) \Leftrightarrow d(o, \bigcap \eta) \leq t_i \Leftrightarrow B_i \cap \bigcap \eta \neq \emptyset \Leftrightarrow B_i \cap \bigcap f_{\alpha,\beta}(\eta) \neq \emptyset \Leftrightarrow d(o, \bigcap f_{\alpha,\beta}(\eta)) \leq t_i \Leftrightarrow \eta \in L_{\alpha,\beta}(t_i)$. Thus, for all $t \geq 0$, there exists i such that $K(t) = K(t_i) = L(t_i)$ and $L_{\alpha,\beta}(t)$ is finer than $K(t)$. ◀

Next lemma ensures that when we replace one of the set C in $\mathcal{C}$ by $C^{\oplus\varepsilon} \cap \bigcup \mathcal{C}$, some of the events in the corresponding filtration happen at an earlier time. Precisely:

▶ **Lemma 12.** *Let $C \in \mathcal{C}$ and $\eta \subseteq \mathcal{C}$ such that $C \notin \eta$. Suppose $\{C\} \cup \eta$ is inclusion-minimal in $\Delta_{o,\mathcal{C}}(t)$. Then, for all $t > d(o, \bigcup \mathcal{C})$ and all $\varepsilon > 0$, we have $d(o, C^{\oplus\varepsilon} \cap \bigcap \eta \cap \bigcup \mathcal{C}) < t$.*

Proof. By assumption, $d(o, C \cap \bigcap \eta) = t$. If $\eta \neq \emptyset$, then η is a proper subset of $\{C\} \cup \eta$ and the minimality of $\{C\} \cup \eta$ in $\Delta_{o,\mathcal{C}}(t)$ implies that $d(o, \bigcap \eta) < t$ and therefore $d(o, \bigcap \eta \cap \bigcup \mathcal{C}) < t$. If $\eta = \emptyset$, observe that the later inequality holds because we have assumed that $d(o, \bigcup \mathcal{C}) < t$. Let x be the point in $C \cap \bigcap \eta$ closest to o and let x' be the point in $\bigcap \eta \cap \bigcup \mathcal{C}$ closest to o. As we go from x to x' on the segment connecting x to x', the distance to o decreases in a sufficiently small neighborhood of x while we remain in both sets $C^{\oplus\varepsilon}$ and $\bigcap \eta \cap \bigcup \mathcal{C}$. Thus, $d(o, C^{\oplus\varepsilon} \cap \bigcap \eta \cap \bigcup \mathcal{C}) < t$. ◀

Before proving Lemma 9, recall that two simplices σ^1 and σ^2 are in conjunction in filtration $K_{o,\mathcal{C}}(t)$ if they are both inclusion-minimal elements of $\Delta_{o,\mathcal{C}}(t)$ for some $t > 0$.

Proof of Lemma 9. The proof consists in applying a sequence of elementary perturbations to set $\mathcal{C}$ while preserving $\operatorname{Nrv} \mathcal{C}$ and $\bigcup \mathcal{C}$ until no two simplices remain in conjunction in the filtration $K_{o,\mathcal{C}}(t)$. Suppose two simplices are in conjunction, say σ^1 and σ^2. Suppose $C \in \sigma^1$ and $C \notin \sigma^2$ and replace the convex set C by $C^{\oplus\varepsilon} \cap \bigcup \mathcal{C}$. Clearly, $\mathcal{C}$ is still a collection of compact convex sets, $\bigcup \mathcal{C}$ is left unchanged by the operation and Lemma 11 implies that for $\varepsilon > 0$ small enough, the nerve of $\mathcal{C}$ is also left unchanged. We prove below that for $\varepsilon > 0$ small enough: (1) σ^1 and σ^2 are not in conjunction anymore after the operation; (2) two simplices η and ν are not in conjunction after the operation unless a face $\eta' \subseteq \eta$ and a face $\nu' \subseteq \nu$ were already in conjunction before the operation. Introduce the map $f : \mathcal{C} \to \mathcal{P}(\mathbb{R}^d)$ defined by $f(C) = C^{\oplus\varepsilon} \cap \bigcup \mathcal{C}$ and $f(C') = C'$ for all $C' \neq C$.

(1) Let us prove that for $\varepsilon > 0$ small enough, $f(\sigma^1)$ and $f(\sigma^2)$ are not in conjunction in $K_{o,f(\mathcal{C})}(t)$. Let $t = d(o, \bigcap \sigma^1) = d(o, \bigcap \sigma^2)$ and suppose $\sigma^1 = \{C\} \cup \eta$ with $C \notin \eta$. By construction, $f(\sigma^1) = C^{\oplus\varepsilon} \cap \bigcap \eta \cap \bigcup \mathcal{C}$. By Lemma 12, we have $d(o, \bigcap f(\sigma^1)) < t = d(o, \bigcap f(\sigma^2))$, showing that $f(\sigma^1)$ and $f(\sigma^2)$ are not in conjunction in $K_{o,f(\mathcal{C})}(t)$.

