Gromov Hyperbolicity, Geodesic Defect, and Apparent Pairs in Vietoris–Rips Filtrations

Ulrich Bauer
Department of Mathematics and Munich Data Science Institute,
Technische Universität München, Germany

Fabian Roll
Department of Mathematics, Technische Universität München, Germany

Abstract

Motivated by computational aspects of persistent homology for Vietoris–Rips filtrations, we generalize a result of Eliyahu Rips on the contractibility of Vietoris–Rips complexes of geodesic spaces for a suitable parameter depending on the hyperbolicity of the space. We consider the notion of geodesic defect to extend this result to general metric spaces in a way that is also compatible with the filtration. We further show that for finite tree metrics the Vietoris–Rips complexes collapse to their corresponding subforests. We relate our result to modern computational methods by showing that these collapses are induced by the apparent pairs gradient, which is used as an algorithmic optimization in Ripser, explaining its particularly strong performance on tree-like metric data.

2012 ACM Subject Classification Mathematics of computing → Geometric topology; Mathematics of computing → Trees; Theory of computation → Computational geometry

Keywords and phrases Vietoris–Rips complexes, persistent homology, discrete Morse theory, apparent pairs, hyperbolicity, geodesic defect, Ripser

Digital Object Identifier 10.4230/LIPIcs.SoCG.2022.15

Related Version *Extended Version*: `https://arxiv.org/abs/2112.06781`

Funding This research has been supported by the DFG Collaborative Research Center SFB/TRR 109 *Discretization in Geometry and Dynamics*.

Acknowledgements We thank Michael Bleher, Lukas Hahn, and Andreas Ott for stimulating discussions about applications of Vietoris–Rips complexes to the topological study of coronavirus evolution motivating our interest in the persistence of tree metrics, the organizers and participants of the AATRN Vietoris–Rips online seminar for sparking our interest in the Contractibility Lemma, and the anonymous reviewers for valuable feedback.

1 Introduction

The *Vietoris–Rips* complex is a fundamental construction in algebraic, geometric, and applied topology. For a metric space X and a threshold $t > 0$, it is defined as the simplicial complex consisting of nonempty and finite subsets of X with diameter at most t:

$$\mathrm{Rips}_t(X) = \{\emptyset \neq S \subseteq X \mid S \text{ finite}, \ \operatorname{diam} S \leq t\}.$$

First introduced by Vietoris [27] in order to make homology applicable to general compact metric spaces, it has also found important applications in geometric group theory [16] and topological data analysis [26]. The role of the threshold in these three application areas is notably different. The homology theory defined by Vietoris arises in the limit $t \to 0$. In contrast, the key applications in geometric group theory rely on the fact that the Vietoris–Rips complex of a hyperbolic geodesic space is contractible for a sufficiently large threshold.

This observation, originally due to Rips and first published in Gromov's seminal paper on hyperbolic groups [16], is a fundamental result about the topology of Vietoris–Rips complexes and plays a central role in the theory of hyperbolic groups.

▶ **Lemma 1** (Contractibility Lemma; Rips, Gromov [16]). *Let X be a δ-hyperbolic geodesic metric space. Then the complex $\mathrm{Rips}_t(X)$ is contractible for every $t > 0$ with $t \geq 4\delta$.*

Here, a metric space (X, d) is called *geodesic* if for any two points $x, y \in X$ there exists an isometric map $[0, d(x, y)] \to X$ such that $0 \mapsto x$ and $d(x, y) \mapsto y$, and it is called *δ-hyperbolic* (in the sense of Gromov [16]) for $\delta \geq 0$ if for any four points $w, x, y, z \in X$ we have

$$d(w, x) + d(y, z) \leq \max\{d(w, y) + d(x, z), d(w, z) + d(x, y)\} + 2\delta. \tag{1}$$

Finally, in applications of Vietoris–Rips complexes to topological data analysis, one is typically interested in the *persistent homology* of the entire filtration of complexes for all possible thresholds. A notable difference to the classical applications is that the metric spaces under consideration are typically finite, and in particular not geodesic. This motivates the interest in a meaningful generalization of the Contractibility Lemma to finite metric spaces. Based on the notion of a *discretely geodesic space* defined by Lang [22], which is the natural setting for hyperbolic groups, and motivated by techniques used in that paper, we consider the following quantitative geometric property (called *ν-almost geodesic* in [10, p. 271]).

▶ **Definition 2.** *A metric space X is* ν-geodesic *if for all $x, y \in X$ and $r, s \geq 0$ with $r + s = d(x, y)$ there exists a point $z \in X$ with $d(x, z) \leq r + \nu$ and $d(y, z) \leq s + \nu$. The* geodesic defect *of X, denoted by $\nu(X)$, is the infimum over all ν such that X is ν-geodesic.*

Our first main result is a generalization of the Contractibility Lemma that also applies to non-geodesic spaces using our notion of geodesic defect, and further produces collapses that are compatible with the Vietoris–Rips filtration above the collapsibility threshold.

▶ **Theorem 3.** *Let X be a finite δ-hyperbolic ν-geodesic metric space. Then there exists a discrete gradient that induces, for every $u > t \geq 4\delta + 2\nu$, a sequence of collapses*

$$\mathrm{Rips}_u(X) \searrow \mathrm{Rips}_t(X) \searrow \{*\}.$$

▶ **Example 4.** An important special case is given by a finite *tree metric space* (V, d), where V is the vertex set of a positively weighted tree $T = (V, E)$, and where the edge weights are taken as lengths and d is the associated path length metric, i.e., for two points $x, y \in V$ their distance is the infimum total weight of any path starting in x and ending in y. The geodesic defect is $\nu(V) = \frac{1}{2} \max_{e \in E} l(e)$, where $l(e)$ is the length of the edge e. Moreover, (V, d) is 0-hyperbolic (see [13, Theorems 3.38 and 3.40] for a characterization of 0-hyperbolic spaces).

This example is of particular relevance in the context of evolutionary biology, where persistent Vietoris–Rips homology has been successfully applied to identify recombinations and recurrent mutations [11, 24, 8]. The metrics arising as genetic distances of aligned RNA or DNA sequences are typically very similar to trees, capturing the phylogeny of the evolution. This motivates our interest in the particular case of tree metrics. These metric spaces are known to have acyclic Vietoris–Rips homology in degree > 0, and so any homology is an indication of some evolutionarily relevant phenomenon.

Our second main result is a strengthened version of Theorem 3 for the special case of tree metric spaces that connects the collapses of the Vietoris–Rips filtration to the construction of *apparent pairs*, which play an important role as a computational shortcut in the software

Ripser [6]. This result depends on a particular ordering of the vertices: we say that a total order of V is *compatible* with the tree T if it extends the unique tree partial order resulting from choosing some arbitrary root vertex as the minimal element.

▶ **Theorem 5.** *Let V be a finite tree metric space for a weighted tree $T = (V, E)$, whose vertices are totally ordered in a compatible way. Then the apparent pairs gradient for the lexicographically refined Vietoris–Rips filtration induces a sequence of collapses*

$$\mathrm{Rips}_u(V) \searrow \mathrm{Rips}_t(V) \searrow T_t$$

for every $u > t > 0$ such that no edge $e \in E$ has length $l(e) \in (t, u]$, where T_t is the subforest with vertices V and all edges of E with length at most t. In particular, the persistent homology of the Vietoris–Rips filtration is trivial in degree > 0.

In the special case of trees with unit edge length, the proofs in [1, Proposition 2.2] and [2, Proposition 3] are similar in spirit to our proof of Theorem 25, which is based on discrete Morse theory. Related results about implications of the geometry of a metric space on the homotopy types of the associated Vietoris–Rips complexes can be found in [3, 4, 23].

