Proximity of Persistence Modules and their Diagrams

Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, Steve Oudot

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Proximity of Persistence Modules and their Diagrams

Frédéric Chazal — David Cohen-Steiner — Marc Glisse — Leonidas Guibas — Steve Oudot

N° 6568

November 2008

Thème COM

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--6568--FR+ENG

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE

INRIA

centre de recherche
SACLAY – ÎLE-DE-FRANCE

Proximity of Persistence Modules and their Diagrams

Frédéric Chazal*, David Cohen-Steiner†, Marc Glisse‡ , Leonidas Guibas§ , Steve Oudot¶

Thème COM — Systèmes communicants
Équipe-Projet Geometrica

Rapport de recherche n° 6568 — November 2008 — 29 pages

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.

Key-words: Topological persistence, Stability, Persistence diagram, Topological and geometric data analysis

* frederic.chazal@inria.fr
† David.Cohen-Steiner@sophia.inria.fr
‡ marc.glisse@gipsa-lab.inpg.fr
§ guibas@cs.stanford.edu
¶ steve.oudot@inria.fr

Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université
4, rue Jacques Monod, 91893 ORSAY Cedex
Téléphone : +33 1 72 92 59 00

Proximité entre les modules de persistance et entre leurs diagrammes

Résumé : La persistance topologique est l'un des concepts clés pour l'étude de fonctions à valeurs réelles définies sur des espaces topologiques. Sa validité repose sur la propriété fondamentale que les diagrammes de persistance de fonctions proches sont similaires. Néanmoins, les résultats de stabilité existants sont restreints au cas des fonctions continues définies sur des espaces triangulables.

Dans cet article nous présentons de nouveaux résultats de stabilité qui ne souffrent pas de cette restriction. De plus, en travaillant directement au niveau algébrique, indépendamment du contexte fonctionnel sous-jacent, notre analyse rend possible la comparaison des diagrammes de persistance de fonctions définies sur des espaces distincts, ouvrant ainsi la voie à une multitude de nouvelles applications de la persistance topologique. Au passage, nous étendons la notion de diagramme de persistance à un contexte plus vaste, nous introduisons les notions de discrétisation d'un module de persistance et de fonction de pixélisation associée, nous définissons une mesure de proximité entre modules de persistance, et enfin nous montrons comment interpoler entre des modules de persistance. Nous donnons ainsi un caractère plus analytique à la persistance topologique. Ces nouveaux outils théoriques constituent une nouvelle approche du sujet, en plus de clarifier et simplifier les preuves, et nous pensons qu'ils peuvent être utiles à la communauté pour les applications à venir.

Mots-clés : Persistance topologique, Stabilité, Diagramme de persistance, Analyse topologique et géométrique de données

1 Introduction

Topological persistence has emerged as a powerful tool for the study of the qualitative and quantitative behavior of real-valued functions defined over topological spaces. Given a space $\mathbb{X}$ equipped with a real-valued function $f : \mathbb{X} \to \mathbb{R}$, persistence encodes the evolution of the topology of the sublevel-sets of f, *i.e.* the sets of type $f^{-1}((-\infty, \alpha]) \subseteq \mathbb{X}$, 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 $\bar{\mathbb{R}}^2$, 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, 18, 19]. 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 $\hat{f} : C \to \mathbb{R}$ 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 persistence diagrams are stable under small perturbations of the functions, since otherwise there would be no way of relating the persistence diagram of $\hat{f}$ (and *a fortiori* the one of its simplification) to the one of 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 $f : \mathbb{X} \to \mathbb{R}$ that filters the underlying space $\mathbb{X}$.

The stability of persistence diagrams was studied in depth by Cohen-Steiner, Edelsbrunner and Harer in their seminal paper [12]. They showed in particular 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, 19]. 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 [21], 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 $\Delta = \{(x, x),\ x \in \bar{\mathbb{R}}\}$. This is all the more a pity 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. Roughly speaking, both continuity and triangulability conditions are removed, and the tameness condition is relaxed. Furthermore, functions are allowed to 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 [18], the persistence diagram of a function $f : \mathbb{X} \to \mathbb{R}$ was derived from the family of homology groups of its sublevel sets $\{H_k(f^{-1}((-\infty, \alpha]))\}_{\alpha \in \mathbb{R}}$, enriched with the family of homomorphisms induced by the canonical inclusion maps $f^{-1}((-\infty, \alpha]) \hookrightarrow f^{-1}((-\infty, \beta])$ for $\alpha \leq \beta$. In [22], 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 $\{f_\alpha^\beta : F_\alpha \to F_\beta\}_{\alpha \leq \beta \in A}$ such that $\forall \alpha \leq \beta \leq \gamma,\ \ f_\alpha^\gamma = f_\beta^\gamma \circ f_\alpha^\beta$ and $f_\alpha^\alpha = \mathrm{id}_{F_\alpha}$, 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 ones used in the 0-dimensional setting of size theory [15], it makes the classical definition of persistence diagram inapplicable. Therefore, we propose a new definition, based on an approximation strategy (Section 3). First, we *discretize* the persistence module $\mathcal{F}_{\mathbb{R}}$ over discrete index sets of type $\alpha_0 + \varepsilon\mathbb{Z}$, for which the persistence diagram can be defined using a variant of the classical construction. Then, the persistence diagram of $\mathcal{F}_{\mathbb{R}}$ is naturally defined as the limit multiset obtained by letting $\varepsilon \to 0$ while keeping α_0 fixed. This limit multiset is shown to be independent of the choice of α_0, and it coincides with the classical notion of persistence diagram 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:

$$\forall \alpha \in \mathbb{R},\ f^{-1}((-\infty, \alpha]) \subseteq g^{-1}((-\infty, \alpha + \varepsilon]) \subseteq f^{-1}((-\infty, \alpha + 2\varepsilon]).$$