(2) Let us prove that for $\varepsilon > 0$ small enough, two simplices $f(\eta)$ and $f(\nu)$ cannot be in conjunction in $K_{o,f(\mathcal{C})}(t)$ unless a face $\eta' \subseteq \eta$ and a face $\nu' \subseteq \nu$ are in conjunction in $K_{o,\mathcal{C}}(t)$. Consider two simplices $f(\eta)$ and $f(\nu)$ in conjunction in $K_{o,f(\mathcal{C})}(t)$. By Lemma 11, the filtration $K_{o,f(\mathcal{C})}(t)$ is finer than the filtration $f(K_{o,\mathcal{C}}(t))$ and the remark in Section 2.4 entails the implication: $d(o, \bigcap f(\eta)) = d(o, \bigcap f(\nu)) \implies d(o, \bigcap \eta) = d(o, \bigcap \nu)$. The result hence follows.

To remove all pair of simplices in conjunction, consider the partial order $\prec$ on pair of simplices defined by $(\nu', \eta') \prec (\nu, \eta)$ if $\dim \nu' \leq \dim \nu$ and $\dim \eta' \leq \dim \eta$. Sort all pair of simplices in conjunction according to a total order compatible with this partial order. Take the smallest pair (σ^1, σ^2) and apply the elementary perturbation described above. Notice that the operation does not create any new pair of simplices in conjunction smaller than (σ^1, σ^2). By repeating this operation a finite number of times, we thus get a new collection of compact convex sets possessing the same nerve and the same union but for which no two simplices are in conjunction anymore. In other words, the filtration associated to the new collection is simple. Since elementary perturbations can be made as small as wanted, their composition can be made smaller than any given $\delta > 0$. ◀

3.4 Studying events in the filtration

Throughout this section, $\mathcal{C}$ designates a finite collection of compact convex sets whose union is a non-empty convex set and o designates a point in $\bigcup \mathcal{C}$. In this section, we establish Theorem 4. Let us start with a few general observations. Consider $t > 0$ such that $\Delta_{o,\mathcal{C}}(t) \neq \emptyset$. If we assume $\{K_{o,\mathcal{C}}(t)\}_{t\in\mathbb{R}}$ to be simple, then by definition $\Delta_{o,\mathcal{C}}(t)$ has a unique inclusion-minimal element σ_t. Let p_t be the point in $\bigcap \sigma_t$ closest to o and let $\tau_t = \{C \in \mathcal{C} \mid p_t \in C\}$.

▶ **Lemma 13.** *If $\{K_{o,\mathcal{C}}(t)\}_{t\in\mathbb{R}}$ is simple, then $\Delta_{o,\mathcal{C}}(t) = \{\eta \subseteq \mathcal{C} \mid \sigma_t \subseteq \eta \subseteq \tau_t\}$.*

Proof. Let us prove that for all $\eta \subseteq \mathcal{C}$, we have the equivalence

$$\sigma_t \subseteq \eta \subseteq \tau_t \iff d(o, \bigcap \eta) = t.$$

Consider first η such that $\sigma_t \subseteq \eta \subseteq \tau_t$. We have the following sequence of inclusions $p_t \in \bigcap \tau_t \subseteq \bigcap \eta \subseteq \bigcap \sigma_t$. Hence,

$$t = \|p_t - o\| \geq d(o, \bigcap \tau_t) \geq d(o, \bigcap \eta) \geq d(o, \bigcap \sigma_t) = t$$

showing that $d(o, \bigcap \eta) = t$. Conversely, consider $\eta \subseteq \mathcal{C}$ such that $d(o, \bigcap \eta) = t$ and let x be the point of $\bigcap \eta$ closest to o. Observe that $\sigma_t \subseteq \eta$ because $\eta \in \Delta_{o,\mathcal{C}}(t)$ and therefore $x \in \bigcap \eta \subseteq \bigcap \sigma_t$. It follows that $\bigcap \sigma_t$ contains two points x and p_t such that $\|x - o\| = \|p_t - o\| = t = d(o, \bigcap \sigma_t)$. Because $\bigcap \sigma_t$ is convex, there is a unique point in $\bigcap \sigma_t$ at which the smallest distance to o is achieved, showing that $p_t = x$ and therefore $\eta \subseteq \tau_t$. ◀

Hence, $\Delta_{o,\mathcal{C}}(t)$ has a unique inclusion-minimal element σ_t and a unique inclusion-maximal element τ_t. More precisely, $\Delta_{o,\mathcal{C}}(t)$ consists of all cofaces of σ_t in $K_{o,\mathcal{C}}(t)$ and these cofaces are all faces of τ_t. To prove that removing $\Delta_{o,\mathcal{C}}(t)$ from $K_{o,\mathcal{C}}(t)$ is a collapse, it suffices to establish that $\tau_t \neq \sigma_t$. Lemma 14 below shows that p_t lies on the boundary of all the convex sets in σ_t.