▶ Remark 6. Given a vertex order $\leq$, the lexicographic order on simplices for the reverse vertex order $\geq$ coincides with the reverse colexicographic order for the original order $\leq$, which is used for computations in Ripser. As a consequence, when the input is a tree metric with the points ordered in reverse order of the distances to some arbitrarily chosen root, then Ripser will identify all non-tree simplices in apparent pairs, requiring not a single column operation to compute its trivial persistent homology. In practice, we observe that on data that is almost tree-like, such as genetic evolution distances, Ripser exhibits exceptionally good computational performance. The results of this paper provide a partial geometric explanation for this behavior and yield a heuristic for preprocessing tree-like data by sorting the points to speed up the computation in such cases. In the application to the study of SARS-CoV-2 described in [8], ordering the genome sequences in reverse chronological order, as an approximation of the reverse tree order for the phylogenetic tree, lead to a huge performance improvement, bringing down the computation time for the persistence barcode from a full day to about 2 minutes.

2 Preliminaries

2.1 Discrete Morse theory and the apparent pairs gradient

A *simplicial complex* K on a vertex set $\mathrm{Vert}\, K$ is a collection of nonempty finite subsets of $\mathrm{Vert}\, K$ such that for any set $\sigma \in K$ and any nonempty subset $\rho \subseteq \sigma$ one has $\rho \in K$. A set $\sigma \in K$ is called a *simplex*, and $\dim \sigma = \mathrm{card}\, \sigma - 1$ is its *dimension*. Moreover, ρ is said to be a *face of* σ and σ a *coface of* ρ. If $\dim \rho = \dim \sigma - 1$, then we call ρ a *facet of* σ and σ a *cofacet of* ρ. The *star of* σ, $\mathrm{St}\, \sigma$, is the set of cofaces of σ in K, and the *closure of* σ, $\mathrm{Cl}\, \sigma$, is the set of its faces. For a subset $E \subseteq K$, we write $\mathrm{St}\, E = \bigcup_{e \in E} \mathrm{St}\, e$.

Generalizing the ideas of Forman [14], a function $f \colon K \to \mathbb{R}$ is a *discrete Morse function* [7, 15] if f is monotonic, i.e., for any $\sigma, \tau \in K$ with $\sigma \subseteq \tau$ we have $f(\sigma) \leq f(\tau)$, and there exists a partition of K into intervals $[\rho, \phi] = \{\psi \in K \mid \rho \subseteq \psi \subseteq \phi\}$ in the face poset such that $f(\sigma) = f(\tau)$ for any $\sigma \subseteq \tau$ if and only if σ and τ belong to a common interval in the partition. The collection of *regular* intervals, $[\rho, \phi]$ with $\rho \neq \phi$, is called the *discrete gradient of* f, and any singleton interval $[\sigma, \sigma]$, as well as the corresponding simplex σ, is called *critical*.

▶ **Proposition 7** (Hersh [18, Lemma 4.1]; Jonsson [20, Lemma 4.2]). *Let K be a finite simplicial complex, and $\{K_\alpha\}_{\alpha \in A}$ a set of subcomplexes covering K, each equipped with a discrete gradient V_α, such that for any simplex of K*

- *there is a unique minimal subcomplex K_α containing that simplex, and*
- *the simplex is critical for the discrete gradients of all other such subcomplexes.*

Then the regular intervals in the V_α are disjoint, and their union is a discrete gradient on K.

An *elementary collapse* $K \searrow K \setminus \{\sigma, \tau\}$ is the removal of a pair of simplices, where σ is a facet of τ, with τ the unique proper coface of σ. A *collapse* $K \searrow L$ onto a subcomplex L is a sequence of elementary collapses starting in K and ending in L. An elementary collapse can be realized continuously by a strong deformation retraction and therefore collapses preserve the homotopy type. A discrete gradient can encode a collapse.

▶ **Proposition 8** (Forman [14]; see also [21, Theorem 10.9]). *Let K be a finite simplicial complex and let $L \subseteq K$ be a subcomplex. Assume that V is a discrete gradient on K such that the complement $K \setminus L$ is the union of intervals in V. Then there exists a collapse $K \searrow L$.*

Let $f\colon K \to \mathbb{R}$ be a monotonic function. Assume that the vertices of K are totally ordered. The *f-lexicographic order* is the total order $\leq_f$ on K given by ordering the simplices

- by their value under f,
- then by dimension,
- then by the lexicographic order induced by the total vertex order.

We call a pair (σ, τ) of simplices in K a *zero persistence pair* if $f(\sigma) = f(\tau)$. An *apparent pair* (σ, τ) with respect to the f-lexicographic order is a pair of simplices in K such that σ is the maximal facet of τ, and τ is the minimal cofacet of σ. The collection of apparent pairs forms a discrete gradient [6, Lemma 3.5], called the *apparent pairs gradient*.

Assume that K is finite and $f\colon K \to \mathbb{R}$ a discrete Morse function with discrete gradient V. Refine V to another discrete gradient

$$\widetilde{V} = \{(\psi \setminus \{v\}, \psi \cup \{v\}) \mid \psi \in [\rho, \phi] \in V,\ v = \min(\phi \setminus \rho)\}$$

by doing a minimal vertex refinement on each interval.

▶ **Lemma 9.** *The zero persistence apparent pairs with respect to the f-lexicographic order are precisely the gradient pairs of $\widetilde{V}$.*

Proof. Let (σ, τ) be a zero persistence apparent pair. Then $f(\sigma) = f(\tau)$, and σ and τ are contained in the same regular interval $I = [\rho, \phi]$ of V. Let v be the minimal vertex in $\phi \setminus \rho$. By assumption, σ is the maximal facet of τ, and τ is the minimal cofacet of σ. Hence, σ is lexicographically maximal among all facets of τ in I, and τ is lexicographically minimal under all cofacets of σ in I. By the assumption that (σ, τ) forms an apparent pair, we cannot have $v \in \sigma$, as otherwise $\tau \setminus \{v\}$ would be a larger facet of τ than σ. Similarly, we cannot have $v \notin \tau$, as otherwise $\sigma \cup \{v\}$ would be a smaller cofacet of σ than τ. This means that $\tau = \sigma \cup \{v\}$ and therefore $\{\sigma, \tau\} \in \widetilde{V}$.

Conversely, assume that $\{\sigma, \tau\} \in \widetilde{V}$ holds. Consider the interval $I = [\rho, \phi]$ of V with $\{\sigma, \tau\} \subseteq I$ and let v be the minimal vertex in $\phi \setminus \rho$. By construction of $\widetilde{V}$, $\sigma = \tau \setminus \{v\}$ is the lexicographically maximal facet of τ in I and $\tau = \sigma \cup \{v\}$ is the lexicographically minimal cofacet of σ in I. Therefore, (σ, τ) is a zero persistence apparent pair. ◀

2.2 Rips' Contractibility Lemma via the injective hull

In this section, we recall some known facts about embeddings of metric spaces into their injective hull. We adapt these results using our notion of geodesic defect to prove a version of the Contractibility Lemma for finite δ-hyperbolic ν-geodesic metric spaces, following [25].

Let Y be a metric space. The *Čech complex* of a subspace $X \subseteq Y$ for radius $r > 0$ is the nerve of the collection of closed balls in Y with radius r centered at points in X:

$$\check{\mathrm{C}}\mathrm{ech}_r(X,Y) = \{\emptyset \neq S \subseteq X \mid S \text{ finite}, \bigcap_{x \in S} D_r(x) \neq \emptyset\},$$

where $D_r(x) = \{y \in Y \mid d(x,y) \leq r\}$ denotes the closed ball in Y of radius r centered at x.

A metric space is *hyperconvex* [12] if it is geodesic and if any collection of closed balls has the Helly property, i.e., if any two of these balls have a nonempty intersection, then all balls have a nonempty intersection. The following lemma is a direct consequence of this definition.

▶ **Lemma 10.** *If Y is hyperconvex and $X \subseteq Y$ is a subspace, then* $\check{\mathrm{C}}\mathrm{ech}_r(X,Y) = \mathrm{Rips}_{2r}(X)$.

Let X be a metric space. We describe its *injective hull* $E(X)$, following Lang [22]. A function $f\colon X \to \mathbb{R}$ with $f(x) + f(y) \geq d(x,y)$ for all $x, y \in X$ is *extremal* if $f(x) = \sup_{y \in X}(d(x,y) - f(y))$ for every $x \in X$. The difference between any two extremal functions turns out to be bounded, and so we can equip the set $E(X)$ of extremal functions with the metric induced by the supremum norm, i.e., $d(f,g) = \sup_{x \in X} |f(x) - g(x)|$. We define an isometric embedding $e\colon X \to E(X)$ by $y \mapsto d_y$, where $d_y(x) = d(y,x)$.