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 $\{\widetilde{\mathcal{H}}^s_{\mathbb{R}}\}_{s\in[0,\varepsilon]}$ of persistence modules, with $\widetilde{\mathcal{H}}^0_{\mathbb{R}} = \mathcal{F}_{\mathbb{R}}$ and $\widetilde{\mathcal{H}}^\varepsilon_{\mathbb{R}} = \mathcal{G}_{\mathbb{R}}$, such that $\forall s, s' \in [0, \varepsilon]$, $\widetilde{\mathcal{H}}^s_{\mathbb{R}}$ and $\widetilde{\mathcal{H}}^{s'}_{\mathbb{R}}$ are $|s - s'|$-interleaved.

Our main results are stated in terms of the above distance: first, we provide a very simple and geometrically-flavored proof that the persistence diagrams of any two ε-interleaved persistence modules lie at most 3ε away from each other in the bottleneck distance (Section 4.2); second, taking advantage of this relaxed stability result, we show how the interpolation argument of [12] can be applied at an algebraic level directly, using the above interpolation technique, to 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

2.1 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 $\Delta = \{(x, x),\ x \in \bar{\mathbb{R}}\}$ 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 (see Figure 1).

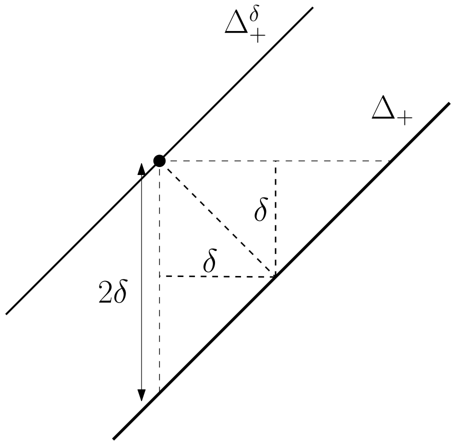

Figure 1: l^∞-distance to the diagonal.

The main mathematical objects considered in topological persistence are the so-called persistence diagrams (or barcodes) that are *multisets* in $\bar{\mathbb{R}}^2$, *i.e.* subsets of $\overline{\mathbb{R}}^2$ where to each point p of D is associated a multiplicity $\mathrm{mult}(p) \in \mathbb{N} \cup \{+\infty\}$. The *support* of D, i.e. the subset considered without the multiplicities, is denoted by $|D|$. Equivalently, D can be represented as a disjoint union

$$D = \bigcup_{p \in |D|} \coprod_{i=1}^{\mathrm{mult}(p)} p.$$

A *multi-bijection* m between two multisets D, D' is a bijection

$$m : \bigcup_{p \in |D|} \coprod_{i=1}^{\mathrm{mult}(p)} p \to \bigcup_{p' \in |D'|} \coprod_{i=1}^{\mathrm{mult}(p')} p'.$$

Given two multisets D and D', the Hausdorff distance between their support is defined by:

$$\mathrm{d}_{\mathcal{H}}^\infty(|D|, |D'|) = \max \left\{ \sup_{p \in |D|} \inf_{p' \in |D'|} \|p - p'\|_\infty,\ \sup_{p' \in |D'|} \inf_{p \in |D|} \|p - p'\|_\infty \right\}$$

For simplicity, we abuse notations and write $\mathrm{d}_{\mathcal{H}}^\infty(D, D')$ for the Hausdorff distance $\mathrm{d}_{\mathcal{H}}^\infty(|D|, |D'|)$. Note that this is not a distance between multisets, but between their supports. A relevant distance between multisets is the *Bottleneck distance* d_B^∞, introduced in [12]:

$$\mathrm{d}_B^\infty(D, D') = \inf_{\substack{m: D \to D' \\ \text{multi-bijection}}} \left(\sup_{p \in D} \|p - m(p)\|_\infty \right)$$

In order to avoid heavy notations in the rest of the paper, we use the same notation for a multiset and for its support. The distinction will be made clear whenever necessary.

2.2 Filtrations and persistence modules

The homology theory used in the paper is singular homology with coefficients in a commutative ring R with unity (see [20] for an introduction to the subject), which will be assumed to be a field and omitted in our notations.

Topological persistence has been introduced as a useful and robust tool to encode the topological behavior of real valued functions defined on topological spaces. Its construction strongly relies on the filtrations defined by the sublevel sets of the considered functions.

Definition 2.1 *Given a topological space* $\mathbb{X}$ *and a function* $f : \mathbb{X} \to \mathbb{R}$*, the* sublevel-sets filtration *of* f *is defined by:*
i) the family of subspaces $\{\mathbb{X}^f_\alpha\}_{\alpha\in\mathbb{R}}$ *where* $\forall \alpha \in \mathbb{R}$, $\mathbb{X}^f_\alpha = f^{-1}((-\infty, \alpha]) \subseteq \mathbb{X}$,
ii) the family of canonical inclusion maps $i_\alpha^{\alpha'} : \mathbb{X}^f_\alpha \hookrightarrow \mathbb{X}^f_{\alpha'}$, $\forall \alpha \leq \alpha'$.

