

Proximity of Persistence Modules and their Diagrams

Frédéric Chazal*– David Cohen-Steiner†– Marc Glisse‡– Leonidas J. Guibas§– Steve Y. Oudot*

November 28, 2008

Abstract

Topological persistence has proven to be a key concept for the study of real-valued functions defined over topological spaces. Its validity relies on the fundamental property that the persistence diagrams of nearby functions are close. However, existing stability results are restricted to the case of continuous functions defined over triangulable spaces.

In this paper, we present new stability results that do not suffer from the above restrictions. Furthermore, by working at an algebraic level directly, we make it possible to compare the persistence diagrams of functions defined over different spaces, thus enabling a variety of new applications of the concept of persistence. Along the way, we extend the definition of persistence diagram to a larger setting, introduce the notions of discretization of a persistence module and associated pixelization map, define a proximity measure between persistence modules, and show how to interpolate between persistence modules, thereby lending a more analytic character to this otherwise algebraic setting. We believe these new theoretical concepts and tools shed new light on the theory of persistence, in addition to simplifying proofs and enabling new applications.

Keywords: Topological persistence, Stability, Persistence diagram, Discretization, Topological and Geometric Data Analysis.

*INRIA Saclay, Geometrica group, 4 rue Jacques Monod, 91893 ORSAY, France. Email: {frederic.chazal, steve.oudot}@inria.fr

†INRIA Sophia, Geometrica group, 2004 route des lucioles, 06902 Sophia Antipolis, France. Email: david.cohen-steiner@sophia.inria.fr

‡Gipsa-Lab, CNRS UMR 5216, Grenoble, France. Email: marc.glisse@gipsa-lab.inpg.fr

§Stanford University, Department of Computer Science, Stanford, CA 94305, USA. Email: guibas@cs.stanford.edu

1 Introduction

Topological persistence has emerged as a powerful tool for the study of the qualitative and quantitative behavior of real-valued functions. Given a topological space \mathbb{X} equipped with a function $f : \mathbb{X} \rightarrow \mathbb{R}$, persistence encodes the evolution of the topology of the sublevel-sets of f , *i.e.* the sets $\mathbb{X}_\alpha^f = f^{-1}((-\infty, \alpha]) \subseteq \mathbb{X}$, as parameter α ranges from $-\infty$ to $+\infty$. Topological changes occur only at critical values of f , which can be paired in a natural way. The outcome is a set of intervals, called a *persistence barcode* [6], where each interval encodes the birth and death times of a topological feature in the sublevel-sets of f . An equivalent representation is by a multiset of points in the extended plane \mathbb{R}^2 , called a *persistence diagram* [12], where the coordinates of each point correspond to the endpoints of some interval in the barcode.

Such representations prove to be useful in a variety of contexts. For instance, in scalar field analysis, they can be used to guide the simplification of a real-valued function by iterative cancellation of critical pairs, ridding the data of its inherent topological noise [1, 17, 18]. In topological data analysis, they can be used to infer the structure of an unknown space \mathbb{X} from a finite point sampling L , through the construction of an intermediate object, called a *filtration*, which consists of an abstract simplicial complex C built on top of the point cloud L together with a filtering function $\hat{f} : C \rightarrow \mathbb{R}$ that encodes the times of appearance of the simplices in the complex — see [7] for a survey. In these contexts as in many others, the validity of the persistence-based approach relies on the fundamental property that persistence diagrams are stable with respect to small perturbations of the functions. In scalar field analysis for instance, the scalar field f under study is usually known through some finite set of measurements, from which a piecewise-linear (PL) approximation \hat{f} of f is built. The simplification is then performed on \hat{f} , and the whole approach makes sense only if the persistence diagram of f can be related to the one of its approximation \hat{f} . In topological data analysis, the need for stability stems from the fact that the space \mathbb{X} underlying the input data set L remains unknown, which implies that the filtering function \hat{f} must be derived solely from the input data set L and shown to be close to some function $f : \mathbb{X} \rightarrow \mathbb{R}$ that filters the underlying space \mathbb{X} .

The stability of persistence diagrams was first studied by Cohen-Steiner, Edelsbrunner and Harer in their seminal paper [12]. In particular, they showed that the persistence diagrams of two real-valued functions f, g defined over a same topological space \mathbb{X} lie at most $\|f - g\|_\infty$ away from each other in the bottleneck distance. However, their result requires that three additional conditions be met: (1.) \mathbb{X} is triangulable, (2.) f and g are continuous, and (3.) f and g are *tame* in the sense that they only have finitely many critical values. Despite these restrictions, the stability result of [12] has found a variety of applications [2, 10, 11, 14, 18]. Interestingly enough, the result has also been applied within contexts where the above conditions are not met: in topological data analysis for instance, the real-valued function \hat{f} used to filter the simplicial complex C is usually taken to be constant over each simplex, and therefore non-continuous. However, as explained *e.g.* in [20], it can be replaced by some PL function with the same persistence diagram, defined over the first barycentric subdivision of C . Thus, a reduction from the piecewise-constant setting to some continuous setting is made. Nevertheless, such reductions may not always exist, and generally speaking the stability result of [12] suffers from the following limitations:

- The triangulability condition (1.), although reasonable in view of practical applications, may not always be satisfied in theory.
- The continuity condition (2.) is a stringent one. In the context of scalar field analysis for instance, if the original function f is not continuous, then its persistence diagram cannot be related to the one of its PL approximation \hat{f} , even though $\|f - \hat{f}\|_\infty$ is small. As mentioned above, although in some specific scenarios the problem can be easily reduced to some continuous setting, it is not clear that such reductions exist in general.
- The tameness condition (3.) requires that persistence diagrams only have finitely many points off the diagonal $\Delta = \{(x, x), x \in \mathbb{R}\}$. This is unfortunate as the zero-dimensional version of persistence, known as *size theory* and studied since the early 90's, does have a stability result that holds for a class

of functions with an infinite number of critical values, albeit defined only over compact connected manifolds [16, Thm. 25].

- Finally, the fact that the functions f, g have to be defined over a same topological space \mathbb{X} is a strong limitation. There indeed exist scenarios requiring to compare the persistence diagrams of functions defined over different spaces that are not related to each other in any obvious way. One such scenario served as the initial motivation for our work: it has to do with the analysis of scalar fields over sampled spaces where no PL approximation \hat{f} is readily available [9].

This paper presents new stability results that do not suffer from the above limitations: both continuity and triangulability conditions are removed, and the tameness condition is relaxed; moreover, functions can be defined over different topological spaces. To achieve this result, we drop the functional setting and work at algebraic level directly. Our analysis differs from the one of [12] in essential ways, has a more geometric flavor, and introduces several novel algebraic and geometric constructions that shed new light on the theory of persistence. On the practical side, our results have led to new algorithms for the analysis of scalar fields over point cloud data [9], thus enabling a variety of new applications of the persistence paradigm.

Details of our contributions. In the original persistence paper [17], the persistence diagram of a function $f : \mathbb{X} \rightarrow \mathbb{R}$ was derived from the family of homology groups of its sublevel sets $\{H_k(\mathbb{X}_\alpha^f)\}_{\alpha \in \mathbb{R}}$, enriched with the family of homomorphisms induced by the canonical inclusion maps $\mathbb{X}_\alpha^f \hookrightarrow \mathbb{X}_\beta^f$ for $\alpha \leq \beta$. In [21], the authors showed that persistence can in fact be defined at algebraic level directly, without the need for an underlying functional setting. Introducing the concept of *persistence module* \mathcal{F}_A as the one of a family $\{F_\alpha\}_{\alpha \in A}$ of vector spaces (or modules over a same commutative ring) indexed by $A \subseteq \mathbb{R}$, together with a family of homomorphisms $\{f_\alpha^\beta : F_\alpha \rightarrow F_\beta\}_{\alpha \leq \beta \in A}$ such that $\forall \alpha \leq \beta \leq \gamma, f_\alpha^\gamma = f_\beta^\gamma \circ f_\alpha^\beta$ and $f_\alpha^\alpha = \text{id}_{F_\alpha}$, they proved that persistence diagrams can be defined for persistence modules satisfying some tameness condition similar to (3.). Keeping persistence modules as our main objects of study, we propose a weaker tameness condition that allows them to have infinitely many critical values (Section 2).

Although this new tameness condition is similar in spirit to the one used in the 0-dimensional setting of size theory [15], it makes the standard definition of persistence diagram inapplicable. We therefore propose a new definition, based on an approximation strategy (Section 3): first, we discretize our persistence module \mathcal{F}_A over arbitrary discrete families of indices with no accumulation point, and show that the persistence diagrams of such discretizations are defined in a similar way as in the classical setting (Section 3.1); second, we obtain the persistence diagram of \mathcal{F}_A as a well-defined limit of the persistence diagrams of its various discretizations (Section 3.2). This new definition coincides with the standard one whenever the tameness condition of [12] is satisfied.

In order to make stability claims, we define a notion of proximity between persistence modules that is inspired from the functional setting (Section 4.1). More precisely, whenever $\|f - g\|_\infty \leq \varepsilon$, the sublevel-sets of functions f, g are ε -interleaved with respect to inclusion, that is: $\forall \alpha \in \mathbb{R}, \mathbb{X}_\alpha^f \subseteq \mathbb{X}_{\alpha+\varepsilon}^g \subseteq \mathbb{X}_{\alpha+2\varepsilon}^f$. Together with the canonical inclusions between sublevel-sets of f (resp. sublevel-sets of g), the above inclusions induce a commutative diagram at homology level that ε -interleaves the persistence modules of f and g . This notion of ε -interleaving of two persistence modules turns out to be independent of the functional setting, and defines a notion of distance between persistence modules. In addition, we show how to *interpolate* between any two ε -interleaved persistence modules $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$, i.e. how to build a family $\{\tilde{\mathcal{H}}_\mathbb{R}^s\}_{s \in [0, \varepsilon]}$ of persistence modules, with $\tilde{\mathcal{H}}_\mathbb{R}^0 \simeq \mathcal{F}_\mathbb{R}$ and $\tilde{\mathcal{H}}_\mathbb{R}^\varepsilon \simeq \mathcal{G}_\mathbb{R}$, such that $\forall s, s' \in [0, \varepsilon], \tilde{\mathcal{H}}_\mathbb{R}^s$ and $\tilde{\mathcal{H}}_\mathbb{R}^{s'}$ are $|s - s'|$ -interleaved.

Our main results are stated in terms of the above distance: first, we provide a simple and geometrically-flavored proof that any tame ε -interleaved persistence modules have 3ε -close persistence diagrams in the bottleneck distance (Section 4.2); then, combining this result with the interpolation technique described above, we reduce the bound on the bottleneck distance between persistence modules from 3ε down to ε , which is the best possible bound (Section 4.3).

2 Background and definitions

Extended plane, multisets and bottleneck distance. Throughout the paper, $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$ denotes the extended real line, and we use the following rules: $\forall x \in \mathbb{R}, x + \infty = +\infty$ and $x - \infty = -\infty$. The extended plane $\bar{\mathbb{R}}^2 = \bar{\mathbb{R}} \times \bar{\mathbb{R}}$ is endowed with the l^∞ norm, noted $\|\cdot\|_\infty$. Since $|x - y| = +\infty$ whenever $x \in \mathbb{R}$ and $y \in \{\pm\infty\}$, the topology induced by $\|\cdot\|_\infty$ on $\bar{\mathbb{R}}^2$ is such that the points of \mathbb{R}^2 , of $\{\pm\infty\} \times \mathbb{R}$, of $\mathbb{R} \times \{\pm\infty\}$, and of $\{\pm\infty\} \times \{\pm\infty\}$ form distinct connected components. Let $\Delta = \{(x, x), x \in \mathbb{R}\}$ be the diagonal, and $\Delta_+ = \{(x, y) \in \bar{\mathbb{R}}^2 : y \geq x\}$ the closed half-plane above Δ . More generally, for any $\delta \geq 0$, let $\Delta_+^\delta = \{(x, y) \in \bar{\mathbb{R}}^2 : y \geq x + 2\delta\}$ be the closed half-plane at l^∞ -distance δ above Δ .

A *multiset* D in $\bar{\mathbb{R}}^2$ is a subset of $\bar{\mathbb{R}}^2$ such that each point $p \in D$ is assigned a multiplicity $\text{mult}(p) \in \mathbb{N} \cup \{+\infty\}$. The *support* of D , noted $|D|$, is the subset considered without the multiplicities. Equivalently, D can be represented as a disjoint union $D = \bigcup_{p \in |D|} \prod_{i=1}^{\text{mult}(p)} p$. A *multi-bijection* m between two multisets D, D' is a bijection $m : \bigcup_{p \in |D|} \prod_{i=1}^{\text{mult}(p)} p \rightarrow \bigcup_{p' \in |D'|} \prod_{i=1}^{\text{mult}(p')} p'$. Given two multisets D and D' , we abuse notation and write $d_{\mathcal{H}}^\infty(D, D')$ for the Hausdorff distance (in the l^∞ metric) between their supports. A relevant distance between multisets is the so-called bottleneck distance $d_B^\infty(D, D')$, introduced in [12] and defined as $\inf_m \sup_{p \in D} \|p - m(p)\|_\infty$, where m ranges over all multi-bijections $D \rightarrow D'$.

Filtrations and persistence modules. The homology theory used in the paper is singular homology with coefficients in a commutative ring R with unity (see [19] for an introduction to the subject), which will be assumed to be a field and omitted in our notations.