▶ Remark 11. $E(X)$ is a hyperconvex space. In particular, $E(X)$ is contractible, and nonempty intersections of closed metric balls are contractible [22, 19]. Moreover, nonempty intersections of open metric balls are also contractible [25, Proposition 2.8 and Lemma 2.15].

The following theorem is essentially due to Lang [22]. Originally, it has been stated for a special case, but the proof applies verbatim to the below statement involving our notion of the geodesic defect, which indeed provided the motivation for our definition. Note that the definition of δ-hyperbolic used in [22] differs from the one used here by a factor of 2.

▶ **Proposition 12** (Lang [22, Proposition 1.3])**.** *Let X be a δ-hyperbolic ν-geodesic metric space. Then the injective hull $E(X)$ is δ-hyperbolic, and every point in $E(X)$ has distance at most $2\delta + \nu$ to $e(X)$.*

Now we prove a generalization of the Contractibility Lemma using the injective hull analogously to the proof for geodesic spaces in [25, Corollary 8.4].

▶ **Theorem 13.** *Let X be a finite δ-hyperbolic ν-geodesic metric space. Then the complex $\mathrm{Rips}_t(X)$ is contractible for every $t \geq 4\delta + 2\nu$.*

Proof. By Proposition 12, we know that for $r > \frac{t}{2} \geq 2\delta + \nu$ the collection of open balls with radius r centered at the points in $e(X)$ covers $E(X)$. By finiteness of X, there exists an $r > \frac{t}{2}$ such that the nerve of this cover is isomorphic to $\check{\mathrm{C}}\mathrm{ech}_{\frac{t}{2}}(e(X), E(X))$. As e is an isometric embedding, Lemma 10, Remark 11, and the Nerve Theorem [17, Section 4.G] imply

$$\mathrm{Rips}_t(X) = \mathrm{Rips}_t(e(X)) = \check{\mathrm{C}}\mathrm{ech}_{\frac{t}{2}}(e(X), E(X)) \simeq E(X) \simeq *. \quad ◀$$

3 Filtered collapsibility of Vietoris–Rips complexes

In this section, we revisit the original proof of the Contractibility Lemma in [16], adapted to the language of discrete Morse theory [14]. Focusing on the finite case, which also constitutes the key part of the original proof, we extend the statement beyond geodesic spaces using our notion of geodesic defect, strengthen the assertion of contractibility to collapsibility, and further extend the result to become compatible with the Vietoris–Rips filtration.

▶ **Theorem 14.** *Let X be a finite δ-hyperbolic ν-geodesic metric space. Then for every $t \geq 4\delta + 2\nu$ there exists a discrete gradient that induces a collapse* $\mathrm{Rips}_t(X) \searrow \{*\}$.

Proof. Without loss of generality, assume that $\delta > 0$; if X is 0-hyperbolic, then it is also ϵ-hyperbolic for any $\epsilon > 0$, and for sufficiently small $\epsilon > 0$ we have $\mathrm{Rips}_{4\epsilon+2\nu}(X) = \mathrm{Rips}_{2\nu}(X)$.

Choose a reference point $p \in X$ and order the points according to their distance to p, choosing a total order $p = x_1 < \cdots < x_n$ on X such that $x_i < x_j$ implies $d(x_i, p) \leq d(x_j, p)$. Let $t \geq 4\delta + 2\nu$ and consider the filtration

$$\{p\} = K_1 \subseteq \cdots \subseteq K_n = \mathrm{Rips}_t(X),$$

where $K_i = \mathrm{Rips}_t(X_i)$ for $X_i := \{x_1, \ldots, x_i\}$. We prove that for $i \in \{2, \ldots, n\}$ there exists a discrete gradient V_i on K_i inducing a collapse $K_i \searrow K_{i-1}$.

First assume $d(x_i, p) < t$. Then for any vertex x_k of K_i we have $k \leq i$ and $d(x_k, p) \leq d(x_i, p) < t$, so K_i is a simplicial cone with apex p. Pairing the simplices containing p with those not containing p, we obtain a discrete gradient inducing a collapse $K_i \searrow K_{i-1}$:

$$V_i = \{(\sigma \setminus \{p\}, \sigma \cup \{p\}) \mid \sigma \in K_i \setminus K_{i-1}\}.$$

Now assume $d(x_i, p) \geq t$. We show that there exists a point $z \in X_{i-1}$ such that for every simplex $\sigma \in K_i \setminus K_{i-1}$, the union $\sigma \cup \{z\}$ is also a simplex in $K_i \setminus K_{i-1}$. To this end, we show that any vertex y of σ has distance $d(y, z) \leq t$ to z. For $r = d(x_i, p) - 2\delta - \nu$ and $s = 2\delta + \nu$ we have $r + s = d(x_i, p)$, and therefore, by the assumption that X is a ν-geodesic space, there exists a point $z \in X$ with $d(z, p) \leq r + \nu = d(x_i, p) - 2\delta$, implying $z < x_i$, and $d(z, x_i) \leq s + \nu = 2\delta + 2\nu$. By assumption $t \geq 4\delta + 2\nu$, and thus we get $d(z, x_i) \leq t - 2\delta$. Note that $y \in X_i$ implies $d(y, p) \leq d(x_i, p)$, and $y, x_i \in \sigma$ implies $d(y, x_i) \leq \operatorname{diam} \sigma \leq t$. The four-point condition (1) now yields

$$\begin{aligned} d(y,z) &\leq \max\{d(y,x_i) + d(z,p),\ d(y,p) + d(z,x_i)\} + 2\delta - d(x_i,p) \\ &= \max\{\underbrace{d(y,x_i)}_{\leq t} + \underbrace{d(z,p) - d(x_i,p)}_{\leq -2\delta},\ \underbrace{d(y,p) - d(x_i,p)}_{\leq 0} + \underbrace{d(z,x_i)}_{\leq t-2\delta}\} + 2\delta \leq t. \end{aligned} \tag{2}$$

Similarly to the above, pairing the simplices containing z with those not containing z yields a discrete gradient inducing a collapse $K_i \searrow K_{i-1}$:

$$V_i = \{(\sigma \setminus \{z\}, \sigma \cup \{z\}) \mid \sigma \in K_i \setminus K_{i-1}\}.$$

Finally, by Proposition 7, the union $V = \bigcup_i V_i$ is a discrete gradient on $\mathrm{Rips}_t(X)$ and by Proposition 8 it induces a collapse $\mathrm{Rips}_t(X) \searrow \{p\}$. ◀

▶ Remark 15. For a simplicial complex K, a particular type of simplicial collapse called an *elementary strong collapse* from K to $K \setminus \mathrm{St}\, v$ is defined in [5] for the case where the link of the vertex v is a simplicial cone. The proof of Theorem 14 actually shows that for $t \geq 4\delta + 2\nu$ there exists a sequence of elementary strong collapses from $\mathrm{Rips}_t(X)$ to $\{*\}$.

We can now extend the proof strategy of Theorem 14 to obtain a filtration-compatible strengthening of the Contractibility Lemma.

Proof of Theorem 3. As in the proof of Theorem 14, we can assume that $\delta > 0$, and order the points in X according to their distance to a chosen reference point $p = x_1 < \cdots < x_n$.

As X is finite, we can enumerate the values of pairwise distances by $0 = r_0 < \cdots < r_l$. For every $r_m > 4\delta + 2\nu$ we construct a discrete gradient W_m inducing a collapse $\mathrm{Rips}_{r_m}(X) \searrow \mathrm{Rips}_{r_{m-1}}(X)$. This will prove the theorem, because it follows from Theorem 14 that there exists a discrete gradient V that induces a collapse $\mathrm{Rips}_{4\delta+2\nu}(X) \searrow \{*\}$, and an application of Proposition 7 assembles these gradients into a single gradient $W = V \cup \bigcup_m W_m$on $\mathrm{Cl}(X)$ inducing collapses $\mathrm{Rips}_u(X) \searrow \mathrm{Rips}_t(X) \searrow \{*\}$ for every $u > t \geq 4\delta + 2\nu$.