The family of inclusion maps between the sublevel-sets of f induces a family of homomorphisms between their kth homology groups, known as the *kth persistence module* of the sublevel-sets filtration of f [22]. Persistence modules can in fact be defined independently of the sublevel sets filtration of a function, in a formal way that generalizes the notion of persistence to a larger setting than the one considered in [17, 22]:

Definition 2.2 *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* $\{f_\alpha^{\alpha'} : F_\alpha \to F_{\alpha'}\}_{\alpha\leq\alpha'\in A}$ *of homomorphisms between the modules, such that:*

$$\forall \alpha \leq \alpha' \leq \alpha'' \in A, \quad f_\alpha^{\alpha''} = f_{\alpha'}^{\alpha''} \circ f_\alpha^{\alpha'} \quad \textit{and} \quad f_\alpha^\alpha = id_{F_\alpha}.$$

In our context, since the ring R is assumed to be a fixed field, 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. An important type of discrete persistence module is the one defined over some periodic index set of the form $A = \alpha_0 + \varepsilon\mathbb{Z}$, where $\alpha_0 \in \mathbb{R}$ and $\varepsilon > 0$ are fixed parameters. Such a persistence module is said to be *ε-periodic.*

2.3 Tameness and topological persistence

Definition 2.3 ([12]) *Let* $\mathbb{X}$ *be a topological space and* $f : \mathbb{X} \to \mathbb{R}$ *a function. A* homological critical value *of* f *is a real number* α *for which there exists an integer* k *such that for all sufficiently small* $\varepsilon > 0$ *the maps* $H_k(\mathbb{X}^f_{\alpha-\varepsilon}) \to H_k(\mathbb{X}^f_{\alpha+\varepsilon})$ *induced by inclusion are not isomorphisms. The function* f *is said to be* tame *if it has only a finite number of homological critical values, and if furthermore the homology groups* $H_k(\mathbb{X}^f_\alpha)$ *are finite-dimensional for all* $k \in \mathbb{N}$ *and all* $\alpha \in \mathbb{R}$.

The persistence diagram of a tame function $f : \mathbb{X} \to \mathbb{R}$ is defined as a multiset of $\bar{\mathbb{R}}^2$ in the following way. Let $\alpha_1 < \alpha_2 < \cdots < \alpha_n$ be the homological critical values of f, and let $\beta_0, \cdots, \beta_n$ be an interleaved sequence of real numbers, that is: for all $i = 1, \cdots n$, $\beta_{i-1} < \alpha_i < \beta_i$. For technical reasons, we let

$\beta_{-1} = \alpha_0 = -\infty$ and $\beta_{n+1} = \alpha_{n+1} = +\infty$. For any $0 \leq i < j \leq n+1$, the multiplicity of point (α_i, α_j) is defined by:

$$\begin{aligned} \mathrm{mult}(\alpha_i, \alpha_j) \quad = \quad & \mathrm{rank}\ (H_*(\mathbb{X}^f_{\beta_{i-1}}) \to H_*(\mathbb{X}^f_{\beta_j})) \\ & - \mathrm{rank}\ (H_*(\mathbb{X}^f_{\beta_i}) \to H_*(\mathbb{X}^f_{\beta_j})) \\ & + \mathrm{rank}\ (H_*(\mathbb{X}^f_{\beta_i}) \to H_*(\mathbb{X}^f_{\beta_{j-1}})) \\ & - \mathrm{rank}\ (H_*(\mathbb{X}^f_{\beta_{i-1}}) \to H_*(\mathbb{X}^f_{\beta_{j-1}})), \end{aligned} \tag{1}$$

where $H_*(.)$ denotes the direct sum of the homology groups of all dimensions: $H_*(.) = \bigoplus_{k \in \mathbb{Z}} H_k(.)$, and where the arrows are the homomorphisms induced by inclusions between the sublevel sets[1] of f. The *persistence diagram* of f, or Df for short, is the multiset defined as the union of all points (α_i, α_j) where $0 \leq i < j \leq n+1$, with multiplicities $\mathrm{mult}(\alpha_i, \alpha_j)$, and of the diagonal Δ with infinite multiplicity. The main prior result on persistence diagrams is the following stability result:

Theorem 2.4 ([12]) *Let $\mathbb{X}$ be a triangulable space and let $f, g : \mathbb{X} \to \mathbb{R}$ be two continuous tame functions. Then*

$$d_B^\infty(Df, Dg) \leq \|f - g\|_\infty = \sup_{x \in \mathbb{X}} |f(x) - g(x)|$$

It turns out that the notions of multiplicity of a point and of persistence diagram do not actually require the number of homological critical values to be finite. It is indeed sufficient to ensure that the homology groups $H_*(\mathbb{X}^f_\alpha)$ are finite-dimensional for all $\alpha \in \mathbb{R}$:

Definition 2.5 *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 we have $\mathrm{rank}\ f_\alpha^{\alpha'} < +\infty$ for all $\alpha' \geq \alpha$. From now on, and until the end of the paper, *tameness* will be understood as in Definition 2.5. 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. In a previous version of this technical report [8], we showed that persistence diagrams can even be defined under a weaker condition, called *δ-tameness*, which states that the rank of $f_\alpha^{\alpha'}$ is finite whenever $\alpha' - \alpha > \delta$. The results of this paper can be extended to the δ-tame setting modulo some additional technicalities [8].