Given a subset $A \subseteq \mathbb{R}$, a *filtration* is a family $\{\mathbb{X}_\alpha\}_{\alpha \in A}$ of topological spaces that are nested with respect to inclusion, that is: $\forall \alpha \leq \alpha' \in A, \mathbb{X}_\alpha \subseteq \mathbb{X}_{\alpha'}$. A special type of filtration is the one formed by the sublevel-sets $\mathbb{X}_\alpha^f = f^{-1}((-\infty, \alpha])$ of some real-valued function $f : \mathbb{X} \rightarrow \mathbb{R}$. Given an arbitrary filtration $\{\mathbb{X}_\alpha\}_{\alpha \in \mathbb{R}}$, the family of inclusion maps $\mathbb{X}_\alpha \hookrightarrow \mathbb{X}_{\alpha'}$ induces a family of homomorphisms between the k th homology groups $H_k(\mathbb{X}_\alpha)$, known as the *kth persistence module* of the filtration. In fact, persistence modules can be defined at algebraic level directly, regardless of any underlying topological or functional setting [21]:

Definition 2.1 *Let R be a commutative ring with unity, and A a subset of \mathbb{R} . A persistence module \mathcal{F}_A is a family $\{F_\alpha\}_{\alpha \in A}$ of R -modules indexed by the elements of A , together with a family $\{f_\alpha^{\alpha'} : F_\alpha \rightarrow F_{\alpha'}\}_{\alpha \leq \alpha' \in A}$ of homomorphisms such that: $\forall \alpha \leq \alpha' \leq \alpha'' \in A, f_\alpha^{\alpha''} = f_{\alpha'}^{\alpha''} \circ f_\alpha^{\alpha'}$ and $f_\alpha^\alpha = \text{id}_{F_\alpha}$.*

In our context, the ring R is assumed to be a fixed field, hence the modules F_α are vector spaces and the homomorphisms $f_\alpha^{\alpha'}$ are linear maps between vector spaces. In particular, the rank of $f_\alpha^{\alpha'}$ is a well-defined integer or $+\infty$. \mathcal{F}_A is said to be *discrete* whenever A is discrete with no accumulation point. This includes for instance all cases where the index set A is finite. Another important case is when A is a periodic set of the form $\alpha_0 + \varepsilon\mathbb{Z}$, where $\alpha_0 \in \mathbb{R}$ and $\varepsilon > 0$ are fixed parameters. In this case, \mathcal{F}_A is said to be *ε -periodic*.

Tameness. In [12], the k th persistence module of a function $f : \mathbb{X} \rightarrow \mathbb{R}$ is characterized as *tame* if (a.) all k th homology groups $H_k(\mathbb{X}_\alpha^f)$ are finite-dimensional, and (b.) there are only finitely many *homological critical values*, i.e. values $\alpha \in \mathbb{R}$ such that for all sufficiently small $\varepsilon > 0$ the maps $H_k(\mathbb{X}_{\alpha-\varepsilon}^f) \rightarrow H_k(\mathbb{X}_{\alpha+\varepsilon}^f)$ induced by inclusions are not isomorphisms. It turns out that condition (b.) is not necessary for our concepts and stability results to hold. Taking a purely algebraic point of view, we redefine tameness as follows:

Definition 2.2 *A persistence module \mathcal{F}_A is said to be tame if $\forall \alpha \in A, \dim F_\alpha < +\infty$.*

The fact that $\dim F_\alpha < +\infty$ implies that $\text{rank } f_\alpha^{\alpha'} < +\infty$ for all $\alpha' \geq \alpha$. From now on, and until the end of the paper, *tameness* will be understood as in Definition 2.2. In Sections 3 and 4 below, we show that this weaker tameness condition is sufficient for defining the persistence diagram of a persistence module, and we exhibit stability results for this class of persistence modules. Modulo some additional technicalities, we show in the full version of the paper [8] that persistence diagrams can be defined and their stability proven under an even weaker condition, called *δ -tameness*, which states that $\text{rank } f_\alpha^{\alpha'} < +\infty$ whenever $\alpha' - \alpha > \delta$.

3 Discretizing persistence modules

Definition 3.1 Let \mathcal{F}_A be a persistence module, and let B be a discrete subset of A with no accumulation points. The discretization of \mathcal{F}_A over B is the persistence module \mathcal{F}_B given by the family $\{\mathcal{F}_\alpha\}_{\alpha \in B}$ of vector spaces together with the family $\{f_\alpha^{\alpha'}\}_{\alpha \leq \alpha' \in B}$ of homomorphisms.

To every discrete set B with no accumulation points corresponds a *pixelization grid* $\Gamma_B \subset \bar{\mathbb{R}}^2$ whose vertices are the points of type (β, β') for β, β' ranging over $\bar{B} = B \cup \{\inf B, +\infty\}$. By convention, every grid cell is the Cartesian product of two right-closed intervals of $\bar{\mathbb{R}}$. Specifically, if $\inf B = -\infty$, then each grid cell is of one of the following forms, where $\beta_i < \beta_{i+1}$ (resp. $\beta_j < \beta_{j+1}$) are consecutive elements of B : $(\beta_i, \beta_{i+1}] \times (\beta_j, \beta_{j+1}]$, or $(\beta_i, \beta_{i+1}] \times \{+\infty\}$, or $\{-\infty\} \times (\beta_j, \beta_{j+1}]$, or $\{-\infty\} \times \{+\infty\}$. If on the contrary we have $\inf B > -\infty$, then each grid cell takes one of the following forms: $(\beta_i, \beta_{i+1}] \times (\beta_j, \beta_{j+1}]$, or $(\beta_i, \beta_{i+1}] \times \{+\infty\}$, or $[-\infty, \beta_i] \times (\beta_j, \beta_{j+1}]$, or $[-\infty, \beta_i] \times \{+\infty\}$. To the grid Γ_B is associated a *B-pixelization map* $\text{pix}_B : \Delta_+ \rightarrow \Gamma_B \cup \Delta$ that performs the following snapping operations: each point of Δ_+ lying in a cell \mathcal{C} of Γ_B that does not intersect the diagonal Δ is snapped onto the upper-right corner of \mathcal{C} , whereas each point lying in a grid cell that intersects Δ is snapped onto its nearest point of Δ — in particular, diagonal points are left unchanged. Figure 1 (left) illustrates the behavior of pix_B when $\inf B > -\infty$.

3.1 Persistence diagrams of discrete tame persistence modules

Let \mathcal{F}_B be a discrete tame persistence module. For clarity, we rewrite $B = \{\beta_i\}_{i \in I}$, where $I \subseteq \mathbb{Z}$ is such that $\beta_i < \beta_j$ for all $i < j \in I$. Such a rewriting is possible because B has no accumulation points. Then, $\bar{I} = I \cup \{\inf I, +\infty\}$ indexes \bar{B} . By convention, when $j = +\infty$, we let $\text{rank } f_{\beta_i}^{\beta_j} = \text{rank } f_{\beta_i}^{\beta_m}$ if I has a maximum element $m \in \mathbb{Z}$, and $\text{rank } f_{\beta_i}^{\beta_j} = \lim_{k \rightarrow +\infty} \text{rank } f_{\beta_i}^{\beta_k}$ otherwise. Such a limit always exists because the general inequality $\text{rank}(g \circ f) \leq \text{rank } f$ implies that, for any fixed i , the map $k \mapsto \text{rank } f_{\beta_i}^{\beta_k}$ is non-increasing and therefore constant for sufficiently large k , the ranks being non-negative integers. Similarly, if $\inf I = -\infty$, then for all $j \in I \cup \{+\infty\}$ we let $\text{rank } f_{\beta_{\inf I}}^{\beta_j} = \lim_{k \rightarrow -\infty} \text{rank } f_{\beta_k}^{\beta_j}$.

Definition 3.2 The persistence diagram of \mathcal{F}_B is the multi-subset $D\mathcal{F}_B$ of $\bar{\mathbb{R}}^2$ defined by:

- (i) $D\mathcal{F}_B$ is contained in $\Gamma_B \cap \Delta_+$,
- (ii) each point on the diagonal Δ has multiplicity $+\infty$,
- (iii) each node (β_i, β_j) with $i < j \in \bar{I}$ has multiplicity $\text{mult}(\beta_i, \beta_j) = \text{rank } f_{\beta_i}^{\beta_{j-1}} - \text{rank } f_{\beta_i}^{\beta_j}$ if $i = \inf I$, and $\text{mult}(\beta_i, \beta_j) = \text{rank } f_{\beta_i}^{\beta_{j-1}} - \text{rank } f_{\beta_i}^{\beta_j} + \text{rank } f_{\beta_{i-1}}^{\beta_j} - \text{rank } f_{\beta_{i-1}}^{\beta_{j-1}}$ if $i > \inf I$.

Condition (iii) is illustrated in Figure 1 (center). It follows from our tameness condition (Definition 2.2) and from standard rank arguments that the multiplicity of each point of $D\mathcal{F}_B \setminus \Delta$ is a finite non-negative integer. Moreover, an elementary computation shows that $D\mathcal{F}_B$ satisfies the following inclusion-exclusion property illustrated in Figure 1 (right):

Lemma 3.3 For all $i_1 < i_2 \leq j_1 < j_2 \in \bar{I}$, we have

$$\sum_{i_1 < i \leq i_2, j_1 < j \leq j_2} \text{mult}(\beta_i, \beta_j) = \text{rank } f_{\beta_{i_2}}^{\beta_{j_1}} - \text{rank } f_{\beta_{i_2}}^{\beta_{j_2}} + \text{rank } f_{\beta_{i_1}}^{\beta_{j_2}} - \text{rank } f_{\beta_{i_1}}^{\beta_{j_1}}.$$

Furthermore, for all $j_1 < j_2 \in \bar{I}$, we have $\sum_{j_1 < j \leq j_2} \text{mult}(\beta_{\inf I}, \beta_j) = \text{rank } f_{\beta_{\inf I}}^{\beta_{j_1}} - \text{rank } f_{\beta_{\inf I}}^{\beta_{j_2}}$.

It follows from this lemma that, for any given half-open upper-left quadrant $Q_\beta^{\beta'} = [-\infty, \beta] \times (\beta', +\infty]$ with $\beta \leq \beta' \in \mathbb{R}$, the total multiplicity (and therefore the support) of the points of $D\mathcal{F}_B$ contained in $Q_\beta^{\beta'}$ is finite.

This does not mean however that $|D\mathcal{F}_B| \setminus \Delta$ is finite. Nevertheless, since B has no accumulation points in \mathbb{R} , the vertices of the grid Γ_B do not accumulate in \mathbb{R}^2 nor in $\{\pm\infty\} \times \mathbb{R}$ nor in $\mathbb{R} \times \{\pm\infty\}$, and therefore $|D\mathcal{F}_B| \setminus \Delta$ has no accumulation points. Moreover, for any $\beta \leq \beta' \in \mathbb{R}$, the points of $D\mathcal{F}_B \cap ([\beta, \beta'] \times \mathbb{R})$ lying above Δ are covered by a finite union of half-open upper-left quadrants, which implies that their total multiplicity is finite. Thus, although $|D\mathcal{F}_B|$ may have infinitely many points off the diagonal Δ , it satisfies some local finiteness properties that will be exploited in the rest of the paper.




Figure 1: *Left: the pixelization map pix_B , where $B = \{\beta_1, \beta_2, \beta_3, \beta_4\}$. Center: the multiplicity of a node (β_i, β_j) is fully determined by the ranks of the homomorphisms corresponding to the corners of the bottom-left cell incident to (β_i, β_j) . The number of such corners is two or four, depending on whether $i = \inf I$ (top) or $i > \inf I$ (bottom). Right: the sum of the multiplicities of the nodes (red disks) contained in the box $(\beta_{i_1}, \beta_{i_2}] \times (\beta_{j_1}, \beta_{j_2}]$ is equal to the alternate sum of the ranks of the homomorphisms corresponding to the corners (black squares) of the box.*

3.2 Persistence diagrams of arbitrary tame persistence modules

In order to be able to define the persistence diagram of an arbitrary tame persistence module \mathcal{F}_A , we first need to compare the persistence diagrams of its various discretizations:

Theorem 3.4 *For any discretizations \mathcal{F}_B and \mathcal{F}_C of \mathcal{F}_A , the restriction of the pixelization map pix_B (resp. pix_C) to $D\mathcal{F}_{B \cup C}$ defines a multi-bijection between $D\mathcal{F}_{B \cup C}$ and $D\mathcal{F}_B$ (resp. $D\mathcal{F}_C$).*

An important special case of this result is when $B \subseteq C$. Then, we have $B \cup C = C$, and the theorem states that the restriction of pix_B to $D\mathcal{F}_C$ defines a multi-bijection between $D\mathcal{F}_C$ and $D\mathcal{F}_B$.