Let m be arbitrary such that $r_m > 4\delta + 2\nu$. Consider the filtration

$$\mathrm{Rips}_{r_{m-1}}(X) = K_1 \subseteq \cdots \subseteq K_n = \mathrm{Rips}_{r_m}(X),$$

where $K_i = \mathrm{Rips}_{r_{m-1}}(X) \cup \mathrm{Rips}_{r_m}(X_i)$ for $X_i := \{x_1, \dots, x_i\}$. We prove that for $i \in \{2, \dots, n\}$ there exists a discrete gradient V_i on K_i inducing a collapse $K_i \searrow K_{i-1}$. Note that $K_i \setminus K_{i-1}$ consists of all simplices of diameter r_m that contain x_i as the maximal vertex.

First assume $d(x_i, p) < r_m$. Let $\sigma \in K_i \setminus K_{i-1}$. As x_i is the maximal vertex of σ, we have $d(v, p) \leq d(x_i, p) < r_m$ for all $v \in \sigma$. Since σ has diameter r_m, this implies that $\sigma \cup \{p\}$ also has diameter r_m. Moreover, this implies that there exists an edge $e \subseteq \sigma \setminus \{p\} \subseteq \sigma$ not containing p with $\operatorname{diam} e = r_m$. Therefore, $\sigma \setminus \{p\}$ also has diameter r_m. As $p < x_i$, both simplices $\sigma \setminus \{p\}$ and $\sigma \cup \{p\}$ contain x_i as the maximal vertex and are thus contained in $K_i \setminus K_{i-1}$. Pairing the simplices containing p with those not containing p, we obtain a discrete gradient inducing a collapse $K_i \searrow K_{i-1}$:

$$V_i = \{(\sigma \setminus \{p\}, \sigma \cup \{p\}) \mid \sigma \in K_i \setminus K_{i-1}\}.$$

Now assume $d(x_i, p) \geq r_m$. We show that there exists a point $z \in X_{i-1}$ such that for every simplex $\sigma \in K_i \setminus K_{i-1}$, the simplices $\sigma \setminus \{z\}$ and $\sigma \cup \{z\}$ are also contained in $K_i \setminus K_{i-1}$. To this end, we show first that any vertex y of σ has distance $d(y, z) \leq r_m$ to z. As in the proof of Theorem 14, there exists a point $z \in X$ with $d(z, p) \leq d(x_i, p) - 2\delta$, implying $z < x_i$, and $d(z, x_i) \leq 2\delta + 2\nu$. By assumption $r_m > 4\delta + 2\nu$, and thus we get $d(z, x_i) < r_m - 2\delta$. Similar to Equation (2), we have the following estimate

$$d(y, z) \leq \max\{\underbrace{d(y, x_i)}_{\leq r_m} + \underbrace{d(z, p) - d(x_i, p)}_{\leq -2\delta}, \underbrace{d(y, p) - d(x_i, p)}_{\leq 0} + \underbrace{d(z, x_i)}_{< r_m - 2\delta}\} + 2\delta \leq r_m,$$

and if $d(y, x_i) < r_m$, then $d(y, z) < r_m$. Hence, $\operatorname{diam}(\sigma \cup \{z\}) = r_m$, and $\operatorname{diam} \sigma = r_m$ implies $\operatorname{diam} \sigma \setminus \{z\} = r_m$, by an argument similar to the above. As $z < x_i$, both simplices $\sigma \setminus \{z\}$ and $\sigma \cup \{z\}$ contain x_i as the maximal vertex and are thus contained in $K_i \setminus K_{i-1}$. Pairing the simplices containing z with those not containing z, we obtain a discrete gradient inducing a collapse $K_i \searrow K_{i-1}$:

$$V_i = \{(\sigma \setminus \{z\}, \sigma \cup \{z\}) \mid \sigma \in K_i \setminus K_{i-1}\}.$$

By Proposition 7 the union $W_m = \bigcup V_i$ is a discrete gradient on $\mathrm{Rips}_{r_m}(X)$, and by Proposition 8 it induces a collapse $\mathrm{Rips}_{r_m}(X) \searrow \mathrm{Rips}_{r_{m-1}}(X)$. ◀

4 Collapsing Vietoris–Rips complexes of trees by apparent pairs

In this section, we analyze the Vietoris–Rips filtration of a tree metric space (V, d) for a positively weighted finite tree $T = (V, E)$, with the goal of proving the collapses in Theorem 5 using the apparent pairs gradient. To this end, we introduce two other discrete gradients: the *canonical gradient*, which is independent of any choices, and the *perturbed gradient*, which coarsens the canonical gradient and can be interpreted as a gradient that arises through a symbolic perturbation of the edge lengths. We then show that the intervals in the perturbed gradient are refined by apparent pairs of the lexicographically refined Vietoris–Rips filtration, with respect to a particular total order on the vertices.

We write $D_r(x) = \{y \in V \mid d(x, y) \leq r\}$ and $S_r(x) = \{y \in V \mid d(x, y) = r\}$.

▶ **Lemma 16.** *Let $x, y \in V$ be two distinct points at distance $d(x, y) = r$. Then we have* $\operatorname{diam} D_r(x) \cap D_r(y) = r$*. Furthermore, if $a, b \in D_r(x) \cap D_r(y)$ are points with $d(a, b) = r$, then these points are contained in the union $S_r(x) \cup S_r(y)$.*

Proof. We start by showing the first claim. Let $a, b \in D_r(x) \cap D_r(y)$ be any two points. We show that $d(a, b) \leq r$ holds, implying $\operatorname{diam} D_r(x) \cap D_r(y) \leq r$. Because $x, y \in D_r(x) \cap D_r(y)$ we also have $\operatorname{diam} D_r(x) \cap D_r(y) \geq r$, proving equality.

Write $[n] = \{1, \dots, n\}$ and let $\gamma\colon ([n], \{\{i, i+1\} \mid i \in [n-1]\}) \to T$ be the unique shortest path $x \rightsquigarrow y$. Moreover, let Ψ_a and Ψ_b be the unique shortest paths $x \rightsquigarrow a$ and $x \rightsquigarrow b$, respectively. Consider the largest numbers $t_a, t_b \in [n]$ with $\gamma(t_a) = \Psi_a(t_a)$ and $\gamma(t_b) = \Psi_b(t_b)$ and assume without loss of generality $t_a \leq t_b$. Note that the unique shortest path $a \rightsquigarrow b$ is then given by the concatenation $a \rightsquigarrow \gamma(t_a) \rightsquigarrow \gamma(t_b) \rightsquigarrow b$, where $\gamma(t_a) \rightsquigarrow \gamma(t_b)$ is the restricted path $\gamma_{|[t_a,t_b]}$. By assumption, we have $d(a, y) \leq r$ and this implies the inequality

$$d(a, \gamma(t_a)) + d(\gamma(t_a), y) = d(a, y) \leq r = d(x, y) = d(x, \gamma(t_a)) + d(\gamma(t_a), y),$$

which is equivalent to $d(a, \gamma(t_a)) \leq d(x, \gamma(t_a))$. Similarly, the assumption $d(x, b) \leq r$ implies $d(\gamma(t_b), b) \leq d(\gamma(t_b), y)$. Thus, the distance $d(a, b)$ satisfies

$$\begin{aligned} d(a, b) &= d(a, \gamma(t_a)) + d(\gamma(t_a), \gamma(t_b)) + d(\gamma(t_b), b) \\ &\leq d(x, \gamma(t_a)) + d(\gamma(t_a), \gamma(t_b)) + d(\gamma(t_b), y) = d(x, y) = r, \end{aligned} \tag{3}$$

which finishes the proof of the first claim.