3 Discretizing persistence modules

To define the persistence diagram of a tame persistence module, we proceed in two steps: first, we consider the discretized persistence modules obtained by restricting the family of vector spaces F_α to discrete sequences of indices. The persistence diagram of such discretizations is defined in a similar way as in the classical setting. Second, we show that the persistence diagram of the whole

[1] By convention, if $i = -\infty$ then $\mathrm{rank}\ (H_*(\mathbb{X}^f_{\beta_i}) \to H_*(\mathbb{X}^f_{\beta_j})) = \lim_{\beta \to -\infty} \mathrm{rank}\ (H_*(\mathbb{X}^f_\beta) \to H_*(\mathbb{X}^f_{\beta_j}))$, and if $j = +\infty$ then $\mathrm{rank}\ (H_*(\mathbb{X}^f_{\beta_i}) \to H_*(\mathbb{X}^f_{\beta_j})) = \lim_{\beta \to +\infty} \mathrm{rank}\ (H_*(\mathbb{X}^f_{\beta_i}) \to H_*(\mathbb{X}^f_\beta))$.

persistence module $\mathcal{F}_A$ is obtained as a well-defined limit of the persistence diagrams of its discretizations.

Definition 3.1 *Let $\mathcal{F}_A$ be a persistence module, and let B be a subset of A. The* restriction *of $\mathcal{F}_A$ to 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. If $\mathcal{F}_B$ is discrete (namely, if B is discrete with no accumulation point), then $\mathcal{F}_B$ is called the B*-discretization *of $\mathcal{F}_A$.*

To every discrete set B with no accumulation point is associated a *pixelization grid* $\Gamma_B \subset \bar{\mathbb{R}}^2$ whose vertices are the points of type $(\beta,\ \beta')$ 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 $[-\infty, \beta_i] \times \{+\infty\}$. Associated to the grid Γ_B is a B-*pixelization map* $\mathrm{pix}_B : \Delta_+ \to \Gamma_B \cup \Delta$, defined in Eq. (2) below, where the notation $\lceil\alpha\rceil_{\bar{B}}$ stands for $\min \bar{B} \cap [\alpha, +\infty]$, the element of $\bar{B}$ immediately following[2] α:

$$\forall \alpha \leq \alpha' \in \bar{\mathbb{R}},\ \mathrm{pix}_B(\alpha,\ \alpha') = \begin{cases} (\lceil\alpha\rceil_{\bar{B}},\ \lceil\alpha'\rceil_{\bar{B}}), & \text{if } \lceil\alpha\rceil_{\bar{B}} < \lceil\alpha'\rceil_{\bar{B}}; \\ \left(\frac{\alpha+\alpha'}{2},\ \frac{\alpha+\alpha'}{2}\right), & \text{if } \lceil\alpha\rceil_{\bar{B}} = \lceil\alpha'\rceil_{\bar{B}}. \end{cases} \tag{2}$$

Figure 2 illustrates the effect of the pixelization map pix_B, which performs the following snapping operations: each point of Δ_+ lying in a cell $\mathcal{C}$ of the grid 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.

3.1 Persistence diagrams of discrete tame persistence modules

Let $\mathcal{F}_B$ be a discrete tame persistence module. For the sake of 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 made possible by the fact that B has no accumulation point. Then, $\bar{I} = I \cup \{\inf I, +\infty\}$ indexes $\bar{B}$. By convention, when $j = +\infty$, we let $\mathrm{rank}\ f_{\beta_i}^{\beta_j} = \mathrm{rank}\ f_{\beta_i}^{\beta_m}$ if I has a maximum element $m \in \mathbb{Z}$, and $\mathrm{rank}\ f_{\beta_i}^{\beta_j} = \lim_{k\to+\infty} \mathrm{rank}\ f_{\beta_i}^{\beta_k}$ otherwise. Such a limit always exists since the general inequality $\mathrm{rank}\,(g \circ f) \leq \mathrm{rank}\ f$ implies that for any fixed i, the map $k \mapsto \mathrm{rank}\ f_{\beta_i}^{\beta_k}$ is non-increasing. In particular, i being fixed, since the ranks are non negative integers, $k \mapsto \mathrm{rank}\ f_{\beta_i}^{\beta_k}$ is constant for sufficiently large k. Similarly, if $\inf I = -\infty$, then for all $j \in I \cup \{+\infty\}$ we let $\mathrm{rank}\ f_{\beta_{\inf I}}^{\beta_j} = \lim_{k\to-\infty} \mathrm{rank}\ f_{\beta_k}^{\beta_j}$.

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:*

[2]This element is well-defined because B has no accumulation point. It is equal to α whenever $\alpha \in \bar{B}$.