▶ **Lemma 14.** *Consider $\eta \in \operatorname{Nrv} \mathcal{C}$ and let x be the point in $\bigcap \eta$ closest to o. Suppose $o \neq x$. If x lies in the interior of some C, then x is also the point in $\bigcap(\eta \setminus \{C\})$ closest to o.*

Proof. Suppose that x lies in the interior of some $C \in \eta$; see Figure 4, left. Because $o \neq x$, we cannot have $\eta = \{C\}$ and therefore $\eta' = \eta \setminus \{C\}$ is non-empty. Let us prove that x is the point of $\bigcap \eta'$ closest to o. Suppose for a contradiction that there exists a point x' in $\bigcap \eta'$ closer to o than x. Since the map $m \mapsto \|m - o\|$ is convex and $\|x - o\| > \|x' - o\|$, the distance to o would be decreasing along the segment $[x, x']$ in the vicinity of x and since this segment, in the vicinity of x, is contained in $\bigcap \eta$ this would contradict the fact that x is the closest point to o in $\bigcap \eta$. ◀

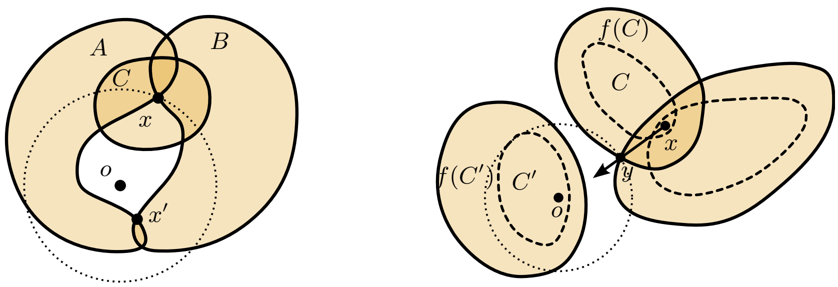

Figure 4 Left: Counterexample to Lemma 14 when sets in $\mathcal{C}$ are not convex. Right: Notation for the proof of Lemma 15.

If we were able to prove that p_t belongs to the interior of one of the convex sets of τ_t, we would be done because we would be sure that $\tau_t \neq \sigma_t$. Unfortunately, this is not true in general (see point x_{ACD} in Figure 3 for a counterexample) but becomes true if we slightly perturb $\mathcal{C}$, as explained in Lemma 15.

▶ **Lemma 15.** *Let $f : \mathcal{C} \to \mathcal{P}(\mathbb{R}^d)$ be a map such that for all $C \in \mathcal{C}$, the subset $f(C)$ is convex, compact and contains C in its interior. Suppose $\operatorname{Nrv} f(\mathcal{C}) = f(\operatorname{Nrv} \mathcal{C})$. Suppose furthermore that all trigger points of $\operatorname{Nrv} f(\mathcal{C})$ lie in the interior of $\bigcup f(\mathcal{C})$. Let y be one of those trigger points. If $y \neq o$, then y lies in the interior of some $f(C) \in f(\mathcal{C})$.*

Proof. Let $\tau = \{C \in \mathcal{C} \mid y \in f(C)\}$. Suppose for a contradiction that $y \neq o$ and that y lies on the boundary of $f(C)$ for all $C \in \tau$; see Figure 4, right. Since $\operatorname{Nrv} f(\mathcal{C}) = f(\operatorname{Nrv} \mathcal{C})$, we have that $y \in \bigcap f(\tau) \neq \emptyset$ implies that $\bigcap \tau \neq \emptyset$. Let x be a point in the latter intersection. For all $C \in \tau$, we have that x belongs to the interior of $f(C)$ while y belongs to the boundary of $f(C)$. Thus, the vector $y - x$ points outward $f(C)$ at y for all $C \in \tau$. Since all convex sets $f(C')$ for which C' not in τ are at some positive distance from y, it follows that, on the segment starting from y in the direction $y - x$, points sufficiently close to y do not belong to $\bigcup f(\mathcal{C})$. In other words, $y \in \partial \bigcup f(\mathcal{C})$. But, y being a trigger point, y lies in the interior of $\bigcup f(\mathcal{C})$, yielding a contradiction. ◀

Next lemma gives an example of perturbation f of $\mathcal{C}$ which ensures that after perturbing $\mathcal{C}$ with f, all trigger points of $\operatorname{Nrv} f(\mathcal{C})$ are in the interior of $\bigcup f(\mathcal{C})$; see Figure 3.

▶ **Lemma 16.** *Consider $\varepsilon > 0$ and a map $\alpha : \mathcal{C} \to \mathbb{R}$ such that $0 \leq \alpha(C) \leq (\sqrt{2} - 1)\varepsilon$. Let $f : \mathcal{C} \to \mathcal{P}(\mathbb{R}^d)$ defined by $f(C) = C^{\oplus(\varepsilon+\alpha(C))} \cap (\bigcup \mathcal{C})^{\oplus\varepsilon}$. Suppose $o \in \bigcup \mathcal{C}$ and $\operatorname{Nrv} f(\mathcal{C}) = f(\operatorname{Nrv} \mathcal{C})$. Then, all trigger points of $\operatorname{Nrv} f(\mathcal{C})$ lie in the interior of $\bigcup f(\mathcal{C})$.*

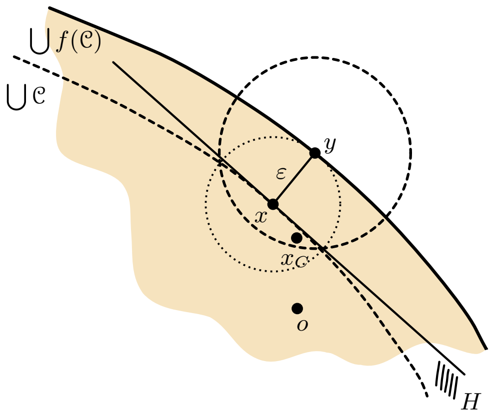

Figure 5 Notation for the proof of Lemma 16.