We now show the second claim; assume $d(a, b) = r$. From the inequalities (3) and $d(a, \gamma(t_a)) \leq d(x, \gamma(t_a))$, $d(\gamma(t_b), b) \leq d(\gamma(t_b), y)$ together with the assumption $d(a, b) = r$, we deduce the equalities $d(a, \gamma(t_a)) = d(x, \gamma(t_a))$ and $d(\gamma(t_b), b) = d(\gamma(t_b), y)$. Hence,

$$d(a, y) = d(a, \gamma(t_a)) + d(\gamma(t_a), y) = d(x, \gamma(t_a)) + d(\gamma(t_a), y) = d(x, y) = r$$

and similarly $d(x, b) = r$, proving the second claim. ◀

Enumerate the values of pairwise distances by $0 = r_0 < \cdots < r_l = \operatorname{diam} V$. Let $K_m := \operatorname{Rips}_{r_{m-1}}(V) \cup T_{r_m}$. We show that the complement $C_m := \operatorname{Rips}_{r_m}(V) \setminus K_m$ is the set of all cofaces of non-tree edges of length r_m. We further show that it is partitioned into regular intervals in the face poset, and that this constitutes a discrete gradient.

▶ **Lemma 17.** *Every edge $e \in \operatorname{Rips}_{r_m}(V) \setminus \operatorname{Rips}_{r_{m-1}}(V)$ is contained in a unique maximal simplex $\Delta_e \in \operatorname{Rips}_{r_m}(V) \setminus \operatorname{Rips}_{r_{m-1}}(V)$. Moreover, if e is a tree edge of length r_m, then $\Delta_e = e$, and if $e \in C_m$, then $\Delta_e \in C_m$ and $e \subsetneq \Delta_e$.*

Proof. By definition, e corresponds to two points $x, y \in V$ at distance $d(x, y) = r_m$. If e is contained in the simplex $\Delta \in \mathrm{Rips}_{r_m}(V)$, then the points in Δ lie in the intersection $D_{r_m}(x) \cap D_{r_m}(y)$, which has diameter r_m by Lemma 16. Hence, the maximal simplex Δ_e is spanned by all the points in $D_{r_m}(x) \cap D_{r_m}(y)$.

If e is a tree edge of length r_m, then this intersection only contains x and y, and hence $\Delta_e = e$. If $e \in C_m$, then this intersection contains at least one vertex different from x and y that lies on the unique shortest path $x \rightsquigarrow y$. This implies $e \subsetneq \Delta_e$. ◀

For every maximal simplex $\Delta \in C_m \subseteq \mathrm{Rips}_{r_m}(V)$, we write E_Δ for the set of edges $e \in C_m$ with $\Delta_e = \Delta$. Note that E_Δ is the set of non-tree edges of length r_m contained in Δ.

4.1 Generic tree metrics

Before dealing with the general case, let us focus on the special case where the metric space (V, d) is *generic*, meaning that the pairwise distances are distinct. In this case, Lemma 17 implies that the diameter function $\mathrm{diam}\colon \mathrm{Cl}(V) \to \mathbb{R}$ is a discrete Morse function, defined on the full simplicial complex on V, with discrete gradient

$$\{[e, \Delta_e] \mid \text{non-tree edge } e \subseteq \mathrm{Cl}(V)\},$$

which we call the *generic gradient*, and only the vertices V and the tree edges E are critical. Together with Proposition 8, this yields the following theorem.

▶ **Theorem 18.** *If the tree metric space (V, d) is generic, then the generic gradient induces, for every $m \in \{1, \ldots, l\}$, a sequence of collapses*

$$\mathrm{Rips}_{r_m}(V) \searrow (\mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m}) \searrow T_{r_m}.$$

Moreover, it follows from Lemma 9 that for the Vietoris–Rips filtration, refined lexicographically with respect to an arbitrary total order on the vertices, the zero persistence apparent pairs refine the generic gradient, and therefore also induce the above collapses.

▶ **Theorem 19.** *If the tree metric space (V, d) is generic, then the apparent pairs gradient induces, for every $m \in \{1, \ldots, l\}$, a sequence of collapses*

$$\mathrm{Rips}_{r_m}(V) \searrow (\mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m}) \searrow T_{r_m}.$$

4.2 Arbitrary tree metrics

We now turn to the general case, where Lemma 9 is not directly applicable anymore, as the diameter function is not necessarily a discrete Morse function. Nevertheless, we show that Theorem 19 is still true without the genericity assumption, if the vertices V are ordered in a compatible way. Let Δ be a maximal simplex $\Delta \in C_m \subseteq \mathrm{Rips}_{r_m}(V)$.

▶ **Lemma 20.** *We have* $\mathrm{St}\, E_\Delta = C_m \cap \mathrm{Cl}\,\Delta$.

Proof. The inclusion $\mathrm{St}\, E_\Delta \subseteq C_m \cap \mathrm{Cl}\,\Delta$ holds by definition of E_Δ. To show the inclusion $\mathrm{St}\, E_\Delta \supseteq C_m \cap \mathrm{Cl}\,\Delta$, let $\sigma \in C_m \cap \mathrm{Cl}\,\Delta$ be any simplex. As the Vietoris–Rips complex is a clique complex, there exists an edge $e \subseteq \sigma \subseteq \Delta$ with $\mathrm{diam}\, e = r_m$. By Lemma 17, this edge can not be a tree edge end hence $e \in C_m$. Therefore, $e \in E_\Delta$ and $\sigma \in \mathrm{St}\, e \subseteq \mathrm{St}\, E_\Delta$. ◀

▶ **Lemma 21.** *If two distinct maximal simplices $\Delta, \Delta' \in C_m = \mathrm{Rips}_{r_m}(V) \setminus K_m$ intersect in a common face $\Delta \cap \Delta'$, then this face is contained in K_m.*

Proof. Assume for a contradiction that $\emptyset \neq \Delta \cap \Delta' \notin K_m$, implying $\Delta \cap \Delta' \in C_m$. By Lemma 20, there exists an edge $e \in E_\Delta \subseteq C_m$ with $e \subseteq \Delta \cap \Delta'$, and therefore $\Delta = \Delta'$ by uniqueness of the maximal simplex containing e (Lemma 17), a contradiction. ◀

We denote by L_Δ the set of all vertices of Δ that are not contained in any edge in E_Δ.

▶ **Lemma 22.** *Let $e = \{u, w\} \in E_\Delta$ be an edge. Then any point $x \in V \setminus \{u, w\}$ on the unique shortest path $u \rightsquigarrow w$ of length r_m in T is contained in L_Δ. In particular, L_Δ is nonempty.*

Proof. By assumption, we have $d(u, x) < r_m$, $d(w, x) < r_m$ and $d(u, w) = r_m$. Therefore, $\operatorname{diam}\{u, w, x\} = r_m$ and $x \in \{u, w, x\} \subseteq \Delta_e = \Delta$. Assume for a contradiction that x is contained in an edge in E_Δ. Then it follows from Lemma 16 that we have $d(u, x) = r_m$ or $d(w, x) = r_m$, contradicting the above. We conclude that $x \in L_\Delta$. ◀

4.2.1 The canonical gradient

We now describe a discrete gradient that is compatible with the diameter function and induces the same collapses as in Theorem 18 even if the tree metric is not generic. This construction is *canonical* in the sense that it does not depend on the choice of an order on the vertices, in contrast to the subsequent constructions.

▶ **Lemma 23.** *For any two edges $f, e \in E_\Delta$ and any vertex $v \in f$ there exists a vertex $z \in e$ such that $\{v, z\} \in E_\Delta$ is an edge in E_Δ.*

Proof. Let $f = \{v, w\}, e = \{x, y\}$; note that $d(v, w) = d(x, y) = r_m$. Since f and e are both contained in the maximal simplex Δ, we have $v, w \in D_{r_m}(x) \cap D_{r_m}(y)$. Both $\{v, x\}$ and $\{v, y\}$ are contained in $\{v, x, y\} \subseteq \Delta$ and Lemma 16 implies that at least one of these two edges is contained in $\operatorname{Rips}_{r_m}(V) \setminus \operatorname{Rips}_{r_{m-1}}(V)$; call this edge e_v. It follows from Lemma 17 that e_v is not a tree edge, and therefore $e_v \in E_\Delta$. ◀

▶ **Lemma 24.** *The set $\operatorname{St} E_\Delta = C_m \cap \operatorname{Cl} \Delta$ is partitioned by the intervals*

$$W_\Delta = \{[\cup S, (\cup S) \cup L_\Delta] \mid \emptyset \neq S \subseteq E_\Delta\}, \tag{4}$$

and these form a discrete gradient on $\operatorname{Cl} \Delta$ inducing a collapse $\operatorname{Cl} \Delta \searrow (K_m \cap \operatorname{Cl} \Delta)$.