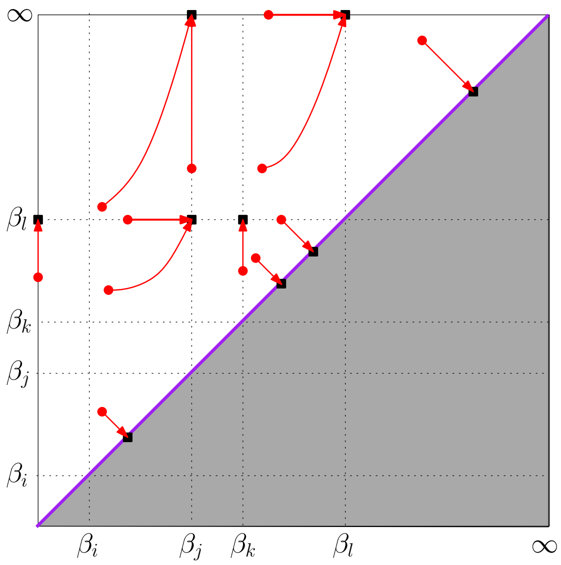

Figure 2: The pixelization map pix_B. Here, $\beta_i < \beta_j < \beta_k < \beta_l$ belong to B, and $\inf B = -\infty$.

(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 $(\beta_i,\ \beta_j)$ with $i < j \in \bar{I}$ has multiplicity*

$$\mathrm{mult}(\beta_i,\ \beta_j) = \mathrm{rank}\ f_{\beta_i}^{\beta_{j-1}} - \mathrm{rank}\ f_{\beta_i}^{\beta_j} + \mathrm{rank}\ f_{\beta_{i-1}}^{\beta_j} - \mathrm{rank}\ f_{\beta_{i-1}}^{\beta_{j-1}}$$

if $i > \inf I$, and

$$\mathrm{mult}(\beta_i,\ \beta_j) = \mathrm{rank}\ f_{\beta_i}^{\beta_{j-1}} - \mathrm{rank}\ f_{\beta_i}^{\beta_j}$$

if $i = \inf I$.

Definition 3.2 is illustrated in Figure 3.

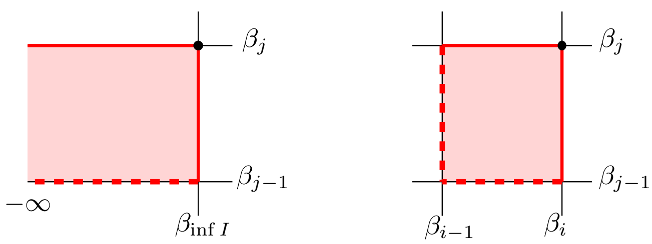

Figure 3: The multiplicity of a node (β_i, β_j) is fully determined by the ranks of the homomorphisms corresponding to the corners of the bottom left cell incident to (β_i, β_j). The number of such corners is two or four, depending on whether $i = \inf I$ (left) or $i > \inf I$ (right).

It follows from our tameness condition (Definition 2.5) 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 4:

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

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

Furthermore, for all $j_1 < j_2 \in \bar{I}$, we have

$$\sum_{j_1 < j \leq j_2} \mathrm{mult}(\beta_{\inf I},\ \beta_j) \ = \ \mathrm{rank}\ f_{\beta_{\inf I}}^{\beta_{j_1}} - \mathrm{rank}\ f_{\beta_{\inf I}}^{\beta_{j_2}}.$$

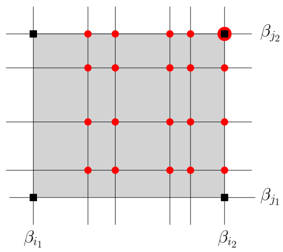

Figure 4: The sum of the multiplicities of the nodes (red disks) contained in the box $(\beta_{i_1},\ \beta_{i_2}] \times (\beta_{j_1},\ \beta_{j_2}]$ is equal to the alternate sum of the ranks of the linear maps corresponding to the vertices (black squares) of the box.

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. Furthermore,

Corollary 3.4 *For any $\beta_1 \leq \beta_2 \in \mathbb{R}$, the total multiplicity (and therefore the support) of the points of $D\mathcal{F}_B$ contained in the vertical band $[\beta_1, \beta_2] \times \bar{\mathbb{R}}$ minus the diagonal Δ is finite.*

Proof. By definition, the part of $[\beta_1, \beta_2] \times \bar{\mathbb{R}}$ lying below Δ contains no point of $D\mathcal{F}_B$. In addition, since B has no accumulation point, the points of $D\mathcal{F}_B \cap ([\beta_1, \beta_2] \times \bar{\mathbb{R}})$ lying above Δ are covered by a finite union of half-open upper-left quadrants. Therefore, their total multiplicity is finite. □

Note finally that, because B has no accumulation point 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 $\mathbb{R} \times \{\pm\infty\}$. Therefore, $|D\mathcal{F}_B| \setminus \Delta$ cannot have any accumulation point.

3.2 Persistence diagrams of arbitrary tame persistence modules

Consider now an arbitrary tame persistence module $\mathcal{F}_A$. In order to define its persistence diagram, we need to compare the persistence diagrams of its various discretizations:

Theorem 3.5 *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[3] $\mathrm{pix}_C \circ \mathrm{pix}_B^{-1}$, which by Theorem 3.5 defines a multi-bijection between $D\mathcal{F}_B$ and $D\mathcal{F}_C$, also moves the points by at most ε. Therefore,

Corollary 3.6 *For any ε-periodic discretizations $\mathcal{F}_B$ and $\mathcal{F}_C$ of $\mathcal{F}_A$, we have* $\mathrm{d}_B^\infty(D\mathcal{F}_B, D\mathcal{F}_C) \leq \varepsilon$.