Proof. Consider $\eta \subseteq \mathcal{C}$ such that $\bigcap f(\eta) \neq \emptyset$ and suppose for a contradiction that the trigger point y of $f(\eta)$ lies on the boundary of $\bigcup f(\mathcal{C})$; see Figure 5. Let x be the point in $\bigcup \mathcal{C}$ closest to y and notice that (1) $\|x - y\| = \varepsilon$ because $\bigcup f(\mathcal{C}) = (\bigcup \mathcal{C})^{\oplus\varepsilon}$ and (2) $x \in \partial \bigcup \mathcal{C}$. We claim that $x \in \bigcap f(\eta)$. Let H be the closed half-space which contains $\bigcup \mathcal{C}$ and whose boundary passes through x while being orthogonal to the vector $y - x$. By construction, for all $C \in \eta$, there is a point x_C in $\bigcup \mathcal{C}$ whose distance to y is equal to or less than $\varepsilon + \alpha(C) \leq \sqrt{2}\varepsilon$. Thus, $x_C \in H \cap B(y, \sqrt{2}\varepsilon) \subseteq B(x, \varepsilon)$. Hence, $\|x - x_C\| \leq \varepsilon$ and therefore x belongs to $f(C)$ for all $C \in \eta$, showing that $x \in \bigcap f(\eta)$ as claimed. Note that $o \in \bigcup \mathcal{C} \subseteq H$. We thus have three points o, x and y such that $y \notin H$, x is the orthogonal projection of y onto H and $o \in H$. It follows that $\|x - o\| \leq \|y - o\|$ and since $x \in \bigcap f(\eta)$, this contradicts the fact that y is the trigger point of η, that is, the point of $\bigcap f(\eta)$ closest to o. ◀

We are now ready to prove Theorem 4.

Proof of Theorem 4. Let $o \in \bigcup \mathcal{C}$. By Lemma 10, for $\varepsilon > 0$ small enough, $\operatorname{Nrv}(\mathcal{C}^{\oplus\varepsilon}) = (\operatorname{Nrv} \mathcal{C})^{\oplus\varepsilon}$. Clearly, $\mathcal{C}^{\oplus\varepsilon}$ is a collection of compact convex sets whose union is convex since $\bigcup(\mathcal{C}^{\oplus\varepsilon}) = (\bigcup \mathcal{C})^{\oplus\varepsilon}$. Applying Lemma 9 to the collection $\mathcal{C}^{\oplus\varepsilon}$, we obtain the existence of a perturbation $f : \mathcal{C} \to \mathcal{P}(\mathbb{R}^d)$ of the form $f(C) = C^{\oplus(\varepsilon+\alpha(C))} \cap (\bigcup \mathcal{C})^{\oplus\varepsilon}$ where $\alpha : \mathcal{C} \to \mathbb{R}$ is a non-negative map such that (1) $\operatorname{Nrv} f(\mathcal{C}) = f(\operatorname{Nrv} \mathcal{C})$ and (2) the filtration $\{K_{o,f(\mathcal{C})}(t)\}_{t\geq 0}$ is simple. Furthermore, the map α can be chosen arbitrarily small and in particular bounded above by $(\sqrt{2} - 1)\varepsilon$. Observe that such an f satisfies the assumptions of Lemma 16 and therefore of Lemma 15. Letting $\mathcal{D} = f(\mathcal{C})$, we prove in the following paragraph that $\operatorname{Nrv} \mathcal{D}$ is collapsible which entails immediately that $\operatorname{Nrv} \mathcal{C}$ is collapsible as the two are isomorphic.

By Lemma 8, $K_{o,\mathcal{D}}(0)$ is collapsible. Let us prove that for all $t > 0$ such that $\Delta_{o,\mathcal{D}}(t) \neq \emptyset$, the operation that removes $\Delta_{o,\mathcal{D}}(t)$ from $K_{o,\mathcal{D}}(t)$ is a collapse. Since the filtration $\{K_{o,\mathcal{D}}(t)\}_{t\geq 0}$ is simple, $\Delta_{o,\mathcal{D}}(t)$ has a unique inclusion-minimal element σ_t. Let p_t be the point in $\bigcap \sigma_t$ closest to o and let $\tau_t = \{D \in \mathcal{D} \mid p_t \in D\}$. By Lemma 13, $\Delta_{o,\mathcal{D}}(t)$ is the set of simplices $\eta \in \operatorname{Nrv} \mathcal{D}$ such that $\sigma_t \subseteq \eta \subseteq \tau_t$. Let us show that $\sigma_t \neq \tau_t$. By Lemma 14, p_t lies on the boundary of all sets in σ_t. By Lemma 15, p_t lies in the interior of at least one set D in τ_t. Hence, $\tau_t \neq \sigma_t$ and the operation that removes $\Delta_{o,\mathcal{D}}(t)$ from $K_{o,\mathcal{D}}(t)$ is a collapse. ◀