Proof. The intervals in W_Δ are disjoint and contained in $\operatorname{St} E_\Delta$ by construction. They are regular, because L_Δ is nonempty (by Lemma 22). By Proposition 8, it remains to show that the intervals in W_Δ partition $\operatorname{St} E_\Delta = \operatorname{Cl} \Delta \setminus (K_m \cap \operatorname{Cl} \Delta)$ and that W_Δ is a discrete gradient.

To show the first claim, it suffices to prove that any simplex $\sigma \in \operatorname{St} E_\Delta$ is contained in a regular interval of W_Δ. Consider the simplex $\tau = \sigma \setminus L_\Delta \subseteq \sigma$. As $\sigma \in \operatorname{St} E_\Delta$, there exists an edge $e \in E_\Delta$ with $e \subseteq \sigma$. By the definition of L_Δ, we have $e \subseteq \sigma \setminus L_\Delta = \tau$. Any other vertex $v \in \tau \setminus e$ is also contained in one of the edges E_Δ. By Lemma 23, there exists an edge $e_v = \{v, w\} \in E_\Delta$, where $w \in e$. Then $\tau = e \cup \bigcup_{v \in \tau \setminus e} e_v$ and $\sigma \in [\tau, \tau \cup L_\Delta] \in W_\Delta$.

The second claim now follows from the observation that the function

$$\sigma \mapsto \begin{cases} \dim(\sigma \cup L_\Delta) & \sigma \in \operatorname{St} E_\Delta \\ \dim \sigma & \sigma \notin \operatorname{St} E_\Delta \end{cases}$$

is a discrete Morse function with discrete gradient W_Δ. ◀

Consider the union $W_m = \bigcup_\Delta W_\Delta$, where Δ runs over all maximal simplices in C_m and W_Δ is as in (4). We call $W = \bigcup_m W_m$ the *canonical gradient*.

▶ **Theorem 25.** *The canonical gradient is a discrete gradient on* $\mathrm{Cl}(V)$*. For every* $m \in \{1, \ldots, l\}$*, it induces a sequence of collapses*

$$\mathrm{Rips}_{r_m}(V) \searrow \mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m} \searrow T_{r_m}.$$

Proof. Let Δ be a maximal simplex in $\Delta \in C_m = \mathrm{Rips}_{r_m}(V) \setminus K_m$, where $K_m = \mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m}$. It follows from Lemma 24 that the set W_Δ is a discrete gradient on the full subcomplex $\mathrm{Cl}\,\Delta \subseteq \mathrm{Rips}_{r_m}(V)$ that partitions $\mathrm{St}\,E_\Delta = \mathrm{Cl}\,\Delta \setminus (K_m \cap \mathrm{Cl}\,\Delta)$ and that induces a collapse $\mathrm{Cl}\,\Delta \searrow (K_m \cap \mathrm{Cl}\,\Delta)$.

It follows directly from Lemma 21 and Proposition 7 that the union $W_m = \bigcup_\Delta W_\Delta$ is a discrete gradient on $\mathrm{Rips}_{r_m}(V)$. Again by Proposition 7, the union $W = \bigcup_m W_m$ is a discrete gradient on $\mathrm{Cl}(V)$.

By construction of the W_Δ, the union W_m partitions the complement $\mathrm{Rips}_{r_m}(V) \setminus K_m$. Hence, by Proposition 8, it induces a collapse $\mathrm{Rips}_{r_m}(V) \searrow K_m = \mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m}$. Since only the vertices and the tree edges are critical for W, this also yields the collapse to T_{r_m}. ◀

4.2.2 The perturbed gradient

Assume that V is totally ordered. We construct a coarsening of the canonical gradient to the *perturbed gradient*, such that under a specific total order of V the perturbed gradient is refined by the zero persistence apparent pairs of the diam-lexicographic order $<$ on simplices.

Consider a maximal simplex $\Delta \in C_m$, where $m \in \{1, \ldots, l\}$. Note that all edges in E_Δ have length r_m and thus are ordered lexicographically. Enumerate them as $e_1 < \cdots < e_q$. Every simplex $\sigma \in C_m \cap \mathrm{Cl}\,\Delta$ contains a maximal edge $e_\sigma \in \mathrm{Cl}\,\sigma \cap E_\Delta$.

▶ **Lemma 26.** *For every edge* $e_i \in E_\Delta$ *the union* $\Sigma_i = \bigcup_{e_\sigma = e_i} \sigma \subseteq \Delta$ *is a simplex in* C_m *and the maximal edge among* $\mathrm{Cl}\,\Sigma_i \cap E_\Delta$ *is* e_i*.*

Proof. Note that $\Sigma_i \subseteq \Delta \in \mathrm{Rips}_{r_m}(V)$ is a simplex and it is contained in C_m, because it is a coface of the non-tree edge e_i of length r_m.

To prove the second claim, let $e_j \in \mathrm{Cl}\,\Sigma_i \cap E_\Delta$ be any edge. Write $e_i = \{x, y\}$ with $x < y$ and $e_j = \{a, b\}$ with $a < b$. By construction of Σ_i, there exist simplices $\sigma_a, \sigma_b \in C_m \cap \mathrm{Cl}\,\Delta$ with $a \in \sigma_a, b \in \sigma_b$ and $e_{\sigma_a} = e_{\sigma_b} = e_i$. Note that $\{x, y, a\} \subseteq \sigma_a$ and $\{x, y, b\} \subseteq \sigma_b$.

By Lemma 16, we have $x, y \in S_{r_m}(a) \cup S_{r_m}(b)$ and therefore $d(a, y) = r_m$ (implying $a \neq y$) or $d(b, y) = r_m$ (implying $b \neq y$). As $\{a, y\} \subseteq \sigma_a$ and $\{b, y\} \subseteq \sigma_b$, this implies $\{a, y\} \leq e_{\sigma_a} = e_i = \{x, y\}$ or $\{b, y\} \leq e_{\sigma_b} = e_i = \{x, y\}$, respectively. In particular, we have $a \leq x$ or $a < b \leq x$, and if $a = x$, then $e_j \subseteq \sigma_b$. In any case, $e_j < e_i = e_{\sigma_b}$ as claimed. ◀

This lemma implies that $N_\Delta = \{[e_i, \Sigma_i]\}_{i=1}^{q}$ is a collection of disjoint intervals. It follows from Lemma 24 that for each $j \in \{1, \ldots, q\}$ the interval $[e_j, \Sigma_j]$ is the union

$$[e_j, \Sigma_j] = \bigcup \{[\cup S, (\cup S) \cup L_\Delta] \mid S \subseteq E_\Delta,\ e_j \text{ maximal element of } \mathrm{Cl}(\cup S) \cap E_\Delta\} \tag{5}$$

and that N_Δ partitions $C_m \cap \mathrm{Cl}\,\Delta$. Moreover, it is the discrete gradient of the function

$$f_\Delta \colon \mathrm{Cl}\,\Delta \to \mathbb{R},\ \sigma \mapsto \begin{cases} i & \sigma \in [e_i, \Sigma_i] \\ \dim \sigma - \dim \Delta & \sigma \in K_m \end{cases} \tag{6}$$

and the intervals are regular, because L_Δ is nonempty (Lemma 22). By Proposition 8, N_Δ induces a collapse $\mathrm{Cl}\,\Delta \searrow K_m \cap \mathrm{Cl}\,\Delta$. Therefore, the total order on V induces a symbolic perturbation scheme on the edges, establishing the situation of a generic tree metric as in Section 4.1.