Another important special case is when B and C are ε -periodic families, of the form $B = \beta_0 + \varepsilon\mathbb{Z}$ and $C = \gamma_0 + \varepsilon\mathbb{Z}$ for fixed parameters $\beta_0, \gamma_0, \varepsilon$. In this case, the pixelization maps pix_B and pix_C move the points of $D\mathcal{F}_A$ by at most ε in the l^∞ norm. Since in addition they only increase the coordinates of the points, the composition $\text{pix}_C \circ \text{pix}_B^{-1}$ (here, pix_B^{-1} is to be understood as the inverse of the restriction of pix_B to $D\mathcal{F}_{B \cup C}$), which by Theorem 3.4 defines a multi-bijection between $D\mathcal{F}_B$ and $D\mathcal{F}_C$, also moves the points by at most ε . Therefore,

Corollary 3.5 *For any ε -periodic discretizations \mathcal{F}_B and \mathcal{F}_C of \mathcal{F}_A , we have $d_B^\infty(D\mathcal{F}_B, D\mathcal{F}_C) \leq \varepsilon$.*

More generally, in view of the definition of pixelization map given at the top of Section 3, we have $d_B^\infty(D\mathcal{F}_B, D\mathcal{F}_C) \leq \varepsilon$ whenever B and C form two right ε -covers of A , that is: $\sup_{\alpha \in A} \inf_{\beta \in B \cap [\alpha, +\infty)} |\alpha - \beta| \leq \varepsilon$ and $\sup_{\alpha \in A} \inf_{\gamma \in C \cap [\alpha, +\infty)} |\alpha - \gamma| \leq \varepsilon$.

Corollary 3.5 suggests that ε can be viewed as a scale parameter at which the persistence module \mathcal{F}_A is considered. In other words, the knowledge of \mathcal{F}_A at a scale of ε leads to the knowledge of its persistence diagram (not yet formally defined) with an uncertainty of ε .

Proof of Theorem 3.4. Consider the discretization of \mathcal{F} over the (discrete) union $B \cup C$. Considering \mathcal{F}_B and \mathcal{F}_C as two discretizations of $\mathcal{F}_{B \cup C}$, we will show that the persistence diagram of \mathcal{F}_B (resp. \mathcal{F}_C) is the image of the persistence diagram of $\mathcal{F}_{B \cup C}$ through the pixelization map pix_B (resp. pix_C).

Let \mathcal{C} be a cell of the grid Γ_B that does not intersect Δ , and let (β_i, β_j) be its upper-right corner. Assume without loss of generality that $i > \inf I$, the case $i = \inf I$ being similar. Denoting by $\text{mult}_B(\beta_i, \beta_j)$ the multiplicity of (β_i, β_j) in the diagram $D\mathcal{F}_B$, we have: $\text{mult}_B(\beta_i, \beta_j) = \text{rank } f_{\beta_i}^{\beta_j-1} - \text{rank } f_{\beta_i}^{\beta_j} + \text{rank } f_{\beta_{i-1}}^{\beta_j} - \text{rank } f_{\beta_{i-1}}^{\beta_j-1}$, which by Lemma 3.3 (applied to $\mathcal{F}_{B \cup C}$) is equal to $\sum_{q \in |D\mathcal{F}_{B \cup C} \cap \mathcal{C}|} \text{mult}_{B \cup C}(q)$. As a result, the restriction of pix_B to the grid cell \mathcal{C} snaps each point of $D\mathcal{F}_{B \cup C} \cap \mathcal{C}$ onto (β_i, β_j) while preserving the total multiplicity, thus defining a multi-bijection between $D\mathcal{F}_{B \cup C} \cap \mathcal{C}$ and $D\mathcal{F}_B \cap \mathcal{C}$.

Let now \mathcal{C} be a cell of Γ_B that intersects Δ . Then, the restriction of pix_B to \mathcal{C} projects the points of $D\mathcal{F}_{B \cup C} \cap \mathcal{C}$ orthogonally onto $\Delta \cap \mathcal{C} = D\mathcal{F}_B \cap \mathcal{C}$, which has infinite multiplicity. Therefore, it defines a multi-bijection between $D\mathcal{F}_{B \cup C} \cap \mathcal{C}$ and $D\mathcal{F}_B \cap \mathcal{C}$.

Applying the above arguments independently on every cell of the grid Γ_B , we obtain that the restriction of pix_B to $D\mathcal{F}_{B \cup C}$ defines a multi-bijection between $D\mathcal{F}_{B \cup C}$ and $D\mathcal{F}_B$. \square

We are now ready to define the persistence diagram of \mathcal{F}_A using a subdivision procedure. For the sake of simplicity, we assume from now on that $A = \mathbb{R}$. Arbitrary index sets $A \subseteq \mathbb{R}$ can be handled in a similar way, at the price of a significant increase in technicality.

We begin our procedure by considering an arbitrary discrete subset $B_0 \subset \mathbb{R}$ with no accumulation points that forms a right 1-cover of \mathbb{R} , that is: $\sup_{\alpha \in \mathbb{R}} \inf_{\beta \in B_0 \cap [\alpha, +\infty)} |\alpha - \beta| \leq 1$. One example of such a subset is $B_0 = \beta_0 + \mathbb{Z}$, for some fixed parameter β_0 . Then, inductively, for any integer $n > 0$ we let B_n be an arbitrary discrete superset of B_{n-1} with no accumulation points that forms a right 2^{-n} -cover of \mathbb{R} , that is: $\sup_{\alpha \in \mathbb{R}} \inf_{\beta \in B_n \cap [\alpha, +\infty)} |\alpha - \beta| \leq 2^{-n}$. In the above example, one can take $B_n = \beta_0 + 2^{-n}\mathbb{Z}$.

By construction, for all $n \in \mathbb{N}$ we have $B_n \subseteq B_{n+1}$, thus \mathcal{F}_{B_n} is a discretization of $\mathcal{F}_{B_{n+1}}$ and therefore the restriction of pix_{B_n} to $D\mathcal{F}_{B_{n+1}}$ defines a multi-bijection between $D\mathcal{F}_{B_n}$ and $D\mathcal{F}_{B_{n+1}}$, by Theorem 3.4. This multi-bijection moves the points by at most 2^{-n} since B_n is a right 2^{-n} -cover of \mathbb{R} . It follows that the sequence $\{D\mathcal{F}_{B_n}\}_{n \in \mathbb{N}}$ of multisets in Δ_+ converges to some limit multiset $M \subset \Delta_+$ in the bottleneck distance. By Corollary 3.5, this limit multiset is independent of the choice of the nested family $\{B_n\}_{n \in \mathbb{N}}$.

Definition 3.6 *The limit multiset M obtained by the above subdivision process is called the persistence diagram of the tame filtration $\mathcal{F}_{\mathbb{R}}$, denoted $D\mathcal{F}_{\mathbb{R}}$.*

An important property deriving from the above subdivision process is that pixelization maps relate the persistence diagram of $\mathcal{F}_{\mathbb{R}}$ to the ones of its discretizations:

Theorem 3.7 *Let $\mathcal{F}_{\mathbb{R}}$ be a tame persistence module. Then, for any discretization \mathcal{F}_B of $\mathcal{F}_{\mathbb{R}}$, the restriction of pix_B to $D\mathcal{F}_{\mathbb{R}}$ defines a multi-bijection between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{F}_B$. In the special case where \mathcal{F}_B is an ε -periodic family, it follows that $d_B^{\infty}(D\mathcal{F}_{\mathbb{R}}, D\mathcal{F}_B) \leq \varepsilon$.*

Proof. Let $B_0 = B \cup (\mathbb{Z} \cap \mathbb{R} \setminus B)$. Inductively, for all $n > 0$, let $B_n = B_{n-1} \cup (\mathbb{Z} \cap \mathbb{R} \setminus B_{n-1})$. The sets B_n are discrete with no accumulation points, and they form a nested family of subsets of \mathbb{R} such that $\sup_{\alpha \in \mathbb{R}} \inf_{\beta \in B_n \cap [\alpha, +\infty)} |\alpha - \beta| \leq 2^{-n}$ for all $n \in \mathbb{N}$. Therefore, according to Definition 3.6, the sequence $\{D\mathcal{F}_{B_n}\}_{n \in \mathbb{N}}$ converges to $D\mathcal{F}_{\mathbb{R}}$ in the bottleneck distance. Furthermore, since by construction we have $B \subseteq B_0 \subseteq B_1 \subseteq \dots \subseteq B_n$, we deduce that $\text{pix}_B = \text{pix}_{B_0} \circ \text{pix}_{B_1} \circ \dots \circ \text{pix}_{B_{n-1}}$. Therefore, by Theorem 3.4, the restriction of pix_B to $D\mathcal{F}_{B_n}$ defines a multi-bijection between $D\mathcal{F}_{B_n}$ and $D\mathcal{F}_B$. Since this is true for all $n \in \mathbb{N}$, the restriction of pix_B to the limit multiset $D\mathcal{F}_{\mathbb{R}}$ defines a multi-bijection between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{F}_B$. \square

In the case where $\mathcal{F}_{\mathbb{R}}$ is the k th persistence module of the sublevel-sets filtration of some function $f : \mathbb{X} \rightarrow \mathbb{R}$ that is tame in the sense of [12], its persistence diagram as defined in [12] coincides with its persistence diagram in the sense of Definition 3.6. See the full version for more details [8].

4 Stability of persistence diagrams

This section provides equivalents to the stability result of [12] in the general setting of tame persistence modules. We first introduce a quantitative notion of proximity between persistence modules in Section 4.1. We propose in fact two notions of proximity: a weaker one and a stronger one, which give rise respectively to a weaker and a stronger stability results, studied in Sections 4.2 and 4.3 respectively. Both results provide tight upper bounds on the stability of persistence diagrams under their respective notions of proximity. In addition, the weaker stability result (Theorem 4.3) has a simple and geometrically-flavored proof, and it is instrumental in proving the stronger stability result (Theorem 4.4).

4.1 Interleaving persistence modules

To emphasize the intuition underlying our definitions, we first consider the case of persistence modules associated with the sublevel sets filtrations of functions. Whenever two functions $f, g : \mathbb{X} \rightarrow \mathbb{R}$ satisfy $\|f - g\|_\infty \leq \varepsilon$, their sublevel sets filtrations are nested as follows: $\forall \alpha \in \mathbb{R}, \mathbb{X}_\alpha^f \subseteq \mathbb{X}_{\alpha+\varepsilon}^g \subseteq \mathbb{X}_{\alpha+2\varepsilon}^f$. This nesting, combined with the canonical inclusions $\mathbb{X}_\alpha^f \subseteq \mathbb{X}_{\alpha'}^f$ and $\mathbb{X}_\alpha^g \subseteq \mathbb{X}_{\alpha'}^g$ for all $\alpha \leq \alpha'$, induces the following commutative diagrams at k th homology level, where $F_\alpha = H_k(\mathbb{X}_\alpha^f)$ and $G_\alpha = H_k(\mathbb{X}_\alpha^g)$ denote the k th homology groups of the sublevel sets, and where the arrows represent the homomorphisms induced by inclusions at k th homology level:

$$\begin{array}{ccc}
 F_{\alpha-\varepsilon} & \xrightarrow{\quad} & F_{\alpha'+\varepsilon} \\
 & \searrow & \nearrow \\
 & G_\alpha & \xrightarrow{\quad} G_{\alpha'} \\
 & \nearrow & \searrow \\
 & F_\alpha & \xrightarrow{\quad} F_{\alpha'} \\
 G_{\alpha-\varepsilon} & \xrightarrow{\quad} & G_{\alpha'+\varepsilon}
 \end{array}
 \qquad
 \begin{array}{ccc}
 & F_{\alpha+\varepsilon} & \xrightarrow{\quad} F_{\alpha'+\varepsilon} \\
 & \nearrow & \nearrow \\
 G_\alpha & \xrightarrow{\quad} G_{\alpha'} & \\
 & \searrow & \searrow \\
 F_\alpha & \xrightarrow{\quad} F_{\alpha'} & \\
 & \nearrow & \nearrow \\
 & G_{\alpha+\varepsilon} & \xrightarrow{\quad} G_{\alpha'+\varepsilon}
 \end{array}
 \tag{1}$$

The two persistence modules $\{F_\alpha\}_{\alpha \in \mathbb{R}}$ and $\{G_\alpha\}_{\alpha \in \mathbb{R}}$ are then said to be *strongly ε -interleaved*. This condition can be relaxed by assuming that the sublevel sets of f and g are only interleaved over a certain ε -periodic set of values of α , that is: $\exists \alpha_0 \in \mathbb{R}$ such that $\forall \alpha \in \alpha_0 + 2\varepsilon\mathbb{Z}, \mathbb{X}_\alpha^f \subseteq \mathbb{X}_{\alpha+\varepsilon}^g \subseteq \mathbb{X}_{\alpha+2\varepsilon}^f$. In this case, the two persistence modules induced at k th homology level are said to be *weakly ε -interleaved*, and the following induced diagram between their 2ε -discretizations commutes:

$$\begin{array}{ccccccc}
 & & \cdots & \xrightarrow{\quad} & F_{\alpha_0+2n\varepsilon} & \xrightarrow{\quad} & F_{\alpha_0+(2n+2)\varepsilon} & \xrightarrow{\quad} & \cdots \\
 & & & \nearrow & & \searrow & & \nearrow & \\
 \cdots & \xrightarrow{\quad} & G_{\alpha_0+(2n-1)\varepsilon} & \xrightarrow{\quad} & G_{\alpha_0+(2n+1)\varepsilon} & \xrightarrow{\quad} & \cdots & &
 \end{array}
 \tag{2}$$

These properties extend directly to arbitrary persistence modules:

Definition 4.1 Two persistence modules \mathcal{F}_A and \mathcal{G}_B are said to be weakly ε -interleaved if:

- (i) there exists $\alpha_0 \in \mathbb{R}$ such that $\alpha_0 + 2\varepsilon\mathbb{Z} \subseteq A$ and $\alpha_0 + \varepsilon + 2\varepsilon\mathbb{Z} \subseteq B$, and
- (ii) there exist two families of homomorphisms $\{\phi_\alpha : F_\alpha \rightarrow G_{\alpha+\varepsilon}\}_{\alpha \in \alpha_0+2\varepsilon\mathbb{Z}}$ and $\{\psi_\alpha : G_\alpha \rightarrow F_{\alpha+\varepsilon}\}_{\alpha \in \alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ such that the diagram of Eq. (2) commutes.