4 Collapsing Rips complexes

In this section, we turn our attention to Rips complexes. Given a point set S and a scale parameter r, the Rips complex $\mathcal{R}(S, r)$ is the simplicial complex whose simplices are subsets of points in S with diameter at most $2r$. Rips complexes are examples of flag completions. Recall that the flag completion of a graph G, denoted $\operatorname{Flag}(G)$, is the maximal simplicial complex whose 1-skeleton is G. Let $G(S, r)$ denote the graph whose vertices are the points S and whose edges connect all pairs of points within distance $2r$. The Rips complex of S with parameter r is $\mathcal{R}(S, r) = \operatorname{Flag}(G(S, r))$. It is the largest simplicial complex sharing with the Čech complex $\mathcal{C}(S, r)$ the same 1-skeleton. However, the Rips complex has the computational advantage over the Čech complex to be a flag completion: it suffices to compute its 1-skeleton to encode the whole complex. In this section, we prove that if S samples sufficiently densely $\operatorname{Hull}(S)$, then the Rips complex is collapsible for a suitable value of the scale parameter.

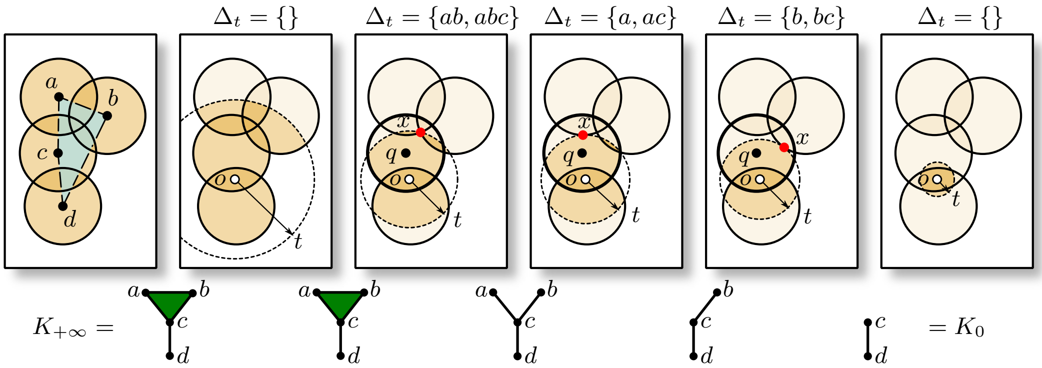

Figure 6 From left to right: Our proof technique (illustrated here when $S = \{a, b, c, d\}$) consists in sweeping space with a sphere centered at $o \in \operatorname{Hull}(S)$ and whose radius t continuously decreases from $+\infty$ to 0. We deduce from the sweep a sequence of collapses reducing $\mathcal{R}(S, \alpha)$ to a vertex.

▶ **Theorem 17.** *Let $S \neq \emptyset$ be a finite set of points in $\mathbb{R}^d$ and $r > 0$. If $\operatorname{Hull}(S) \subseteq S^{\oplus(2-\sqrt{3})r}$, then there exists a sequence of extended collapses reducing $\mathcal{R}(S, r)$ to a vertex in such a way that the result of each extended collapse is a flag complex.*

Proof. Set $r = 1$ and write B_x for the closed unit ball centered at x. Fix a point o in the convex hull of S. We construct a sequence of collapses by sweeping the space with a sphere centered at o and whose radius $t \geq 0$ continuously decreases from $+\infty$ to 0. Specifically, let G_t be the graph whose vertices are points $s \in S$ such that $B(o, t) \cap B_s \neq \emptyset$ and whose edges connect all pair of points $a, b \in S$ such that $B(o, t) \cap B_a \cap B_b \neq \emptyset$. Let $K_t = \operatorname{Flag} G_t$. Clearly, $G_{+\infty} = G(S, 1)$ and $K_{+\infty} = \operatorname{Flag} G_{+\infty} = \mathcal{R}(S, 1)$. We claim that K_0 is collapsible. Indeed, the vertex set of K_0 is the set of points $\tau_0 = \{s \in S \mid o \in B_s\} = S \cap B_o$ which is non-empty since $o \in \operatorname{Hull}(S) \subseteq S^{\oplus 1}$. It follows that $K_0 = \operatorname{Flag} G_0$ consists of τ_0 and all its faces and is collapsible. As we continuously decreases t from $+\infty$ to 0, changes in the simplicial complex K_t occur whenever a vertex or an edge disappears from the graph G_t; see Figure 6. Generically, we may assume that these events do not happen simultaneously.