Consider the union $N_m = \bigcup_\Delta N_\Delta$, where Δ runs over all maximal simplices in C_m. We call $N = \bigcup_m N_m$ the *perturbed gradient*. By (5), the perturbed gradient N coarsens the canonical gradient W. Analogously to Theorem 25, we obtain the following result.

▶ **Theorem 27.** *The perturbed gradient is a discrete gradient on* $\mathrm{Cl}(V)$*. For every* $m \in \{1, \ldots, l\}$*, it induces a sequence of collapses*

$$\mathrm{Rips}_{r_m}(V) \searrow \mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m} \searrow T_{r_m}.$$

▶ Remark 28. As the lower bounds of the intervals in the perturbed gradient are edges, it follows from Theorem 27 that these collapses can be expressed as *edge collapses* [9], a notion that is similar to the elementary strong collapses described in Remark 15.

4.2.3 The apparent pairs gradient

Finally, we show that for a specific total order of V, which we describe next, the perturbed gradient is refined by the *zero persistence apparent pairs* of the diam-lexicographic order.

From now on, assume that the tree T is rooted at an arbitrary vertex and orient every edge away from this point. Let $\leq_V$ be the partial order on V where u is smaller than w if there exists an oriented path $u \rightsquigarrow w$. In particular, we have the identity path $\mathrm{id}\colon u \rightsquigarrow u$. Note that for any two vertices $u, w \in V$ the unique shortest unoriented path $u \leftrightsquigarrow w$ can be written uniquely as a zig-zag $u \overset{\gamma}{\leftsquigarrow} z \overset{\eta}{\rightsquigarrow} w$, where z is the greatest point with $z \leq_V u$, $z \leq_V w$, and γ, η are oriented paths in T that intersect only in z. If $w \leftrightsquigarrow p$ is another unique shortest unoriented path with the zig-zag $w \overset{\varphi}{\leftsquigarrow} z' \overset{\mu}{\rightsquigarrow} p$, then we can form the following diagram

$$\begin{tikzcd}
& & z'' \arrow[dl, squiggly, "\xi"'] \arrow[dr, squiggly, "\lambda"] & & \\
& z \arrow[dl, squiggly, "\gamma"'] \arrow[dr, squiggly, "\eta"] & & z' \arrow[dl, squiggly, "\varphi"'] \arrow[dr, squiggly, "\mu"] & \\
u & & w & & p,
\end{tikzcd} \tag{7}$$

where z'' is the greatest point with $z'' \leq_V z, z'' \leq_V z'$. Moreover, as T has no cycles, it follows that either ξ or λ is the identity path and $\varphi \circ \lambda = \eta$ or $\eta \circ \xi = \varphi$, respectively.

Extend the partial order $\leq_V$ on V to a total order $<$ and consider the diam-lexicographic order on simplices. As this total order on the simplices extends $<$ under the identification $v \mapsto \{v\}$, we will also denote it by $<$. The following lemma directly implies Theorem 5.

▶ **Lemma 29.** *The intervals in the perturbed gradient* N *are refined by apparent pairs with respect to* $<$*. For every* $m \in \{1, \ldots, l\}$*, the zero persistence apparent pairs induce the collapse*

$$\mathrm{Rips}_{r_m}(V) \searrow \mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m}.$$

Proof. Consider a maximal simplex $\Delta \in C_m$. Recall that N_Δ is the discrete gradient of the function $f_\Delta\colon \mathrm{Cl}\,\Delta \to \mathbb{R}$ defined in (6), using the same vertex order as above. By Lemma 9, the zero persistence apparent pairs with respect to the f_Δ-lexicographic order $<_{f_\Delta}$ are precisely the gradient pairs of the minimal vertex refinement of N_Δ.

We next show that each apparent pair $(\sigma, \tau = \sigma \cup \{v\}) \subseteq [e_i, \Sigma_i]$ with respect to $<_{f_\Delta}$, where v is the minimal vertex in $\Sigma_i \setminus e_i$, is an apparent pair with respect to $<$. Clearly, these pairs have persistence zero with respect to the diameter function, as they appear in the same interval of the perturbed gradient. As the apparent pairs of $<_{f_\Delta}$, taken over all Δ, yield a partition of $C_m = \mathrm{Rips}_{r_m}(V) \setminus (\mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m})$, the same is then true for the apparent pairs of $<$. Thus, by Proposition 8, the apparent pairs gradient induces a collapse $\mathrm{Rips}_{r_m}(V) \searrow \mathrm{Rips}_{r_{m-1}}(V) \cup T_{r_m}$.

First, let $\sigma \cup \{p\} \in C_m$ be a cofacet of σ not equal to τ. We show that we must have $\tau < \sigma \cup \{p\}$, proving that τ is the minimal cofacet of σ with respect to $<$: If $p \in \Sigma_i$, then $p \in \Sigma_i \setminus e_i$, as $p \notin \sigma \supseteq e_i$, and the statement is true by minimality of v in the minimal vertex refinement. Now assume that $p \notin \Sigma_i$ and write $e_i = \{u, w\}$ with $u < w$. By (5), we have $L_\Delta \subseteq \Sigma_i$ and hence it follows that $p \notin L_\Delta$ and that the point p is contained in an edge in E_Δ, by definition of L_Δ. It follows from Lemma 16 that p together with at least one vertex of e_i forms an edge in E_Δ. Call this edge g; if there are two such edges, consider the larger one, and call it g. From $\{u, w, p\} \subseteq \Delta$ and $p \notin \Sigma_i$ we get $e_i < g$: The edge e_i is not the maximal edge of the two simplex $\{u, w, p\}$, since otherwise p would be contained in Σ_i. Hence, one of the two other edges is maximal, and that edge is g by definition. Considering the two possible cases $g = \{u, p\}$ and $g = \{w, p\}$, we must have $u < p$. We will argue that $v < p$ holds, which proves $\tau = \sigma \cup \{v\} < \sigma \cup \{p\}$.

Consider the diagram (7). If $\gamma \neq \mathrm{id}$, then it follows from the fact that $e_i = \{u, w\}$ is not a tree edge that along the unique shortest path $u \leftrightsquigarrow w$ there exists a vertex x distinct from u and w with $x < u < p$. Then $x \in L_\Delta \subseteq \Sigma_i \setminus e_i$ by Lemma 22, and as v is the minimal element in $\Sigma_i \setminus e_i$, we get $v \leq x < p$.

If $\gamma = \mathrm{id}$, then $u = z$, and it follows from $d(w, p) \leq r_m$ and $p \notin e_i = \{u, w\}$ that we must have $\lambda \neq \mathrm{id}$ and $\xi = \mathrm{id}$: Otherwise $\lambda = \mathrm{id}$ and $u = z$ lies on φ. Therefore, u lies on the unique shortest path from w to p and $d(w, p) = d(w, u) + d(u, p) = r_m + d(u, p) > r_m$, yielding a contradiction. Thus, the unique shortest path $(u = z) \leftrightsquigarrow p$ decomposes as $u \rightsquigarrow z' \rightsquigarrow p$, where $u \rightsquigarrow z'$ is contained in $u \rightsquigarrow z' \rightsquigarrow w$. Note that $u \neq z'$, because $\lambda \neq \mathrm{id}$. Hence, as e_i is not a tree edge, the immediate successor x of u on the path $u \rightsquigarrow w$ is distinct from u and w with $x \leq z'$. This point satisfies $x \leq z' \leq p$, and it follows from Lemma 22 that we have $x \in L_\Delta \subseteq \Sigma_i \setminus e_i$. Because $p \notin L_\Delta$ we even have $x < p$. Therefore, as v is the minimal vertex in $\Sigma_i \setminus e_i$, it follows that $v \leq x < p$.