For the strong notion of proximity, we require that the index sets satisfy $A = B = \mathbb{R}$:

Definition 4.2 Two persistence modules $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ are said to be strongly ε -interleaved if there exist two families of homomorphisms $\{\phi_\alpha : F_\alpha \rightarrow G_{\alpha+\varepsilon}\}_{\alpha \in \mathbb{R}}$ and $\{\psi_\alpha : G_\alpha \rightarrow F_{\alpha+\varepsilon}\}_{\alpha \in \mathbb{R}}$ such that the diagrams of Eq. (1) commute for all $\alpha \leq \alpha' \in \mathbb{R}$.

Clearly, if $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ are strongly ε -interleaved, then they are also weakly ε -interleaved. Conversely, if \mathcal{F}_A and \mathcal{G}_B are weakly ε -interleaved, with $A = B = \mathbb{R}$, then they are strongly 3ε -interleaved, and this bound is tight — see the full version [8]. Nevertheless, \mathcal{F}_A and \mathcal{G}_B cannot be strongly interleaved when $A, B \subsetneq \mathbb{R}$.

4.2 Persistence diagrams of weakly interleaved persistence modules

Theorem 4.3 (Weak Stability Theorem) *Let \mathcal{F}_A and \mathcal{G}_B be two tame persistence modules. If \mathcal{F}_A and \mathcal{G}_B are weakly ε -interleaved, then $d_B^\infty(D\mathcal{F}_A, D\mathcal{G}_B) \leq 3\varepsilon$, and this bound is tight.*

Proof. Let $\alpha_0 \in \mathbb{R}$ be as in Definition 4.1 (i). Consider the persistence module $\mathcal{H}_{\alpha_0+\varepsilon\mathbb{Z}}$ defined by:

$$\forall n \in \mathbb{Z}, \begin{cases} H_{\alpha_0+2n\varepsilon} = F_{\alpha_0+2n\varepsilon} \text{ and } H_{\alpha_0+(2n+1)\varepsilon} = G_{\alpha_0+(2n+1)\varepsilon} \\ h_{\alpha_0+2n\varepsilon}^{\alpha_0+(2n+1)\varepsilon} = \phi_{\alpha_0+2n\varepsilon} \text{ and } h_{\alpha_0+(2n+1)\varepsilon}^{\alpha_0+(2n+2)\varepsilon} = \psi_{\alpha_0+(2n+1)\varepsilon} \end{cases}$$

By commutativity of the diagram of Eq. (2), $\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}$ and $\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ are two discretizations of $\mathcal{H}_{\alpha_0+\varepsilon\mathbb{Z}}$ over 2ε -periodic sets. Since $\mathcal{H}_{\alpha_0+\varepsilon\mathbb{Z}}$ itself is discrete, Corollary 3.5 implies that $d_B^\infty(D\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}, D\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}) \leq 2\varepsilon$. In addition, $\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}$ and $\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ are discretizations of \mathcal{F}_A and \mathcal{G}_B respectively, therefore Theorem 3.7 implies that $d_B^\infty(D\mathcal{F}_A, D\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}) \leq 2\varepsilon$ and $d_B^\infty(D\mathcal{G}_B, D\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}) \leq 2\varepsilon$. It follows then, by the triangle inequality, that $d_B^\infty(D\mathcal{F}_A, D\mathcal{G}_B) \leq 6\varepsilon$.

In order to reduce the bound from 6ε to 3ε , we need to study how the points of the above diagrams are moved by the multi-bijections induced by the pixelization maps. Let m_1 (resp. m_2) denote the multi-bijection induced by $\text{pix}_{\alpha_0+2\varepsilon\mathbb{Z}}$ between $D\mathcal{F}_A$ and $D\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}$ (resp. between $D\mathcal{H}_{\alpha_0+\varepsilon\mathbb{Z}}$ and $D\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}$). Similarly, let m_3 (resp. m_4) denote the multi-bijection induced by $\text{pix}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ between $D\mathcal{H}_{\alpha_0+\varepsilon\mathbb{Z}}$ and $D\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ (resp. between $D\mathcal{G}_B$ and $D\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$). The map $m = m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1$ is then a multi-bijection between $D\mathcal{F}_A$ and $D\mathcal{G}_B$. Let us track the various possible images of a point $p \in D\mathcal{F}_A$ through this multi-bijection — please refer to Figure 2 (left):

- $m_1(p)$ is at a vertex (u, v) of the grid $\Gamma_{\alpha_0+2\varepsilon\mathbb{Z}}$, marked by a blue disc in the figure;
- $m_2^{-1} \circ m_1(p)$ lies among the four corners of the cell of the grid $\Gamma_{\alpha_0+\varepsilon\mathbb{Z}}$ that contains $m_1(p)$, namely: (u, v) , $(u, v - \varepsilon)$, $(u - \varepsilon, v)$, and $(u - \varepsilon, v - \varepsilon)$, marked by red crosses in the figure;
- the images of these four corners through m_3 are among the four points $(u - \varepsilon, v - \varepsilon)$, $(u - \varepsilon, v + \varepsilon)$, $(u + \varepsilon, v - \varepsilon)$ and $(u + \varepsilon, v + \varepsilon)$, marked by blue squares in the figure;
- since m_4 is the restriction of $\text{pix}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ to $D\mathcal{G}_B$, the possible pre-images of $m_3 \circ m_2^{-1} \circ m_1(p)$ are contained in the union of the bottom left cells of $(u - \varepsilon, v - \varepsilon)$, $(u - \varepsilon, v + \varepsilon)$, $(u + \varepsilon, v - \varepsilon)$ and $(u + \varepsilon, v + \varepsilon)$ in the grid $\Gamma_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ (the gray area in the figure).

All in all, $m(p)$ belongs to the box $(u - 3\varepsilon, u + \varepsilon] \times (v - 3\varepsilon, v + \varepsilon]$. Since $p \in (u - 2\varepsilon, u] \times (v - 2\varepsilon, v]$, we conclude that $\|p - m(p)\|_\infty < 3\varepsilon$. The tightness of this bound is proven in the full version [8]. \square




Figure 2: Left: for the proof of Theorem 4.3. Right: for the proof of Lemma 4.5.

4.3 Persistence diagrams of strongly interleaved persistence modules

Theorem 4.4 (Strong Stability Theorem) *Let $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ be two tame persistence modules. If $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ are strongly ε -interleaved, then $d_B^\infty(D\mathcal{F}_\mathbb{R}, D\mathcal{G}_\mathbb{R}) \leq \varepsilon$.*

The entire Section 4.3 is devoted to the proof of this result. At a high level, our analysis follows the same scheme as in [12]. First, we bound the Hausdorff distance between the persistence diagrams of strongly ε -interleaved persistence modules (Section 4.3.1). The key ingredient for this stage is the so-called Box Lemma from [12], stated as Lemma 4.5 below, to which we provide a new proof based solely on pixelization arguments. Second, we move from Hausdorff to bottleneck distance by means of an interpolation argument (Section 4.3.2). However, differently from [12], we do not interpolate at functional level, but rather at algebraic level directly (Lemma 4.6), since in our context persistence modules are the only available data. In addition to being more general, our strategy is interesting from a technical point of view since it produces tame families of interpolating persistence modules, which a naive function interpolation does not always do.

4.3.1 Bound on the Hausdorff distance

Lemma 4.5 (Box Lemma) *Let $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ be two tame, strongly ε -interleaved persistence modules. Given any $\alpha < \beta < \gamma < \delta$, let \square denote the box $(\alpha, \beta] \times (\gamma, \delta]$, and \square_{ε} the box $(\alpha - \varepsilon, \beta + \varepsilon] \times (\gamma - \varepsilon, \delta + \varepsilon]$ obtained by inflating \square by ε . Then, the sum of the multiplicities of the points of $D\mathcal{F}_{\mathbb{R}}$ contained in \square is at most the sum of the multiplicities of the points of $D\mathcal{G}_{\mathbb{R}}$ contained in \square_{ε} .*

Proof. If $\beta + \varepsilon > \gamma - \varepsilon$, then \square_{ε} intersects the diagonal Δ , hence the total multiplicity of $D\mathcal{G}_{\mathbb{R}} \cap \square_{\varepsilon}$ is infinite and thus at least the total multiplicity of $D\mathcal{F}_{\mathbb{R}} \cap \square$. Assume now that $\beta + \varepsilon \leq \gamma - \varepsilon$. Let $A = \{\alpha, \beta, \gamma, \delta\}$ and $B = \{\alpha - \varepsilon, \beta + \varepsilon, \gamma - \varepsilon, \delta + \varepsilon\}$. Consider the A -discretization of $\mathcal{F}_{\mathbb{R}}$ and the B -discretization of $\mathcal{G}_{\mathbb{R}}$. Since $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ are strongly interleaved, the following diagram commutes (where diagonal arrows stand for the homomorphisms $\phi_{\alpha}, \psi_{\alpha}$ introduced in Definition 4.2):

$$\begin{array}{ccccccc}
 & & F_{\alpha} & \longrightarrow & F_{\beta} & \longrightarrow & F_{\gamma} & \longrightarrow & F_{\delta} \\
 & \nearrow & & & & & & & \searrow \\
 G_{\alpha-\varepsilon} & \longrightarrow & & & G_{\beta+\varepsilon} & \longrightarrow & G_{\gamma-\varepsilon} & \longrightarrow & G_{\delta+\varepsilon}
 \end{array}$$

It follows that \mathcal{F}_A and \mathcal{G}_B are two discretizations of the mixed persistence module $\mathcal{H}_{A \cup B}$ defined by the path $G_{\alpha-\varepsilon} \rightarrow F_{\alpha} \rightarrow F_{\beta} \rightarrow G_{\beta+\varepsilon} \rightarrow G_{\gamma-\varepsilon} \rightarrow F_{\gamma} \rightarrow F_{\delta} \rightarrow G_{\delta+\varepsilon}$ in the above diagram. Then, the same tracking approach as in the proof of Theorem 4.3 shows that the pixelization maps send the points $p \in D\mathcal{F}_{\mathbb{R}} \cap \square$ injectively to a subset of the points of $D\mathcal{G}_{\mathbb{R}} \cap \square_{\varepsilon}$, thus concluding the proof of the lemma. The tracking is illustrated in Figure 2 (right) and detailed in the full version of the paper [8]. \square

By invoking the Box Lemma in the vicinity of every point of $D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}$ separately, we deduce that $d_{\mathcal{H}}^{\infty}(D\mathcal{F}_{\mathbb{R}}, D\mathcal{G}_{\mathbb{R}}) \leq \varepsilon$ whenever $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ are tame and strongly ε -interleaved.

4.3.2 Bound on the bottleneck distance

Lemma 4.6 *Let $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ be two strongly ε -interleaved persistence modules. Then, for all $s \in [0, \varepsilon]$, there exists a persistence module $\tilde{\mathcal{H}}_{\mathbb{R}}^s$ that is strongly s -interleaved with $\mathcal{F}_{\mathbb{R}}$ and strongly $(\varepsilon - s)$ -interleaved with $\mathcal{G}_{\mathbb{R}}$. Furthermore, this persistence module is tame whenever $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ are.*

Proof. We present here an intuitive version of the construction of $\tilde{\mathcal{H}}_{\mathbb{R}}^s$; a detailed proof is available in Appendix A. For clarity, we let $\varepsilon_1 = s$ and $\varepsilon_2 = \varepsilon - s$. Denote by $\phi_{\alpha}^{\mathcal{F}, \mathcal{G}} : F_{\alpha} \rightarrow G_{\alpha+\varepsilon}$ and $\phi_{\alpha}^{\mathcal{G}, \mathcal{F}} : G_{\alpha} \rightarrow F_{\alpha+\varepsilon}$ the homomorphisms provided by Definition 4.2. We want to define a persistence module $\mathcal{H}_{\mathbb{R}}$ that is close to both $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$. The first idea is to consider their sum. However, we do not define H_a as $F_a \oplus G_a$ because, although there is a natural application from $F_{a-\varepsilon_1}$ to $F_a \oplus G_a$ (using $f_{a-\varepsilon_1}^a$), there is no natural way to define an application from $F_a \oplus G_a$ to $F_{a+\varepsilon_1}$. Instead, we let $\mathcal{H}_{\mathbb{R}}$ be a translated sum of $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$: $H_a = F_{a-\varepsilon_1} \oplus G_{a-\varepsilon_2}$. There is an obvious injection from $F_{a-\varepsilon_1}$ into H_a , and now $f_{a-\varepsilon_1}^{a+\varepsilon_1} \oplus \phi_{a-\varepsilon_2}^{\mathcal{G}, \mathcal{F}}$ provides a natural homomorphism from H_a to $F_{a+\varepsilon_1}$. The situation with \mathcal{G} is symmetric.