When a vertex a disappears from G_t at time t, the intersection $B(o,t) \cap B_a$ reduces to a single point x; see Figure 7, left. In this situation, we claim that the link of a in K_t is the closure of the simplex $\tau_x = S \cap B_x \setminus \{a\}$. First, note that τ_x is non-empty since x lies on the segment connecting o to a and therefore belongs to the convex hull of S which is contained in $S^{\oplus 2-\sqrt{3}}$. Hence, there is a point $s \in S$ in the interior of B_x and $\tau_x \neq \emptyset$. Furthermore, τ_x is precisely the vertex set of the link since an edge au belongs to K_t if and only if $B(o,t) \cap B_a \cap B_u \neq \emptyset$ with $u \in S \setminus \{a\}$ which can be reformulated as $u \in S \cap B_x \setminus \{a\}$. Finally, any two vertices u and v in the link are connected by an edge since clearly $u, v \in \tau_x$ implies $B(o,t) \cap B_u \cap B_v \supset \{x\}$. We have just proved that the link of a in K_t is the closure of a simplex. Thus, removing the star of a from K_t is a classical collapse.

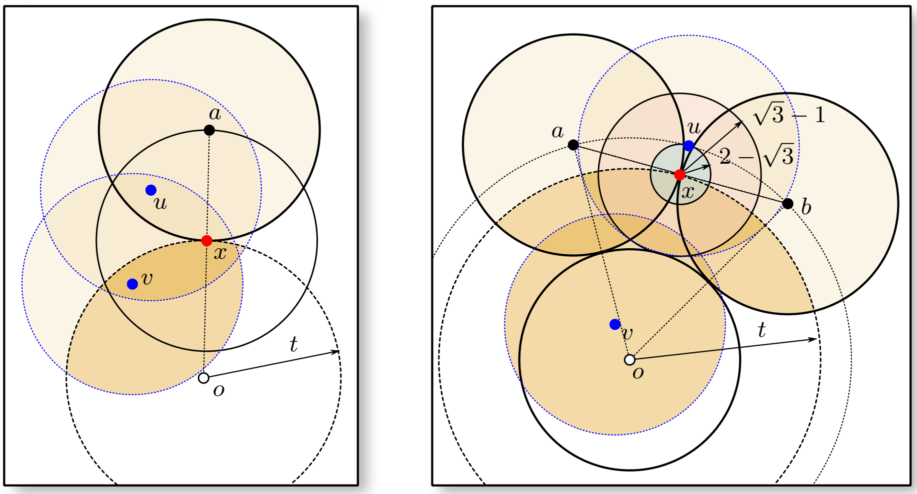

Figure 7 Notation for the proof of Theorem 17. Two kinds of events may occur: either a vertex collapse (on the left) or an edge collapse (on the right). The edge collapse is illustrated when triangle oab is equilateral.

When an edge ab disappears from G_t at time t, there exists a point x such that $\{x\} = B(o,t) \cap B_a \cap B_b$; see Figure 7, right. Note that x lies in the convex hull of $\{a,b,o\}$ and therefore lies in the convex hull of S. Since $\text{Hull}(S) \subseteq S^{\oplus 2-\sqrt{3}}$, there exists $u \in S$ such that $\|u - x\| \leq 2 - \sqrt{3}$. In particular, $x \in B_u$ which ensures that both $x \in B(o,t) \cap B_a \cap B_u \neq \emptyset$ and $x \in B(o,t) \cap B_b \cap B_u \neq \emptyset$ and therefore u belongs to the link of ab in K_t. We claim that the link of ab is a cone with apex u. Consider a point $v \in S$ which belongs to the link of ab in K_t. Equivalently, both $B(o,t) \cap B_a \cap B_v \neq \emptyset$ and $B(o,t) \cap B_b \cap B_v \neq \emptyset$ and Lemma 19 implies that $B(o,t) \cap B(x, \sqrt{3}-1) \cap B_v \neq \emptyset$. Since $\|u - x\| \leq 2 - \sqrt{3}$, we have $B(x, \sqrt{3}-1) \subseteq B_u$ and therefore we also have $B(o,t) \cap B_u \cap B_v \neq \emptyset$, showing that uv also belongs to the link of ab in K_t. We have just proved that the link of ab in K_t is a cone. Thus, removing the star of ab from K_t is an extended collapse. ◀

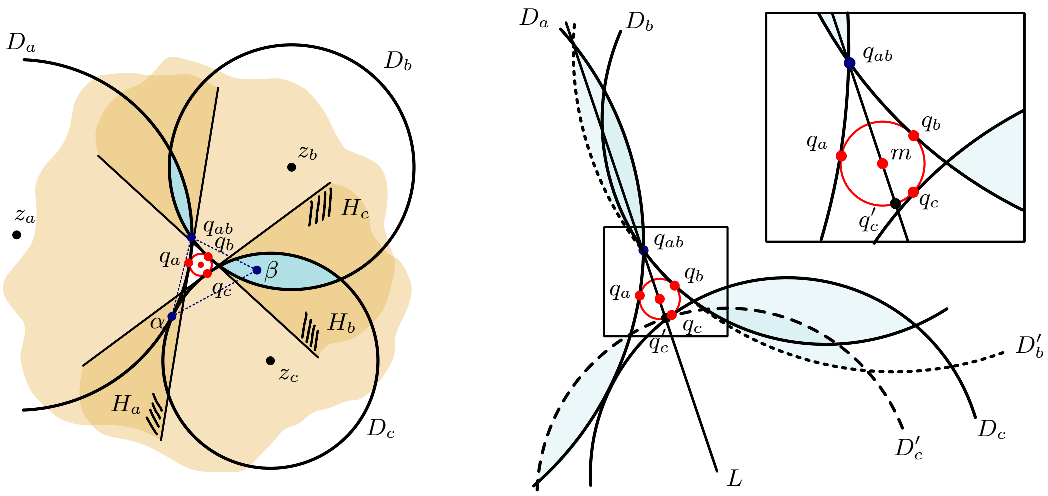

Figure 8 Notation for the proof of Lemma 18.