It remains to prove that σ is the maximal facet of τ with respect to $<$. We write $e_i = \{u, w\}$ with $u < w$ and $\tau = \{b_0, \ldots, b_{\dim \tau}\}$ with $b_0 < \cdots < b_{\dim \tau}$. As $e_i \subseteq \tau$, there are indices $k_1 < k_2$ with $u = b_{k_1} < b_{k_2} = w$. If $k_1 > 0$, then $v = b_0$, so σ is of the form $\{b_1, \ldots, b_{\dim \tau}\}$ and is the maximal facet of τ with respect to $<$ as claimed. Now assume $k_1 = 0$. If τ contains no edges $e \in E_\Delta$ other than e_i, then the facets $\tau \setminus \{u\}$ and $\tau \setminus \{w\}$ are both contained in $\mathrm{Rips}_{r_{m-1}}(V)$, because they do not contain any edge of length r_m, and the maximal facet of τ is $\tau \setminus \{x\}$ with x the minimal vertex in $\tau \setminus e_i$. By assumption, we have $x = v$ and hence $\tau \setminus \{x\} = \tau \setminus \{v\} = \sigma$. If τ contains other edges $e \neq e_i$ with $e \in E_\Delta$, label them $s_1, \ldots, s_a$. As $e_i \subseteq \tau \subseteq \Sigma_i$, it follows from Lemma 26 that we have $s_b < e_i$ for all b. Because of this and our assumption $k_1 = 0$, i.e., u is the minimal vertex of τ, we have $s_b = \{u, x_b\} < \{u, w\} = e_i$ with $u < x_b < w$. Therefore, the facet $\{b_1, \ldots, b_{\dim \tau}\}$ contains no edges in E_Δ and hence it is contained in $\mathrm{Rips}_{r_{m-1}}(V)$. The facet $\{b_0, b_2, \ldots, b_{\dim \tau}\}$ of τ contains e_i, hence it is an element of C_m, and so it is maximal among the facets containing b_0, implying that it is the maximal facet of τ with respect to $<$. Because b_1 is the minimal vertex in $\tau \setminus e_i$ and $v \in \tau \setminus e_i$, it follows from the minimality of $v \in \Sigma_i \setminus e_i$ that we have $b_1 = v$, implying $\{b_0, b_2, \ldots, b_{\dim \tau}\} = \sigma$. Therefore, σ is the maximal facet of τ with respect to $<$. ◀

▶ Remark 30. The preceding Lemma 29 also implies Theorem 3 in the special case of tree metrics: if $u > t \geq 2\nu(V) = \max_{e \in E} l(e)$ are real numbers, then $T_t = T$ is the entire tree, and we obtain collapses $\mathrm{Rips}_u(V) \searrow \mathrm{Rips}_t(V) \searrow T \searrow \{*\}$. If all edges of T have the same length, it turns out that the collapse $T \searrow \{*\}$ is also induced by the apparent pairs gradient for the same order $<$.

References

1 Michał Adamaszek. Clique complexes and graph powers. *Israel J. Math.*, 196(1):295–319, 2013. doi:10.1007/s11856-012-0166-1.

2 Michał Adamaszek, Henry Adams, Ellen Gasparovic, Maria Gommel, Emilie Purvine, Radmila Sazdanovic, Bei Wang, Yusu Wang, and Lori Ziegelmeier. On homotopy types of Vietoris–Rips complexes of metric gluings. *J. Appl. Comput. Topol.*, 4(3):425–454, 2020. doi:10.1007/s41468-020-00054-y.

3 Dominique Attali, André Lieutier, and David Salinas. Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes. *Comput. Geom.*, 46(4):448–465, 2013. doi:10.1016/j.comgeo.2012.02.009.

4 Dominique Attali, André Lieutier, and David Salinas. When convexity helps collapsing complexes. In *35th International Symposium on Computational Geometry*, volume 129 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 11, 15. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. doi:10.4230/LIPIcs.SoCG.2019.11.

5 Jonathan Ariel Barmak and Elias Gabriel Minian. Strong homotopy types, nerves and collapses. *Discrete Comput. Geom.*, 47(2):301–328, 2012. doi:10.1007/s00454-011-9357-5.

6 Ulrich Bauer. Ripser: efficient computation of Vietoris-Rips persistence barcodes. *J. Appl. Comput. Topol.*, 5(3):391–423, 2021. doi:10.1007/s41468-021-00071-5.

7 Ulrich Bauer and Herbert Edelsbrunner. The Morse theory of Čech and Delaunay complexes. *Trans. Amer. Math. Soc.*, 369(5):3741–3762, 2017. doi:10.1090/tran/6991.

8 Michael Bleher, Lukas Hahn, Juan Angel Patino-Galindo, Mathieu Carriere, Ulrich Bauer, Raul Rabadan, and Andreas Ott. Topology identifies emerging adaptive mutations in SARS-CoV-2. Preprint, 2021. arXiv:2106.07292.

9 Jean-Daniel Boissonnat and Siddharth Pritam. Edge collapse and persistence of flag complexes. In *36th International Symposium on Computational Geometry*, volume 164 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 19, 15. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020. doi:10.4230/LIPIcs.SoCG.2020.19.

10 M. Bonk and O. Schramm. Embeddings of Gromov hyperbolic spaces. In *Selected works of Oded Schramm. Volume 1, 2*, Sel. Works Probab. Stat., pages 243–284. Springer, New York, 2011. With a correction by Bonk. doi:10.1007/978-1-4419-9675-6_10.

11 Joseph Minhow Chan, Gunnar Carlsson, and Raul Rabadan. Topology of viral evolution. *Proceedings of the National Academy of Sciences*, 110(46):18566–18571, 2013. doi:10.1073/pnas.1313480110.

12 R. Espínola and M. A. Khamsi. Introduction to hyperconvex spaces. In *Handbook of metric fixed point theory*, pages 391–435. Kluwer Acad. Publ., Dordrecht, 2001. doi:10.1007/978-94-017-1748-9_13.

13 Steven N. Evans. *Probability and real trees*, volume 1920 of *Lecture Notes in Mathematics*. Springer, Berlin, 2008. Lectures from the 35th Summer School on Probability Theory held in Saint-Flour, July 6–23, 2005. doi:10.1007/978-3-540-74798-7.

14 Robin Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998. doi:10.1006/aima.1997.1650.

15 Ragnar Freij. Equivariant discrete Morse theory. *Discrete Math.*, 309(12):3821–3829, 2009. doi:10.1016/j.disc.2008.10.029.

16 M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987. doi:10.1007/978-1-4613-9586-7_3.

17 Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

18 Patricia Hersh. On optimizing discrete Morse functions. *Adv. in Appl. Math.*, 35(3):294–322, 2005. doi:10.1016/j.aam.2005.04.001.

19 J. R. Isbell. Six theorems about injective metric spaces. *Comment. Math. Helv.*, 39:65–76, 1964. doi:10.1007/BF02566944.

20 Jakob Jonsson. *Simplicial complexes of graphs*, volume 1928 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2008. doi:10.1007/978-3-540-75859-4.

21 Dmitry N. Kozlov. *Organized collapse: an introduction to discrete Morse theory*, volume 207 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2020.

22 Urs Lang. Injective hulls of certain discrete metric spaces and groups. *J. Topol. Anal.*, 5(3):297–331, 2013. `doi:10.1142/S1793525313500118`.

23 Janko Latschev. Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold. *Arch. Math. (Basel)*, 77(6):522–528, 2001. `doi:10.1007/PL00000526`.

24 Michael Lesnick, Raúl Rabadán, and Daniel I. S. Rosenbloom. Quantifying genetic innovation: mathematical foundations for the topological study of reticulate evolution. *SIAM J. Appl. Algebra Geom.*, 4(1):141–184, 2020. `doi:10.1137/18M118150X`.

25 Sunhyuk Lim, Facundo Memoli, and Osman Berat Okutan. Vietoris-Rips Persistent Homology, Injective Metric Spaces, and The Filling Radius. Preprint, 2020. `arXiv:2001.07588`.

26 Vin de Silva and Gunnar Carlsson. Topological estimation using witness complexes. In Markus Gross, Hanspeter Pfister, Marc Alexa, and Szymon Rusinkiewicz, editors, *SPBG'04 Symposium on Point - Based Graphics 2004*. The Eurographics Association, 2004. `doi:10.2312/SPBG/SPBG04/157-166`.

27 Leopold Vietoris. Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen. *Math. Ann.*, 97(1):454–472, 1927. `doi:10.1007/BF01447877`.