More generally, in view of the definition of pixelization map given in Eq. (2), we have $\mathrm{d}_B^\infty(D\mathcal{F}_B, D\mathcal{F}_C) \leq \varepsilon$ 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 3.6 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 3.5. Consider the discretization of $\mathcal{F}$ over the discrete union family $B \cup C$. Our approach consists in considering $\mathcal{F}_B$ and $\mathcal{F}_C$ as two discretizations of $\mathcal{F}_{B\cup C}$, and in showing 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_B). The following lemma, illustrated in Figure 5, presents the key argument (for clarity, $\mathrm{mult}_B(p)$ denotes the multiplicity of p in the diagram $D\mathcal{F}_B$):

Lemma 3.7 *Let $\mathcal{C}$ be a cell of the grid Γ_B that does not intersect Δ, and let p be its upper-right corner. Then,*

$$\mathrm{mult}_B(p) = \sum_{q\in|D\mathcal{F}_{B\cup C}|\cap\mathcal{C}} \mathrm{mult}_{B\cup C}(q),$$

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 p while preserving the total multiplicity, thus defining a multi-bijection between $D\mathcal{F}_{B\cup C} \cap \mathcal{C}$ and $D\mathcal{F}_B \cap \mathcal{C}$.*

Proof. Let $\beta_i < \beta_j$ be the coordinates of p. Assume without loss of generality that $i > \inf I$, the case $i = \inf I$ being similar. By definition of the multiplicity, we have:

$$\mathrm{mult}_B(\beta_i, \beta_j) = \mathrm{rank}\ f_{\beta_i}^{\beta_{j-1}} - \mathrm{rank}\ f_{\beta_i}^{\beta_j} + \mathrm{rank}\ f_{\beta_{i-1}}^{\beta_j} - \mathrm{rank}\ f_{\beta_{i-1}}^{\beta_{j-1}},$$

which by Lemma 3.3 (applied to $\mathcal{F}_{B\cup C}$) is equal to $\sum_{q\in|D\mathcal{F}_{B\cup C}|\cap\mathcal{C}} \mathrm{mult}_{B\cup C}(q)$. □

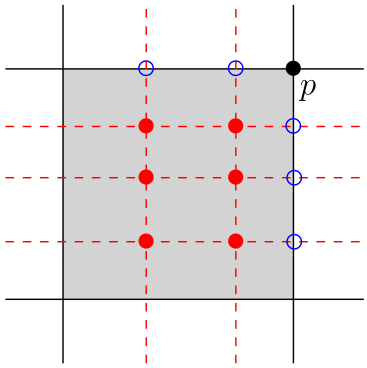

Figure 5: The snapping argument: the multiplicity of point p in $D\mathcal{F}_B$ is equal to the sum of multiplicities of the marked points of $D\mathcal{F}_{B\cup C}$ (including p itself). The pixelization map pix_B snaps all marked points onto p, thus defining a multi-bijection between $D\mathcal{F}_{B\cup C}$ and $D\mathcal{F}_B$ within the cell.

Observe now that, for every cell $\mathcal{C}$ of the grid Γ_B that intersects Δ, 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 Lemma 3.7 or the above remark 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 $D\mathcal{F}_B$. This concludes the proof of Theorem 3.5. □

We are now ready to define the persistence diagram of $\mathcal{F}_A$. For the sake of simplicity, we assume from now on and until the end of Section 3 that $A = \mathbb{R}$. The case of an arbitrary index set $A \subseteq \mathbb{R}$ can be handled in a similar way, at the price of a significant increase in technicality.

We define the persistence diagram of $\mathcal{F}_\mathbb{R}$ by a subdivision process, starting with an arbitrary discrete subset $B_0 \subset \mathbb{R}$ with no accumulation point 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 point 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 3.5. 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.

Consider now any other nested family $\{B'_n\}_{n\in\mathbb{N}}$ of discrete subsets of $\mathbb{R}$ with no accumulation point, such that $\forall n \in \mathbb{N}$ the set B_n forms a right 2^{-n}-cover of $\mathbb{R}$. Then, Corollary 3.6 implies that $\mathrm{d}_B^\infty(D\mathcal{F}_{B_n}, D\mathcal{F}_{B'_n}) \leq 2^{-n}$. As a result, the limit multiset obtained from $\{B'_n\}_{n\in\mathbb{N}}$ is identical to M.

[3]Here, pix_B^{-1} is to be understood as the inverse of the restriction of pix_B to $D\mathcal{F}_{B\cup C}$.

Definition 3.8 *The limit multiset M obtained by the above subdivision process is called the* persistence diagram *of the tame filtration $\mathcal{F}_{\mathbb{R}}$, denoted $D\mathcal{F}_{\mathbb{R}}$. It is independent of the choice of the nested family $\{B_n\}_{n\in\mathbb{N}}$ of discrete subsets of $\mathbb{R}$.*

Another easy yet important consequence of the above subdivision process is that pixelization maps relate the persistence diagram of $\mathcal{F}_{\mathbb{R}}$ to the ones of its discretizations:

Theorem 3.9 *Let $\mathcal{F}_{\mathbb{R}}$ be a tame persistence module, and let B be a discrete subset of $\mathbb{R}$ with no accumulation point. Then, the restriction of pix_B to $D\mathcal{F}_{\mathbb{R}}$ defines a multi-bijection between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{F}_B$.*