$$\begin{array}{ccccc}
F_{a-\varepsilon_1-2\varepsilon_2} & \xrightarrow{\quad} & F_{a-\varepsilon_1} & \xrightarrow{\quad} & F_{a+\varepsilon_1} \\
& \searrow & \nearrow & \searrow & \nearrow \\
& & & H_a & \\
& \nearrow & \searrow & \nearrow & \searrow \\
G_{a-2\varepsilon_1-\varepsilon_2} & \xrightarrow{\quad} & G_{a-\varepsilon_2} & \xrightarrow{\quad} & G_{a+\varepsilon_2}
\end{array}$$

The persistence module $\mathcal{H}_{\mathbb{R}}$ has the features of both $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$. However, the features that $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ have in common appear twice in $\mathcal{H}_{\mathbb{R}}$, and we want to identify them. Consider an element of $F_{a-\varepsilon_1-2\varepsilon_2}$. There are two ways of sending this element into H_a : either through $F_{a-\varepsilon_1}$ with $f_{a-\varepsilon_1-2\varepsilon_2}^{a-\varepsilon_1}$, or through $G_{a-\varepsilon_2}$ with $\phi_{a-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}}$. We want to identify these two images. Similarly, we want to identify the two images of an element of $G_{a-2\varepsilon_1-\varepsilon_2}$ through $F_{a-\varepsilon_1}$ with $\phi_{a-2\varepsilon_1-\varepsilon_2}^{\mathcal{G},\mathcal{F}}$ or through $G_{a-\varepsilon_2}$ with $g_{a-2\varepsilon_1-\varepsilon_2}^{a-\varepsilon_2}$. These two sets of identifications are required for the diagrams to commute. On the other hand, it only makes sense to identify two elements of H_a if they have the same image in both $F_{a+\varepsilon_1}$ and $G_{a+\varepsilon_2}$. A key property is that pairs that we are allowed to identify include all pairs that we want to identify to make the diagrams commute. We can therefore define a *quotient persistence module* $\tilde{\mathcal{H}}_{\mathbb{R}}$, where \tilde{H}_a is the quotient of H_a by either the minimum set of identifications (defined by the images of $F_{a-\varepsilon_1-2\varepsilon_2}$ and $G_{a-2\varepsilon_1-\varepsilon_2}$) or the maximal set of identifications (defined by the kernels of the homomorphisms $H_a \rightarrow F_{a+\varepsilon_1}$ and $H_a \rightarrow G_{a+\varepsilon_2}$). By construction, $\tilde{\mathcal{H}}_{\mathbb{R}}$ is strongly ε_1 -interleaved with $\mathcal{F}_{\mathbb{R}}$ and strongly ε_2 -interleaved with $\mathcal{G}_{\mathbb{R}}$. \square

The family $(\tilde{\mathcal{H}}_{\mathbb{R}}^s)_{s \in [0, \varepsilon]}$ of persistence modules *interpolates* between $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ in the following sense (see Appendix B for a purely technical proof of this otherwise intuitive result):

Lemma 4.7 $\tilde{\mathcal{H}}_{\mathbb{R}}^0 \simeq \mathcal{F}_{\mathbb{R}}$, $\tilde{\mathcal{H}}_{\mathbb{R}}^\varepsilon \simeq \mathcal{G}_{\mathbb{R}}$, and $\forall s, s' \in [0, \varepsilon]$, $\tilde{\mathcal{H}}_{\mathbb{R}}^s$ and $\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ are strongly $|s - s'|$ -interleaved.

Lemmas 4.6 and 4.7 provide the necessary ingredients for the interpolation argument of [12] to apply, thereby concluding the proof of Theorem 4.4. In cases where the persistence diagrams only have finitely many points off the diagonal Δ , the argument of [12] applies directly. In all other cases, additional technical details must be handled using the Weak Stability Theorem 4.3, as shown in Appendix C.

5 Conclusion

We have shown that the notion of persistence diagram can be extended, and its stability proven, beyond the framework of [12]. Working at algebraic level directly, we have provided a mean of comparing the persistence diagrams of functions defined over different spaces, thus giving a positive answer to an open question from [13]. To achieve our goals, we have introduced several novel concepts and constructions that could become useful theoretical tools. On the practical side, we believe our results may enable new applications of the concept of persistence, as they have already done in the context of scalar field analysis [9].

An important question arising from our work is whether the structure theorem of [21] still holds under our weaker tameness condition (Definition 2.2): is it true that persistence modules with identical persistence diagrams are isomorphic, even if the diagrams have infinitely (yet countably) many points off the diagonal?

Our notion of proximity between persistence modules (Definition 4.2) satisfies the axioms of a metric. In particular, two persistence modules $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ are strongly 0-interleaved if and only if (iff) they are isomorphic. Combined with the correspondence and structure theorems of [21], this fact implies that $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ have identical persistence diagrams iff they are strongly 0-interleaved. An approximate version of this result would be that $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ have ε -close persistence diagrams in the bottleneck distance iff they are strongly ε -interleaved. Theorem 4.4 proves one direction, but the other direction remains open.

Another possible extension of this work would be to multi-dimensional persistence, where generalizations of the concept of persistence diagram have been proposed, most notably the rank invariant of [5]. Stability results for this descriptor exist in several restricted contexts [3, 4], and it would be interesting to see whether they can be extended to a larger setting similar to the one of this paper.

References

- [1] D. Attali, M. Glisse, S. Hornus, F. Lazarus, and D. Morozov. Persistence-sensitive simplification of functions on surfaces in linear time. Submitted to SoCG'09.
- [2] P. Bendich, D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and D. Morozov. Inferring local homology from sampled stratified spaces. In *Proc. 48th Annu. IEEE Sympos. Foundations of Computer Science*, pages 536–546, 2007.
- [3] S. Biasotti, A. Cerri, P. Frosini, D. Giorgi, and C. Landi. Multidimensional size functions for shape comparison. *Journal of Mathematical Imaging and Vision*, 32:161–179, 2008.
- [4] F. Cagliari, B. Di Fabio, and M. Ferri. Title: One-dimensional reduction of multidimensional persistent homology. Research report, University of Bologna, February 2007. LANL arXiv:math/0702713v2 [Math.AT].
- [5] G. Carlsson and A. Zomorodian. The theory of multi-dimensional persistence. In *Proc. 23rd ACM Sympos. on Comput. Geom.*, pages 184–193, 2007.
- [6] G. Carlsson, A. Zomorodian, A. Collins, and L. Guibas. Persistence barcodes for shapes. *International Journal of Shape Modeling*, 11:149–187, 2005.
- [7] F. Chazal and D. Cohen-Steiner. *Geometric Inference*. submitted as a book chapter, 2007.
- [8] F. Chazal, D. Cohen-Steiner, M. Glisse, L. J. Guibas, and S. Y. Oudot. Proximity of persistence modules and their diagrams. Research Report 6568, INRIA, November 2008. <http://hal.inria.fr/inria-00292566/en/>.
- [9] F. Chazal, L. J. Guibas, S. Y. Oudot, and P. Skraba. Analysis of scalar fields over point cloud data. In *Proc. 19th ACM-SIAM Sympos. on Discrete Algorithms*, 2009. Full version available as INRIA research report RR-6576, July 2008 (<http://hal.inria.fr/inria-00294591/en/>).
- [10] C. Chen and D. Freedman. Quantifying homology classes. In *Proc. 25th International Symposium on Theoretical Aspects of Computer Science (STACS)*, 2008.
- [11] D. Cohen-Steiner and H. Edelsbrunner. Inequalities for the curvature of curves and surfaces. In *Proc. 21st Annu. Sympos. on Comput. Geom.*, pages 272–277, 2005.
- [12] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. In *Proc. 21st ACM Sympos. Comput. Geom.*, pages 263–271, 2005.
- [13] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.*, 2008. To appear.
- [14] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. Vines and vineyards by updating persistence in linear time. In *Proc. 22nd Sympos. on Comput. Geom.*, pages 119–126, 2006.
- [15] M. d'Amico, P. Frosini, and C. Landi. Using matching distance in size theory: a survey. *International Journal of Imaging Systems and Technology*, 16(5):154–161, 2006.
- [16] M. d'Amico, P. Frosini, and C. Landi. Natural pseudo-distance and optimal matching between reduced size functions. *Acta Appl. Math.*, 2008. Online: <http://www.springerlink.com/content/cj84327h4n280144/>.

- [17] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.
- [18] H. Edelsbrunner, D. Morozov, and V. Pascucci. Persistence-sensitive simplification of functions on 2-manifolds. In *Proc. 22nd Sympos. on Comput. Geom.*, pages 127–134, 2006.
- [19] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
- [20] D. Morozov. *Homological Illusions of Persistence and Stability*. Ph.d. dissertation, Department of Computer Science, Duke University, August 2008.
- [21] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.

A Proof of Lemma 4.6

For clarity, we let $\varepsilon_1 = s$ and $\varepsilon_2 = \varepsilon - s$. Denote by $\phi_\alpha^{\mathcal{F},\mathcal{G}} : F_\alpha \rightarrow G_{\alpha+\varepsilon}$ and $\phi_\alpha^{\mathcal{G},\mathcal{F}} : G_\alpha \rightarrow F_{\alpha+\varepsilon}$ the homomorphisms provided by Definition 4.2.

We first define the *translated sum* $\mathcal{H}_\mathbb{R}$ by $H_\alpha = F_{\alpha-\varepsilon_1} \oplus G_{\alpha-\varepsilon_2}$. The applications $h_\alpha^\beta : H_\alpha \rightarrow H_\beta$ defined for all $x \in F_{\alpha-\varepsilon_1}$ and $y \in G_{\alpha-\varepsilon_2}$ by $h_\alpha^\beta((x, y)) = (f_{\alpha-\varepsilon_1}^{\beta-\varepsilon_1}(x), g_{\alpha-\varepsilon_2}^{\beta-\varepsilon_2}(y))$ make \mathcal{H} a persistence module.

We define $\phi_\alpha^{\mathcal{F},\mathcal{H}} : F_\alpha \rightarrow H_{\alpha+\varepsilon_1}$ by $\phi_\alpha^{\mathcal{F},\mathcal{H}}(x) = (x, 0)$ and $\phi_\alpha^{\mathcal{G},\mathcal{H}} : G_\alpha \rightarrow H_{\alpha+\varepsilon_2}$ by $\phi_\alpha^{\mathcal{G},\mathcal{H}}(y) = (0, y)$. We also define $\phi_\alpha^{\mathcal{H},\mathcal{F}} : H_\alpha \rightarrow F_{\alpha+\varepsilon_1}$ by $\phi_\alpha^{\mathcal{H},\mathcal{F}}((x, y)) = f_{\alpha-\varepsilon_1}^{\alpha+\varepsilon_1}(x) + \phi_\alpha^{\mathcal{G},\mathcal{F}}(y)$ and $\phi_\alpha^{\mathcal{H},\mathcal{G}} : H_\alpha \rightarrow G_{\alpha+\varepsilon_2}$ by $\phi_\alpha^{\mathcal{H},\mathcal{G}}((x, y)) = \phi_\alpha^{\mathcal{F},\mathcal{G}}(x) + g_{\alpha-\varepsilon_2}^{\alpha+\varepsilon_2}(y)$.


These homomorphisms satisfy some easy properties. In particular, $\phi_{\alpha+\varepsilon_1}^{\mathcal{H},\mathcal{F}} \circ \phi_\alpha^{\mathcal{F},\mathcal{H}} = f_{\alpha+2\varepsilon_1}^{\alpha+\varepsilon_1}$, $\phi_{\alpha+\varepsilon_2}^{\mathcal{H},\mathcal{G}} \circ \phi_\alpha^{\mathcal{G},\mathcal{H}} = g_{\alpha+2\varepsilon_2}^{\alpha+\varepsilon_2}$, $\phi_{\alpha+\varepsilon_1}^{\mathcal{H},\mathcal{G}} \circ \phi_\alpha^{\mathcal{F},\mathcal{H}} = \phi_\alpha^{\mathcal{F},\mathcal{G}}$ and $\phi_{\alpha+\varepsilon_2}^{\mathcal{H},\mathcal{F}} \circ \phi_\alpha^{\mathcal{G},\mathcal{H}} = \phi_\alpha^{\mathcal{G},\mathcal{F}}$. Also notice that $h_{\alpha+\varepsilon_1}^{\beta+\varepsilon_1} \circ \phi_\alpha^{\mathcal{F},\mathcal{H}} = \phi_\beta^{\mathcal{F},\mathcal{H}} \circ f_\alpha^\beta$ and $h_{\alpha+\varepsilon_2}^{\beta+\varepsilon_2} \circ \phi_\alpha^{\mathcal{G},\mathcal{H}} = \phi_\beta^{\mathcal{G},\mathcal{H}} \circ g_\alpha^\beta$. We can also prove that $f_{\alpha+\varepsilon_1}^{\beta+\varepsilon_1} \circ \phi_\alpha^{\mathcal{H},\mathcal{F}} = \phi_\beta^{\mathcal{H},\mathcal{F}} \circ h_\alpha^\beta$. For this, we first verify the relation on the image of $\phi_\alpha^{\mathcal{F},\mathcal{H}}$:

$$\begin{aligned} f_{\alpha+\varepsilon_1}^{\beta+\varepsilon_1} \circ \phi_\alpha^{\mathcal{H},\mathcal{F}} \circ \phi_\alpha^{\mathcal{F},\mathcal{H}} &= f_{\alpha+\varepsilon_1}^{\beta+\varepsilon_1} \circ f_{\alpha-\varepsilon_1}^{\alpha+\varepsilon_1} = f_{\alpha-\varepsilon_1}^{\beta+\varepsilon_1} = f_{\beta-\varepsilon_1}^{\beta+\varepsilon_1} \circ f_{\alpha-\varepsilon_1}^{\beta-\varepsilon_1} \\ &= \phi_\beta^{\mathcal{H},\mathcal{F}} \circ \phi_{\beta-\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ f_{\alpha-\varepsilon_1}^{\beta-\varepsilon_1} = \phi_\beta^{\mathcal{H},\mathcal{F}} \circ h_{\alpha-\varepsilon_1}^\beta \circ \phi_\alpha^{\mathcal{F},\mathcal{H}}. \end{aligned}$$