4.1 Two geometric lemmas

The proof of Theorem 17 relies on two geometric lemmas. The first one states facts about three disks in the plane that intersect pairwise but have no common intersection (Lemma 18). It will allow us to deduce facts about the way four balls intersect in $\mathbb{R}^d$ (Lemma 19). As before, B_x denotes the unit closed ball centered at x.

▶ **Lemma 18.** *Let D_a, D_b and D_c be three disks with radius equal to or less than one and such that any two disks have a non-empty intersection while the three together have no common intersection. Let q_{ab} be the point of $D_a \cap D_b$ closest to the center of D_c. There exists a point $q_c \in D_c$ such that:*

- *for all points $\alpha \in D_a \cap D_c$ and $\beta \in D_b \cap D_c$, the point q_c is in the convex hull of α, β and q_{ab};*
- $\|q_c - q_{ab}\| \leq \sqrt{3} - 1$.

Proof. Consider the disk D_m whose boundary is tangent to the boundaries of the three disks D_a, D_b and D_c and whose interior intersects none of the three disks D_a, D_b and D_c. For $x \in \{a, b, c\}$, the two disks D_x and D_m intersect in a single point q_x; see Figure 8, left. Let $\alpha \in D_a \cap D_c$ and $\beta \in D_b \cap D_c$. We claim that q_c belongs to the convex hull of α, β and q_{ab}. Indeed, for $x \in \{a, b, c\}$, let H_x be the half-plane that contains D_x and avoids the interior of D_m. We have $\alpha \in H_a \cap H_c$, $\beta \in H_b \cap H_c$, and $q_{ab} \in H_a \cap H_b$. The triangle $\alpha\beta q_{ab}$ covers the closure of $\mathbb{R}^2 \setminus (H_a \cup H_b \cup H_c)$ and therefore q_c.

Let us prove that $\|q_c - q_{ab}\| \leq \sqrt{3} - 1$. For $x \in \{a, b, c\}$, we denote the center of D_x by z_x and its radius by ρ_x. We are going to transform the three disks D_a, D_b and D_c in such a way that after the transformation:

(i) the three disks intersect pairwise but have no common intersection;
(ii) the distance between q_c and q_{ab} is at least as large as it was before the transformation;
(iii) $\rho_x \leq 1$ for $x \in \{a, b, c\}$;
(iv) the centers z_a, z_b, and z_c form an equilateral triangle of side length two.

Let q_c' be the point on the boundary of D_m that is farthest away from q_{ab}; see Figure 8, right. Clearly, $\|q_c' - q_{ab}\| \geq \|q_c - q_{ab}\|$. The two tangency points q_a and q_b decompose the boundary of D_m in two arcs and it is not difficult to see that one of them contains both

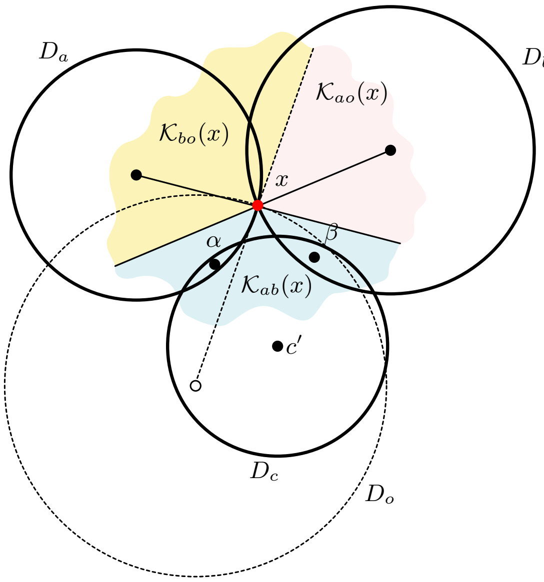

Figure 9 Notation for the proof of Lemma 19.

q_c and q_c'. Consider the disk D_c' obtained by rotating D_c around m until it meets q_c'. As we do so, the rotated disk maintains a contact with at least one of the two disks D_a or D_b. Without loss of generality, we may assume that $D_a \cap D_c' \neq \emptyset$. Let L be the straight-line passing through q_{ab}, q_c' and m. Let D_b' be the symmetric of D_a with respect to L. We have $D_b' \cap D_c' \neq \emptyset$. The two boundaries of D_a and D_b' meet in two points, one of them being q_{ab}. If we replace D_b by D_b' and D_c by D_c', it is easy to check that now the three disks D_a, D_b and D_c satisfies (i), (ii) and (iii) and their centers form an isosceles triangle. We can then further transform the three disks in such a way that after the transformation, they satisfy in addition (iv). When this is the case, we clearly have $\|q_c - q_{ab}\| = \sqrt{3} - 1$. ◀