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 point, 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 3.8, 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 $B\subseteq B_0\subseteq B_1\subseteq\cdots\subseteq B_n$, we deduce that

$$\forall n\in\mathbb{N},\ \mathrm{pix}_B = \mathrm{pix}_B\circ\mathrm{pix}_{B_0}\circ\mathrm{pix}_{B_1}\circ\cdots\circ\mathrm{pix}_{B_{n-1}}.$$

Therefore, by Theorem 3.5, 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 special case where $\mathcal{F}_B$ is an ε-periodic family, we obtain the following key result:

Corollary 3.10 *Let $\mathcal{F}_{\mathbb{R}}$ be a tame persistence module. For any ε-periodic discretization $\mathcal{F}_B$ of $\mathcal{F}_{\mathbb{R}}$, we have $\mathrm{d}_B^\infty(D\mathcal{F}_{\mathbb{R}}, D\mathcal{F}_B)\leq\varepsilon$.*

Remark 1 *In the case where $\mathcal{F}_{\mathbb{R}}$ is the persistence module of the sublevel-sets filtration of some function $f:\mathbb{X}\to\mathbb{R}$ that is tame in the sense of Definition 2.3, its persistence diagram Df as defined in [12] coincides with its persistence diagram $D\mathcal{F}_{\mathbb{R}}$ in the sense of Definition 3.8. To prove this, it is enough to build a growing family $\{B_n\}_{n\in\mathbb{N}}$ of discrete sets such that the homological critical values $\alpha_1,\cdots\alpha_m$ of f are included in none of the B_i. Such a construction is feasible because the set of homological critical values of f is finite and the sets B_i are discrete. Thus, at any step n of the subdivision process, every point (α_i,α_j) of $Df\setminus\Delta$ is contained in the interior of some cell $\mathcal{C}_{i,j,n}$ of the grid Γ_{B_n}, while $D\mathcal{F}_{B_n}\cap\mathcal{C}_{i,j,n}$ is just a finite number of copies of the upper-right corner of $\mathcal{C}_{i,j,n}$. Furthermore, for sufficiently large n, $\mathcal{C}_{i,j,n}$ contains only one point of the support of Df, namely (α_i,α_j), and the definitions of multiplicity given in Eq. (1) and Definition 3.2 imply that the multiplicity of (α_i,α_j) is equal to the multiplicity of the upper right corner of $\mathcal{C}_{i,j,n}$. Since this is true for all sufficiently large $n\in\mathbb{N}$, we have $Df = D\mathcal{F}_{\mathbb{R}}$.*

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 4.6) has a simple and geometrically-flavored proof, and it is instrumental in proving the stronger stability result (Theorem 4.9).

4.1 Interleaving persistence modules

To emphasize the intuition underlying our definitions, we first consider the case of persistence modules associated with sublevel sets filtrations of functions. This special case will motivate our notion of proximity between general persistence modules.

Definition 4.1 *Given two real-valued functions $f, g : \mathbb{X} \to \mathbb{R}$ and $\varepsilon > 0$,*

- *f, g are* weakly ε-interleaved *if $\exists \alpha_0 \in \mathbb{R}$ such that $\forall \alpha \in \alpha_0 + 2\varepsilon\mathbb{Z}$, $X^f_\alpha \subseteq \mathbb{X}^g_{\alpha+\varepsilon} \subseteq \mathbb{X}^f_{\alpha+2\varepsilon}$,*
- *f, g are* strongly ε-interleaved *if $\forall \alpha \in \mathbb{R}$, $\mathbb{X}^f_\alpha \subseteq \mathbb{X}^g_{\alpha+\varepsilon} \subseteq \mathbb{X}^f_{\alpha+2\varepsilon}$.*

The following lemma relates the above notions of interleaved filtrations to the distance (in the l^∞ norm) between the functions:

Lemma 4.2 *Let $f, g : \mathbb{X} \to \mathbb{R}$ and $\varepsilon > 0$. Then,*

(i) if f, g are strongly ε-interleaved, then they are weakly ε-interleaved,

(ii) if f, g are weakly ε-interleaved, then they are strongly 3ε-interleaved,

(iii) f, g are strongly ε-interleaved if and only if $\|f - g\|_\infty \leq \varepsilon$.

Proof. Assertion (i) is a direct consequence of Definition 4.1.

Assume now that f, g are weakly ε-interleaved, and let $\alpha_0 \in \mathbb{R}$ be as in Definition 4.1 (i). Given $\alpha \in \mathbb{R}$, let $n \in \mathbb{Z}$ be such that $\alpha_0 + 2(n-1)\varepsilon < \alpha \leq \alpha_0 + 2n\varepsilon$. Then, we have $\mathbb{X}^f_\alpha \subseteq \mathbb{X}^f_{\alpha_0+2n\varepsilon} \subseteq \mathbb{X}^g_{\alpha_0+(2n+1)\varepsilon}$, which is included in $\mathbb{X}^g_{\alpha+3\varepsilon}$ since $\alpha > \alpha_0 + 2(n-1)\varepsilon$. Symmetrically, we have $\mathbb{X}^g_\alpha \subseteq \mathbb{X}^f_{\alpha+3\varepsilon}$. Since this is true for all $\alpha \in \mathbb{R}$, f and g are strongly 3ε-interleaved, which proves (ii).