We then also check it on the image of $\phi_\alpha^{\mathcal{G},\mathcal{H}}$:

$$\begin{aligned} f_{\alpha+\varepsilon_1}^{\beta+\varepsilon_1} \circ \phi_\alpha^{\mathcal{H},\mathcal{F}} \circ \phi_\alpha^{\mathcal{G},\mathcal{H}} &= f_{\alpha+\varepsilon_1}^{\beta+\varepsilon_1} \circ \phi_\alpha^{\mathcal{G},\mathcal{F}} = \phi_{\beta-\varepsilon_2}^{\mathcal{G},\mathcal{F}} \circ g_{\alpha-\varepsilon_2}^{\beta-\varepsilon_2} = \phi_\beta^{\mathcal{H},\mathcal{F}} \circ \phi_{\beta-\varepsilon_2}^{\mathcal{G},\mathcal{H}} \circ g_{\alpha-\varepsilon_2}^{\beta-\varepsilon_2} \\ &= \phi_\beta^{\mathcal{H},\mathcal{F}} \circ h_{\alpha-\varepsilon_2}^\beta \circ \phi_\alpha^{\mathcal{G},\mathcal{H}}. \end{aligned}$$

Since by definition H_α is generated by the images of $\phi_\alpha^{\mathcal{F},\mathcal{H}}$ and $\phi_\alpha^{\mathcal{G},\mathcal{H}}$, these two equalities prove that $f_{\alpha+\varepsilon_1}^{\beta+\varepsilon_1} \circ \phi_\alpha^{\mathcal{H},\mathcal{F}} = \phi_\beta^{\mathcal{H},\mathcal{F}} \circ h_\alpha^\beta$. Similarly, $g_{\alpha+\varepsilon_2}^{\beta+\varepsilon_2} \circ \phi_\alpha^{\mathcal{H},\mathcal{G}} = \phi_\beta^{\mathcal{H},\mathcal{G}} \circ h_\alpha^\beta$.



The only property that is missing for \mathcal{H} to be strongly ε_1 -interleaved with \mathcal{F} is that $\phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \phi_{\alpha}^{\mathcal{H},\mathcal{F}} = h_{\alpha}^{\alpha+2\varepsilon_1}$. And indeed, this equality is usually not satisfied. However, we still try to prove it to identify the obstruction.

As was done previously, we can study this relation separately on the images of $\phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}}$ and $\phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{H}}$. On the image of $\phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}} : \phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \phi_{\alpha}^{\mathcal{H},\mathcal{F}} \circ \phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}} = \phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ f_{\alpha-\varepsilon_1}^{\alpha+\varepsilon_1} = h_{\alpha}^{\alpha+2\varepsilon_1} \circ \phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}}$, the relation is satisfied. On the image of $\phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{H}} : (\phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \phi_{\alpha}^{\mathcal{H},\mathcal{F}} - h_{\alpha}^{\alpha+2\varepsilon_1}) \circ \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{H}} = \phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \phi_{\alpha}^{\mathcal{H},\mathcal{F}} \circ \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{H}} - h_{\alpha}^{\alpha+2\varepsilon_1} \circ \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{H}} = \phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{F}} - \phi_{\alpha+2\varepsilon_1-\varepsilon_2}^{\mathcal{G},\mathcal{H}} \circ g_{\alpha-\varepsilon_2}^{\alpha+2\varepsilon_1-\varepsilon_2}$, so $(\phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \phi_{\alpha}^{\mathcal{H},\mathcal{F}} - h_{\alpha}^{\alpha+2\varepsilon_1})((x, y)) = (\phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{F}}(y), -g_{\alpha-\varepsilon_2}^{\alpha+2\varepsilon_1-\varepsilon_2}(y))$, which is not necessarily null. We thus define

$$\mathbb{F}_{\alpha} = \left\{ (f_{\alpha-\varepsilon_1-2\varepsilon_2}^{\alpha-\varepsilon_1}(x), -\phi_{\alpha-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}}(x)), x \in F_{\alpha-\varepsilon_1-2\varepsilon_2} \right\} \subset H_{\alpha},$$

$$\mathbb{G}_{\alpha} = \left\{ (\phi_{\alpha-2\varepsilon_1-\varepsilon_2}^{\mathcal{G},\mathcal{F}}(x), -g_{\alpha-2\varepsilon_1-\varepsilon_2}^{\alpha-\varepsilon_2}(x)), x \in G_{\alpha-2\varepsilon_1-\varepsilon_2} \right\} \subset H_{\alpha}$$

and let $\tilde{\mathcal{H}}$ be the persistence module defined by $\tilde{H}_{\alpha} = H_{\alpha}/(\mathbb{F}_{\alpha} + \mathbb{G}_{\alpha})$. Call $\pi_{\alpha} : H_{\alpha} \rightarrow \tilde{H}_{\alpha}$ the quotient map.

Let y be an element of \mathbb{F}_{α} . Then $y = (f_{\alpha-\varepsilon_1-2\varepsilon_2}^{\alpha-\varepsilon_1}(x), -\phi_{\alpha-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}}(x))$ for some $x \in F_{\alpha-\varepsilon_1-2\varepsilon_2}$. $h_{\alpha}^{\beta}(y) = (f_{\alpha-\varepsilon_1}^{\beta-\varepsilon_1}(f_{\alpha-\varepsilon_1-2\varepsilon_2}^{\alpha-\varepsilon_1}(x)), g_{\alpha-\varepsilon_2}^{\beta-\varepsilon_2}(-\phi_{\alpha-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}}(x)))$. $h_{\alpha}^{\beta}(y) = (f_{\beta-\varepsilon_1-2\varepsilon_2}^{\beta-\varepsilon_1}(z), -\phi_{\beta-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}}(z))$ where $z = f_{\alpha-\varepsilon_1-2\varepsilon_2}^{\beta-\varepsilon_1-2\varepsilon_2}(x)$, i.e. $h_{\alpha}^{\beta}(y) \in \mathbb{F}_{\beta}$.

Thus $h_{\alpha}^{\beta}(\mathbb{F}_{\alpha}) \subseteq \mathbb{F}_{\beta}$, $h_{\alpha}^{\beta}(\mathbb{G}_{\alpha}) \subseteq \mathbb{G}_{\beta}$ (same proof) and h_{α}^{β} induces a homomorphism $\tilde{h}_{\alpha}^{\beta} : \tilde{H}_{\alpha} \rightarrow \tilde{H}_{\beta}$ such that $\tilde{h}_{\alpha}^{\beta} \circ \pi_{\alpha} = \pi_{\beta} \circ h_{\alpha}^{\beta}$ and the $\tilde{h}_{\alpha}^{\beta}$ satisfy the persistence module property: $\tilde{h}_{\beta}^{\gamma} \circ \tilde{h}_{\alpha}^{\beta} \circ \pi_{\alpha} = \tilde{h}_{\beta}^{\gamma} \circ \pi_{\beta} \circ h_{\alpha}^{\beta} = \pi_{\gamma} \circ h_{\beta}^{\gamma} \circ h_{\alpha}^{\beta} = \pi_{\gamma} \circ h_{\alpha}^{\gamma} = \tilde{h}_{\alpha}^{\gamma} \circ \pi_{\alpha}$, so $\tilde{h}_{\beta}^{\gamma} \circ \tilde{h}_{\alpha}^{\beta} = \tilde{h}_{\alpha}^{\gamma}$.

We define $\tilde{\phi}_{\alpha}^{\mathcal{F},\mathcal{H}} : F_{\alpha} \rightarrow \tilde{H}_{\alpha+\varepsilon_1}$ as $\pi_{\alpha+\varepsilon_1} \circ \phi_{\alpha}^{\mathcal{F},\mathcal{H}}$ and $\tilde{\phi}_{\alpha}^{\mathcal{G},\mathcal{H}} : G_{\alpha} \rightarrow \tilde{H}_{\alpha+\varepsilon_2}$ as $\pi_{\alpha+\varepsilon_2} \circ \phi_{\alpha}^{\mathcal{G},\mathcal{H}}$.

A key property¹ of \mathbb{F}_{α} and \mathbb{G}_{α} is that they are included in the kernels of $\phi_{\alpha}^{\mathcal{H},\mathcal{F}}$ and $\phi_{\alpha}^{\mathcal{H},\mathcal{G}}$. Indeed, $\phi_{\alpha}^{\mathcal{H},\mathcal{F}}((f_{\alpha-\varepsilon_1-2\varepsilon_2}^{\alpha-\varepsilon_1}(x), -\phi_{\alpha-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}}(x))) = f_{\alpha-\varepsilon_1}^{\alpha+\varepsilon_1} \circ f_{\alpha-\varepsilon_1-2\varepsilon_2}^{\alpha-\varepsilon_1}(x) - \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{F}} \circ \phi_{\alpha-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}}(x) = (f_{\alpha-\varepsilon_1-2\varepsilon_2}^{\alpha+\varepsilon_1} - \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{F}} \circ \phi_{\alpha-\varepsilon_1-2\varepsilon_2}^{\mathcal{F},\mathcal{G}})(x) = 0$, $\phi_{\alpha}^{\mathcal{H},\mathcal{F}}((\phi_{\alpha-2\varepsilon_1-\varepsilon_2}^{\mathcal{G},\mathcal{F}}(x), -g_{\alpha-2\varepsilon_1-\varepsilon_2}^{\alpha-\varepsilon_2}(x))) = f_{\alpha-\varepsilon_1}^{\alpha+\varepsilon_1} \circ \phi_{\alpha-2\varepsilon_1-\varepsilon_2}^{\mathcal{G},\mathcal{F}}(x) - \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{F}} \circ g_{\alpha-2\varepsilon_1-\varepsilon_2}^{\alpha-\varepsilon_2}(x) = 0$ and the proof is identical for $\phi_{\alpha}^{\mathcal{H},\mathcal{G}}$. This implies that $\phi_{\alpha}^{\mathcal{H},\mathcal{F}}$ and $\phi_{\alpha}^{\mathcal{H},\mathcal{G}}$ induce homomorphisms $\tilde{\phi}_{\alpha}^{\mathcal{H},\mathcal{F}} : \tilde{H}_{\alpha} \rightarrow F_{\alpha+\varepsilon_1}$ and $\tilde{\phi}_{\alpha}^{\mathcal{H},\mathcal{G}} : \tilde{H}_{\alpha} \rightarrow G_{\alpha+\varepsilon_2}$ such that $\phi_{\alpha}^{\mathcal{H},\mathcal{F}} = \tilde{\phi}_{\alpha}^{\mathcal{H},\mathcal{F}} \circ \pi_{\alpha}$ and $\phi_{\alpha}^{\mathcal{H},\mathcal{G}} = \tilde{\phi}_{\alpha}^{\mathcal{H},\mathcal{G}} \circ \pi_{\alpha}$.

It is easy to check that all the diagrams that commute for \mathcal{H} also commute when \mathcal{H} is replaced by $\tilde{\mathcal{H}}$ (for instance $\tilde{\phi}_{\alpha+\varepsilon_1}^{\mathcal{H},\mathcal{F}} \circ \tilde{\phi}_{\alpha}^{\mathcal{F},\mathcal{H}} = \tilde{\phi}_{\alpha+\varepsilon_1}^{\mathcal{H},\mathcal{F}} \circ \pi_{\alpha+\varepsilon_1} \circ \phi_{\alpha}^{\mathcal{F},\mathcal{H}} = \phi_{\alpha+\varepsilon_1}^{\mathcal{H},\mathcal{F}} \circ \phi_{\alpha}^{\mathcal{F},\mathcal{H}} = f_{\alpha}^{\alpha+2\varepsilon_1}$). The only thing left to prove for \mathcal{F} and $\tilde{\mathcal{H}}$ to be strongly ε_1 -interleaved is that $\tilde{\phi}_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \tilde{\phi}_{\alpha}^{\mathcal{H},\mathcal{F}} = \tilde{h}_{\alpha}^{\alpha+2\varepsilon_1}$. As was shown above, this is equivalent to the nullity of all the elements of \mathbb{G}_{α} , which is obviously the case in \tilde{H}_{α} . We have thus proved that $\tilde{\mathcal{H}}$ is strongly ε_1 -interleaved with \mathcal{F} and (similarly) strongly ε_2 -interleaved with \mathcal{G} . Note also that if \mathcal{F} and \mathcal{G} are tame, then \mathcal{H} is tame as well, and therefore so is $\tilde{\mathcal{H}}$.