▶ **Lemma 19.** *Let a and b be two points such that B_a and B_b have a non-empty intersection. Let o be a point such that $d(o, B_a \cap B_b) = t > 0$. Let x be the (unique) point of $B_a \cap B_b$ closest to o. Any unit ball which has a non-empty intersection with both $B_a \cap B(o,t)$ and $B_b \cap B(o,t)$ has a non-empty intersection with $B(x, \sqrt{3}-1) \cap B(o,t)$.*

Proof. Note that $x \in B(o,t)$. Let c such that $B_c \cap B_a \cap B(o,t) \neq \emptyset$ and $B_c \cap B_b \cap B(o,t) \neq \emptyset$. If $x \in B_c$, then the claim holds trivially since $x \in B(x, \sqrt{3}-1) \cap B(o,t)$. Let us assume from now on that $x \notin B_c$. Take $\alpha \in B_c \cap B_a \cap B(o,t)$ and $\beta \in B_c \cap B_b \cap B(o,t)$ and consider a 2-plane Π that contains the three points x, α and β. This 2-plane intersects the four balls B_a, B_b, B_c and $B(o,t)$ in four disks that we denote respectively D_a, D_b, D_c and D_o; see Figure 9. The three disks D_a, D_b and D_o have a non-empty intersection reduced to point x. We have $\alpha \in D_c \cap D_a \cap D_o \neq \emptyset$, $\beta \in D_c \cap D_b \cap D_o \neq \emptyset$ and $x \notin D_c$.

Let z_c be the center of D_c. We claim that x is the point of $D_a \cap D_b$ closest to z_c. Define the *outer cone* of $D_a \cap D_o$ at x as the set of points:

$$\mathcal{K}_{ao}(x) = \{y \in \Pi \mid \forall z \in D_a \cap D_o,\ \langle y - x, z - x \rangle \leq 0\}.$$

Equivalently, $\mathcal{K}_{ao}(x)$ is the set of points whose distance from x is less than or equal to the distance from any other point of $D_a \cap D_o$. The fact that $x \notin D_c$ while $\alpha \in D_c$ implies that $\|z_c - \alpha\| \leq \|z_c - x\|$. Thus, α is a point in the intersection $D_a \cap D_o$ closer to z_c than x. Equivalently, $z_c \notin \mathcal{K}_{ao}(x)$. Similarly, $z_c \notin \mathcal{K}_{bo}(x)$. Since $\mathcal{K}_{ao}(x) \cup \mathcal{K}_{bo}(x) \cup \mathcal{K}_{ab}(x) = \Pi$, it follows that $z_c \in \mathcal{K}_{ab}(x)$. In other words, x is the point of $D_a \cap D_b$ closest to z_c as claimed.

Therefore, Lemma 18 can be applied and shows the existence of a point $x' \in D_c$ in the convex hull of α, β and x such that $x' \in B(x, \sqrt{3} - 1)$. Since all three points α, β and x belong to $B(o, t)$, it follows that $x' \in B(o, t)$, yielding the result. ◀

5 Future work

We envision that our work could create a bridge towards a non-smooth discrete Morse theory. Continuous Morse theory studies how the homology of a smooth manifold is determined by the critical points of a Morse function. Robin Forman has introduced a discrete Morse Theory [7] enjoying similar properties but defined on simplicial complexes. This has been generalized in [8, 4]. The filtrations that we built can be interpreted as defining a Morse function with a single critical event creating a connected component. By removing the convexity assumption on the union of convex sets, or by varying the function defining the filtration (here the distance to o), we get a Morse Theory that may characterize the homology of any non-smooth set that can be expressed as a finite union of convex sets (such as embedded simplical or polyhedral complexes for examples). While this paper does not explore this generalization, it opens the possibility for a non-smooth, discrete Morse Theory.

References

1. 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.
2. 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. `doi:10.1016/j.comgeo.2012.02.009`.
3. Dominique Attali and André Lieutier. Geometry-driven collapses for converting a Čech complex into a triangulation of a nicely triangulable shape. *Discrete & Computational Geometry*, 54(4):798–825, 2015.
4. Ulrich Bauer and Herbert Edelsbrunner. The Morse theory of Čech and Delaunay complexes. *Transactions of the American Mathematical Society*, 369(5):3741–3762, 2017.
5. Bruno Benedetti. Discrete Morse theory for manifolds with boundary. *Transactions of the American Mathematical Society*, 364(12):6631–6670, 2012.
6. Karol Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fundamenta Mathematicae*, 35(1):217–234, 1948. URL: `http://eudml.org/doc/213158`.
7. Robin Forman. Morse theory for cell complexes, 1998.
8. Ragnar Freij. Equivariant discrete Morse theory. *Discrete Mathematics*, 309(12):3821–3829, 2009.
9. A. Hatcher. *Algebraic topology.* Cambridge University Press, 2002.