B Proof of Lemma 4.7

To avoid confusion, for all $\varepsilon_1 + \varepsilon_2 = \varepsilon$ we denote by $\tilde{\mathcal{H}} = \mathcal{F} \otimes_{\varepsilon_1, \varepsilon_2} \mathcal{G}$ the interpolating persistence module. We shall prove here that the various interpolating modules are related by: $(\mathcal{F} \otimes_{\varepsilon_1, \varepsilon_2 + \varepsilon_3} \mathcal{G}) \otimes_{\varepsilon_2, \varepsilon_3} \mathcal{G} \simeq \mathcal{F} \otimes_{\varepsilon_1 + \varepsilon_2, \varepsilon_3} \mathcal{G}$ when $\varepsilon_1 + \varepsilon_2 + \varepsilon_3 = \varepsilon$.

Let \mathcal{H}, \mathcal{K} and \mathcal{J} be the persistence modules defined by $H_{\alpha} = F_{\alpha-\varepsilon_1} \oplus G_{\alpha-\varepsilon_2-\varepsilon_3}$, $K_{\alpha} = F_{\alpha-\varepsilon_1-\varepsilon_2} \oplus G_{\alpha-\varepsilon_3}$, $J_{\alpha} = H_{\alpha-\varepsilon_2} \oplus G_{\alpha-\varepsilon_3} = (F_{\alpha-\varepsilon_1-\varepsilon_2} \oplus G_{\alpha-2\varepsilon_2-\varepsilon_3}) \oplus G_{\alpha-\varepsilon_3}$ and $\psi_{\alpha} : K_{\alpha} \rightarrow J_{\alpha}$ the function defined by $\psi_{\alpha}((x, y), z) = (x, g_{\alpha-2\varepsilon_2-\varepsilon_3}^{\alpha-\varepsilon_3}(y) + z)$. Define

$$\mathbb{F}_{\alpha}^K = \left\{ (f_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3}^{\alpha-\varepsilon_1-\varepsilon_2}(x), -\phi_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3}^{\mathcal{F},\mathcal{G}}(x)), x \in F_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3} \right\} \subset K_{\alpha}.$$

¹Indeed, if we define \tilde{H}_{α} as $H_{\alpha}/(\text{Ker}(\phi_{\alpha}^{\mathcal{H},\mathcal{F}}) \cap \text{Ker}(\phi_{\alpha}^{\mathcal{H},\mathcal{G}}))$, then the construction works as well.

$$\mathbb{G}_\alpha^K = \left\{ (\phi_{\alpha-2\varepsilon_1-2\varepsilon_2-2\varepsilon_3}^{\mathcal{G},\mathcal{F}}(x), -g_{\alpha-2\varepsilon_1-2\varepsilon_2-2\varepsilon_3}^{\alpha-\varepsilon_3}(x)), x \in G_{\alpha-2\varepsilon_1-2\varepsilon_2-2\varepsilon_3} \right\} \subset K_\alpha,$$

$$\mathbb{F}_\alpha^H = \left\{ (f_{\alpha-\varepsilon_1-2\varepsilon_2-2\varepsilon_3}^{\alpha-\varepsilon_1}(x), -\phi_{\alpha-\varepsilon_1-2\varepsilon_2-2\varepsilon_3}^{\mathcal{F},\mathcal{G}}(x)), x \in F_{\alpha-\varepsilon_1-2\varepsilon_2-2\varepsilon_3} \right\} \subset H_\alpha,$$

$$\mathbb{G}_\alpha^H = \left\{ (\phi_{\alpha-2\varepsilon_1-\varepsilon_2-2\varepsilon_3}^{\mathcal{G},\mathcal{F}}(x), -g_{\alpha-2\varepsilon_1-\varepsilon_2-2\varepsilon_3}^{\alpha-\varepsilon_2-\varepsilon_3}(x)), x \in G_{\alpha-2\varepsilon_1-\varepsilon_2-2\varepsilon_3} \right\} \subset H_\alpha,$$

$$\mathbb{H}_\alpha^J = \left\{ (h_{\alpha-\varepsilon_2-2\varepsilon_3}^{\alpha-\varepsilon_2}(x), -\phi_{\alpha-\varepsilon_2-2\varepsilon_3}^{\mathcal{H},\mathcal{G}}(x)), x \in H_{\alpha-\varepsilon_2-2\varepsilon_3} \right\} \subset J_\alpha,$$

$$\mathbb{G}_\alpha^J = \left\{ (\phi_{\alpha-2\varepsilon_2-2\varepsilon_3}^{\mathcal{G},\mathcal{H}}(x), -g_{\alpha-2\varepsilon_2-2\varepsilon_3}^{\alpha-\varepsilon_3}(x)), x \in G_{\alpha-2\varepsilon_2-2\varepsilon_3} \right\} \subset J_\alpha.$$

The interpolating persistence modules are: $\tilde{\mathcal{H}} = \mathcal{F} \otimes_{\varepsilon_1, \varepsilon_2 + \varepsilon_3} \mathcal{G} = (\tilde{H}_\alpha = H_\alpha / (\mathbb{F}_\alpha^H + \mathbb{G}_\alpha^H))$, $\tilde{\mathcal{K}} = \mathcal{F} \otimes_{\varepsilon_1 + \varepsilon_2, \varepsilon_3} \mathcal{G} = (\tilde{K}_\alpha = K_\alpha / (\mathbb{F}_\alpha^K + \mathbb{G}_\alpha^K))$ and $\tilde{\mathcal{J}} = \tilde{\mathcal{H}} \otimes_{\varepsilon_2, \varepsilon_3} \mathcal{G} = (\tilde{J}_\alpha = J_\alpha / (\mathbb{F}_{\alpha-\varepsilon_2}^H \times 0 + \mathbb{G}_{\alpha-\varepsilon_2}^H \times 0 + \mathbb{H}_\alpha^J + \mathbb{G}_\alpha^J))$ and our goal is to prove that ψ_α induces an isomorphism between \tilde{J}_α and \tilde{K}_α .

\mathbb{H}_α^J can be rewritten as $\mathbb{F}_\alpha^J + \mathbb{G}_\alpha^{J,2}$ where

$$\mathbb{F}_\alpha^J = \left\{ ((f_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3}^{\alpha-\varepsilon_1-\varepsilon_2}(x), 0), -\phi_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3}^{\mathcal{F},\mathcal{G}}(x)), x \in F_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3} \right\},$$

$$\mathbb{G}_\alpha^{J,2} = \left\{ ((0, g_{\alpha-2\varepsilon_2-3\varepsilon_3}^{\alpha-2\varepsilon_2-\varepsilon_3}(x)), -g_{\alpha-2\varepsilon_2-3\varepsilon_3}^{\alpha-\varepsilon_3}(x)), x \in G_{\alpha-2\varepsilon_2-3\varepsilon_3} \right\}.$$

Since $\mathbb{G}_\alpha^{J,2} \subset \mathbb{G}_\alpha^J$, we have $\tilde{J}_\alpha = J_\alpha / (\mathbb{F}_{\alpha-\varepsilon_2}^H \times 0 + \mathbb{G}_{\alpha-\varepsilon_2}^H \times 0 + \mathbb{F}_\alpha^J + \mathbb{G}_\alpha^J)$.

\mathbb{G}_α^J is the kernel of ψ_α . Besides, $\psi_\alpha(J_\alpha) = K_\alpha$, $\psi_\alpha(\mathbb{F}_\alpha^J) = \mathbb{F}_\alpha^K$, $\psi_\alpha(\mathbb{G}_{\alpha-\varepsilon_2}^H \times 0) = \mathbb{G}_\alpha^K$ and $\psi_\alpha(\mathbb{F}_{\alpha-\varepsilon_2}^H \times 0) = \mathbb{F}_\alpha^{K,2}$ where $\mathbb{F}_\alpha^{K,2} \subset \mathbb{F}_\alpha^K$. This proves that ψ_α does induce an isomorphism $\tilde{\psi}_\alpha$ between \tilde{J}_α and \tilde{K}_α . It is easy to check that the persistence homomorphisms $J_\alpha \rightarrow J_\beta$ and $K_\alpha \rightarrow K_\beta$ are the same (through ψ_α and ψ_β). We have thus proved that $(\mathcal{F} \otimes_{\varepsilon_1, \varepsilon_2 + \varepsilon_3} \mathcal{G}) \otimes_{\varepsilon_2, \varepsilon_3} \mathcal{G} \simeq \mathcal{F} \otimes_{\varepsilon_1 + \varepsilon_2, \varepsilon_3} \mathcal{G}$.

This implies that $\mathcal{F} \otimes_{\varepsilon_1, \varepsilon - \varepsilon_1} \mathcal{G}$ and $\mathcal{F} \otimes_{\varepsilon_1', \varepsilon - \varepsilon_1'} \mathcal{G}$ are strongly $|\varepsilon_1 - \varepsilon_1'|$ -interleaved. To conclude the proof, we just need to notice that $\mathcal{F}_\mathbb{R}$ and $\tilde{\mathcal{H}}_\mathbb{R}^0$ are strongly 0-interleaved, which means that $\tilde{\phi}^{\mathcal{F}, \mathcal{H}}$ defines an isomorphism between $\mathcal{F}_\mathbb{R}$ and $\tilde{\mathcal{H}}_\mathbb{R}^0$ with inverse $\tilde{\phi}^{\mathcal{H}, \mathcal{F}}$. The same holds for $\mathcal{G}_\mathbb{R}$ and $\tilde{\mathcal{H}}_\mathbb{R}^\varepsilon$.

$$\begin{array}{ccccccc}
& & F_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3} & \longrightarrow & F_{\alpha-\varepsilon_1-\varepsilon_2} & & \\
& & \searrow & & \searrow & & \\
& & & & & & \\
& & H_{\alpha-\varepsilon_2-2\varepsilon_3} & \longrightarrow & H_{\alpha-\varepsilon_2} & & \\
& & \searrow & & \searrow & & \\
& & & & & & J_\alpha \simeq K_\alpha \\
& & & & & & \\
G_{\alpha-2\varepsilon_2-3\varepsilon_3} & & G_{\alpha-2\varepsilon_1-2\varepsilon_2-2\varepsilon_3} & \longrightarrow & G_{\alpha-2\varepsilon_2-2\varepsilon_3} & \longrightarrow & G_{\alpha-\varepsilon_3}
\end{array}$$

C Proof of Theorem 4.4

Let $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ be two tame persistence modules that are strongly ε -interleaved, and let $(\tilde{\mathcal{H}}_\mathbb{R}^s)_{s \in [0, \varepsilon]}$ be the interpolating family of persistence modules defined in Lemma 4.6. Choose two arbitrary numbers $\alpha \leq \beta$, as well as an arbitrarily small value $\eta > 0$. As mentioned at the end of Section 3.1, any discretization $\tilde{\mathcal{H}}_A^s$ of $\tilde{\mathcal{H}}_\mathbb{R}^s$ has the property that the support of $D\tilde{\mathcal{H}}_A^s$ contains finitely many points in the vertical band $[\alpha, \beta] \times \mathbb{R}$ minus the diagonal Δ . Then, Theorem 3.7 (applied with a sufficiently small discretization step) guarantees

that the support of $D\tilde{\mathcal{H}}_{\mathbb{R}}^s$ contains finitely many points in the area $B_{\eta}^{\alpha,\beta} = ([\alpha, \beta] \times \bar{\mathbb{R}}) \cap \Delta_{+}^{\eta} = \{(u, v) \in \bar{\mathbb{R}}^2 \mid \alpha \leq u \leq \beta \text{ and } v \geq u + 2\eta\}$. Let

$$\delta_{\eta}^{\alpha,\beta}(s) = \frac{1}{4} \min\{\|p - q\|_{\infty} \mid p, q \in B_{\eta}^{\alpha,\beta} \text{ and } p \neq q\} > 0. \quad (3)$$

For any $s, s' \in [0, \varepsilon]$, $\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ is said to be η -close to $\tilde{\mathcal{H}}_{\mathbb{R}}^s$ if $|s - s'| < \delta_{\eta}^{\alpha,\beta}(s)$. The following result, adapted from the Easy Bijection Lemma of [12] to our context, provides a tight bound on the bottleneck distance in the area $B_{\eta}^{\alpha,\beta}$:

Lemma C.1 (Easy Bijection Lemma) *Let $s, s' \in [0, \varepsilon]$ be such that $\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ is η -close to $\tilde{\mathcal{H}}_{\mathbb{R}}^s$. Then, there exists a multi-bijection m between $D\tilde{\mathcal{H}}_{\mathbb{R}}^s$ and $D\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ such that:*

- (i) $\|p - m(p)\|_{\infty} \leq |s' - s|$ for all $p \in D\tilde{\mathcal{H}}_{\mathbb{R}}^s \cap B_{\eta+4|s-s'|}^{\alpha+4|s-s'|, \beta-4|s-s'|}$,
- (ii) $\|p - m(p)\|_{\infty} \leq 3|s' - s|$ for any other point p of $D\tilde{\mathcal{H}}_{\mathbb{R}}^s$.

Proof. By Lemmas 4.6 and 4.7, $\tilde{\mathcal{H}}_{\mathbb{R}}^s$ and $\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ are tame and strongly $|s - s'|$ -interleaved. Therefore, the Weak Stability Theorem 4.3 implies that there exists a multi-bijection $m : D\tilde{\mathcal{H}}_{\mathbb{R}}^s \rightarrow D\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ that moves the points of $D\tilde{\mathcal{H}}_{\mathbb{R}}^s$ by at most $3|s - s'|$ in the l^{∞} -distance, thus proving (ii).

Consider now an arbitrary point $p \in D\tilde{\mathcal{H}}_{\mathbb{R}}^s \cap B_{\eta+4|s-s'|}^{\alpha+4|s-s'|, \beta-4|s-s'|}$. By Eq. (3), for all $q \in D\tilde{\mathcal{H}}_{\mathbb{R}}^s \cap B_{\eta}^{\alpha,\beta}$ we have $\|p - q\|_{\infty} \geq 4\delta_{\eta}^{\alpha,\beta}(s)$, which is greater than $4|s - s'|$ since $\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ is η -close to $\tilde{\mathcal{H}}_{\mathbb{R}}^s$. Furthermore, we have $\|p - q\|_{\infty} > 4|s - s'|$ for all $q \in \bar{\mathbb{R}}^2 \setminus \Delta_{+}^{\eta}$ and all $q \in \bar{\mathbb{R}}^2 \setminus ([\alpha, \beta] \times \bar{\mathbb{R}})$. As a result, the l^{∞} -distance of p to $D\tilde{\mathcal{H}}_{\mathbb{R}}^s \setminus \{p\}$ is greater than $4|s - s'|$. Since $\|p - m(p)\|_{\infty} \leq 3|s - s'|$, the triangle inequality implies that the l^{∞} -distance of $m(p)$ to $D\tilde{\mathcal{H}}_{\mathbb{R}}^s \setminus \{p\}$ is greater than $|s - s'|$. Now, according to Section 4.3.1, we have $d_{\mathcal{H}}^{\infty}(D\tilde{\mathcal{H}}_{\mathbb{R}}^s, D\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}) \leq |s - s'|$, therefore the l^{∞} -distance of $m(p)$ to $D\tilde{\mathcal{H}}_{\mathbb{R}}^{s'}$ is at most $|s - s'|$, which implies that $\|p - m(p)\|_{\infty} \leq |s - s'|$. \square

Consider now an arbitrary positive function $r : [0, \varepsilon] \rightarrow \mathbb{R}$ that is bounded from above by $\delta_{\eta}^{\alpha,\beta}$, that is: $\forall s \in [0, \varepsilon], 0 < r(s) \leq \delta_{\eta}^{\alpha,\beta}(s)$. The family $\{(s - \frac{r(s)}{2}, s + \frac{r(s)}{2})\}_{s \in [0, \varepsilon]}$ forms an open cover of $[0, \varepsilon]$. Since the latter is compact, there exists a finite subcover $\{(s_i - \frac{r(s_i)}{2}, s_i + \frac{r(s_i)}{2})\}_{1 \leq i \leq k}$. Assume without loss of generality that this subcover is minimal, which implies that for all i the open intervals $(s_i - \frac{r(s_i)}{2}, s_i + \frac{r(s_i)}{2})$ and $(s_{i+1} - \frac{r(s_{i+1})}{2}, s_{i+1} + \frac{r(s_{i+1})}{2})$ intersect each other, yielding $|s_{i+1} - s_i| < \frac{r(s_i)}{2} + \frac{r(s_{i+1})}{2} \leq \max\{r(s_i), r(s_{i+1})\}$. In other words, either $\tilde{\mathcal{H}}_{\mathbb{R}}^{s_i}$ is η -close to $\tilde{\mathcal{H}}_{\mathbb{R}}^{s_{i+1}}$, or the other way around. It follows that there exists a multi-bijection $m_i : D\tilde{\mathcal{H}}_{\mathbb{R}}^{s_i} \rightarrow D\tilde{\mathcal{H}}_{\mathbb{R}}^{s_{i+1}}$ satisfying assertions (i) and (ii) of the Easy Bijection Lemma C.1 either with $s = s_i$ and $s' = s_{i+1}$ or the other way around. In addition, the subcover being minimal, we can put $s_0 = 0$ and $s_{k+1} = \varepsilon$ and get that $|s_1 - s_0| < \frac{r(s_1)}{2}$ and $|s_{k+1} - s_k| < \frac{r(s_k)}{2}$, which implies that there also exist multi-bijections $m_0 : D\mathcal{F}_{\mathbb{R}} \rightarrow D\tilde{\mathcal{H}}_{\mathbb{R}}^{s_1}$ and $m_k : D\tilde{\mathcal{H}}_{\mathbb{R}}^{s_k} \rightarrow D\mathcal{G}_{\mathbb{R}}$ satisfying conditions (i) and (ii) above. Let

$$m = m_k \circ m_{k-1} \circ \dots \circ m_1 \circ m_0 \quad (4)$$

be the induced multi-bijection between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{G}_{\mathbb{R}}$. Combining assertion (ii) of Lemma C.1 with the triangle inequality, we obtain that m moves the points of $D\mathcal{F}_{\mathbb{R}}$ by at most 3ε . Moreover, the pairing between points of $D\mathcal{F}_{\mathbb{R}}$ and points of $D\mathcal{G}_{\mathbb{R}}$ defined by m has the following property:

Lemma C.2 *For any point $p \in (D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}) \cap B_{\eta+7\sup r+\varepsilon}^{\alpha+8\varepsilon, \beta-8\varepsilon}$, p and its pair q satisfy $\|p - q\|_{\infty} \leq \varepsilon$.*

Proof. Let $p \in (D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}) \cap B_{\eta+7\sup r+\varepsilon}^{\alpha+8\varepsilon, \beta-8\varepsilon}$. Assume without loss of generality that $p \in D\mathcal{F}_{\mathbb{R}}$, the case $p \in D\mathcal{G}_{\mathbb{R}}$ being symmetric. Let $p = p_0, p_1, \dots, p_k, p_{k+1} = m(p)$ be the images of p through the

sequence of multi-bijections m_i introduced in Eq. (4). We will show that $\|p_i - p_{i+1}\|_\infty \leq |s_i - s_{i+1}|$ for all $i = 0, \dots, k$, which by the triangle inequality implies that $\|p - m(p)\|_\infty \leq \varepsilon$. Assume for a contradiction that there exist some indices $0 \leq i \leq k$ such that $\|p_i - p_{i+1}\|_\infty > |s_i - s_{i+1}|$, and let l be the smallest such index. Then, by the triangle inequality we have $\|p - p_l\|_\infty \leq \sum_{0 \leq i \leq l-1} |s_i - s_{i+1}| \leq \varepsilon$. As a consequence, p_l belongs to $B_{\eta+7 \sup r}^{\alpha+7\varepsilon, \beta-7\varepsilon}$, which is included in $B_{\eta+7|s_l-s_{l+1}|}^{\alpha+7|s_l-s_{l+1}|, \beta-7|s_l-s_{l+1}|}$ since as we saw above $|s_l - s_{l+1}| \leq \min\{\varepsilon, \max\{r(s_l), r(s_{l+1})\}\} \leq \min\{\varepsilon, \sup r\}$. Now, by Lemma C.1 (ii), we have $\|p_l - p_{l+1}\|_\infty \leq 3|s_l - s_{l+1}|$, which implies that p_{l+1} belongs to $B_{\eta+4|s_l-s_{l+1}|}^{\alpha+4|s_l-s_{l+1}|, \beta-4|s_l-s_{l+1}|}$. As a result, whether $\tilde{\mathcal{H}}_{\mathbb{R}}^{s_l}$ be η -close to $\tilde{\mathcal{H}}_{\mathbb{R}}^{s_{l+1}}$ or the other way around, Lemma C.1 (i) implies that $\|p_l - p_{l+1}\|_\infty \leq |s_l - s_{l+1}|$, which contradicts our assumption and thus proves the lemma. \square

We now define a variant m' of the multi-bijection $m : D\mathcal{F}_{\mathbb{R}} \rightarrow D\mathcal{G}_{\mathbb{R}}$ as follows:

1. every point $p \in D\mathcal{F}_{\mathbb{R}}$ such that $\|p - m(p)\|_\infty \leq \varepsilon$ is paired with $m(p)$,
2. every remaining point of $D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}$ is paired with its closest point on the diagonal Δ .

Since the points of Δ have infinite multiplicity, m' defines a multi-bijection between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{G}_{\mathbb{R}}$. In addition, by Lemma C.2, every point of $D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}$ considered at step 2. lies outside $B_{\eta+7 \sup r + \varepsilon}^{\alpha+8\varepsilon, \beta-8\varepsilon}$. Thus, the points in the vertical band $[\alpha + 8\varepsilon, \beta - 8\varepsilon] \times \bar{\mathbb{R}}$ that lie in the closed half-plane $\Delta_+^{\eta+7 \sup r + \varepsilon}$ are moved by at most ε , while the points of $[\alpha + 8\varepsilon, \beta - 8\varepsilon] \times \bar{\mathbb{R}}$ lying below $\Delta_+^{\eta+7 \sup r + \varepsilon}$ are less than $(\eta + 7 \sup r + \varepsilon)$ away from Δ and are therefore moved by less than $\eta + 7 \sup r + \varepsilon$. Since η and r can be chosen arbitrarily small, we obtain:

Corollary C.3 *For any $\nu > 0$, there exists a multi-bijection $D\mathcal{F}_{\mathbb{R}} \rightarrow D\mathcal{G}_{\mathbb{R}}$ that moves the points of $D\mathcal{F}_{\mathbb{R}} \cap ([\alpha + 8\varepsilon, \beta - 8\varepsilon] \times \bar{\mathbb{R}})$ by at most $\varepsilon + \nu$.*

There remains to take care of the bounds $\alpha < \beta$ of the vertical band. Let $(\alpha_n)_{n \in \mathbb{N}}$ and $(\beta_n)_{n \in \mathbb{N}}$ be two sequences of real numbers such that:

- $(\alpha_n)_{n \in \mathbb{N}}$ is strictly decreasing and $\lim \alpha_n = -\infty$,
- $(\beta_n)_{n \in \mathbb{N}}$ is strictly increasing and $\lim \beta_n = +\infty$.

Letting $\nu > 0$ be fixed, for all $n \in \mathbb{N}$ we denote by $m_n : D\mathcal{F}_{\mathbb{R}} \rightarrow D\mathcal{G}_{\mathbb{R}}$ the multi-bijection given by Corollary C.3 for the band $[\alpha_n + 8\varepsilon, \beta_n - 8\varepsilon] \times \bar{\mathbb{R}}$. Since $D\mathcal{G}_{\mathbb{R}} \cap B_{\nu}^{\alpha_0+8\varepsilon, \beta_0-8\varepsilon}$ contains a finite number of points (counted with multiplicities), the set of the restrictions of the multi-bijections m_n to $D\mathcal{F}_{\mathbb{R}} \cap B_{2\nu+\varepsilon}^{\alpha_0+9\varepsilon+\nu, \beta_0-9\varepsilon-\nu}$ is finite. So, taking a subsequence of (α_n, β_n) if necessary, we can assume without loss of generality that all the restrictions of m_n to $D\mathcal{F}_{\mathbb{R}} \cap B_{2\nu+\varepsilon}^{\alpha_0+9\varepsilon+\nu, \beta_0-9\varepsilon-\nu}$ are equal. Symmetrically, since $D\mathcal{F}_{\mathbb{R}} \cap B_{\nu}^{\alpha_0+8\varepsilon, \beta_0-8\varepsilon}$ has finite total multiplicity, there exist only finitely many possible restrictions of the multi-bijections m_n^{-1} to $D\mathcal{G}_{\mathbb{R}} \cap B_{2\nu+\varepsilon}^{\alpha_0+9\varepsilon+\nu, \beta_0-9\varepsilon-\nu}$. Therefore, taking a subsequence if necessary, we can assume without loss of generality that all the restrictions of m_n^{-1} to $D\mathcal{G}_{\mathbb{R}} \cap B_{2\nu+\varepsilon}^{\alpha_0+9\varepsilon+\nu, \beta_0-9\varepsilon-\nu}$ are equal. By the same argument, taking a subsequence if necessary, we can assume that all the restrictions of m_n to $D\mathcal{F}_{\mathbb{R}} \cap B_{2\nu+\varepsilon}^{\alpha_1+9\varepsilon+\nu, \beta_1-9\varepsilon-\nu}$ are equal, and that all the restrictions of m_n^{-1} to $D\mathcal{G}_{\mathbb{R}} \cap B_{2\nu+\varepsilon}^{\alpha_1+9\varepsilon+\nu, \beta_1-9\varepsilon-\nu}$ are equal. We can iterate this process for all $n \in \mathbb{N}$, and by a diagonal argument we obtain a subsequence $(\alpha_{f(n)}, \beta_{f(n)})$ of (α_n, β_n) such that the pairings (p, q) with $p \in D\mathcal{F}_{\mathbb{R}} \cap \Delta_+^{2\nu+\varepsilon}$ or $q \in D\mathcal{G}_{\mathbb{R}} \cap \Delta_+^{2\nu+\varepsilon}$ remain constant throughout and are therefore well-defined at the limit. This set of pairings can now be extended to a multi-bijection between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{G}_{\mathbb{R}}$ by snapping onto the diagonal Δ the points of $D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}$ that lie below $\Delta_+^{2\nu+\varepsilon}$. This multi-bijection moves the points by at most $2\nu + \varepsilon$. Since ν can be chosen arbitrarily small, we obtain that the bottleneck distance between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{G}_{\mathbb{R}}$ is at most ε , thus concluding the proof of Theorem 4.4.