Assume now that f, g are strongly ε-interleaved. For all points $x \in \mathbb{X}$, we have

$$x \in f^{-1}((-\infty, f(x)]) = \mathbb{X}^f_{f(x)} \subseteq \mathbb{X}^g_{f(x)+\varepsilon} = g^{-1}((-\infty, f(x) + \varepsilon]),$$

which implies that $g(x) \leq f(x) + \varepsilon$. By symmetry, we also have $f(x) \leq g(x) + \varepsilon$. Hence, $|f(x) - g(x)| \leq \varepsilon$. Since this is true for all $x \in \mathbb{X}$, we deduce that $\|f - g\|_\infty \leq \varepsilon$. Conversely, assume that $\|f - g\|_\infty \leq \varepsilon$. Then, for all $\alpha \in \mathbb{R}$ and all $x \in \mathbb{X}^f_\alpha$, we have $f(x) \leq \alpha$ and therefore $g(x) \leq f(x) + \varepsilon \leq \alpha + \varepsilon$, which implies that $x \in \mathbb{X}^g_{\alpha+\varepsilon}$. Symmetrically, we also have $\mathbb{X}^g_\alpha \subseteq \mathbb{X}^f_{\alpha+\varepsilon}$. Since this is true for all $\alpha \in \mathbb{R}$, we conclude that f, g are strongly ε-interleaved, which proves (iii). □

When f and g are weakly ε-interleaved, the canonical inclusions between their sublevel sets induce the following commutative diagram between the 2ε-discretizations of their kth persistence modules, where $F_\alpha = H_k(\mathbb{X}^f_\alpha)$ and $G_\alpha =$

$H_k(\mathbb{X}^g_\alpha)$ denote the kth homology groups of the sublevel sets, and where the arrows represent the homomorphisms induced by inclusions at kth homology level:

$$\begin{tikzcd}
 & \cdots \arrow[r] & F_{\alpha_0+2n\varepsilon} \arrow[rr] \arrow[dr] & & F_{\alpha_0+(2n+2)\varepsilon} \arrow[r] & \cdots \\
\cdots \arrow[r] & G_{\alpha_0+(2n-1)\varepsilon} \arrow[ur] \arrow[rr] & & G_{\alpha_0+(2n+1)\varepsilon} \arrow[ur] \arrow[r] & \cdots &
\end{tikzcd} \tag{3}$$

The persistence modules of f and g are said to be *weakly ε-interleaved* if the above diagram is commutative for some fixed $\alpha_0 \in \mathbb{R}$.

In addition, when f, g are strongly ε-interleaved, the following diagrams induced at kth homology level commute for all $\alpha \leq \alpha' \in \mathbb{R}$:

$$\begin{tikzcd}
F_{\alpha-\varepsilon} \arrow[rrr] \arrow[dr] & & & F_{\alpha'+\varepsilon} \\
 & G_\alpha \arrow[r] & G_{\alpha'} \arrow[ur] &
\end{tikzcd}
\qquad
\begin{tikzcd}
 & F_{\alpha+\varepsilon} \arrow[r] & F_{\alpha'+\varepsilon} \\
G_\alpha \arrow[ur] \arrow[r] & G_{\alpha'} \arrow[ur] &
\end{tikzcd}$$

$$\begin{tikzcd}
 & F_\alpha \arrow[r] & F_{\alpha'} \arrow[dr] & \\
G_{\alpha-\varepsilon} \arrow[ur] \arrow[rrr] & & & G_{\alpha'+\varepsilon}
\end{tikzcd}
\qquad
\begin{tikzcd}
F_\alpha \arrow[r] \arrow[dr] & F_{\alpha'} \arrow[dr] & \\
 & G_{\alpha+\varepsilon} \arrow[r] & G_{\alpha'+\varepsilon}
\end{tikzcd} \tag{4}$$

These properties extend directly to arbitrary persistence modules, thus providing weak and strong notions of proximity:

Definition 4.3 *Two persistence modules $\mathcal{F}_A$ and $\mathcal{G}_B$ are said to be* weakly ε-interleaved *if the following two conditions are satisfied:*

(i) *there exists $\alpha_0 \in \mathbb{R}$ such that $\alpha_0 + 2\varepsilon\mathbb{Z} \subseteq A$ and $\alpha_0 + \varepsilon + 2\varepsilon\mathbb{Z} \subseteq B$,*

(ii) *there exist two families of homomorphisms $\{\phi_\alpha : F_\alpha \to G_{\alpha+\varepsilon}\}_{\alpha\in\alpha_0+2\varepsilon\mathbb{Z}}$ and $\{\psi_\alpha : G_\alpha \to F_{\alpha+\varepsilon}\}_{\alpha\in\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ such that the diagram of Eq. (3) commutes.*

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

Definition 4.4 *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 $\{\phi_\alpha : F_\alpha \to G_{\alpha+\varepsilon}\}_{\alpha\in\mathbb{R}}$ and $\{\psi_\alpha : G_\alpha \to F_{\alpha+\varepsilon}\}_{\alpha\in\mathbb{R}}$ such that the diagrams of Eq. (4) 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, as shown in Lemma 4.5 below. Nevertheless, $\mathcal{F}_A$ and $\mathcal{G}_B$ cannot be strongly interleaved when $A, B \subsetneq \mathbb{R}$.

Lemma 4.5 *If two persistence modules $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ are weakly ε-interleaved, then they are strongly 3ε-interleaved, and this bound is tight.*

Proof. The approach is the same in spirit as in the proof of Lemma 4.2 (ii). Let α_0, $\{\phi_\alpha : F_\alpha \to G_{\alpha+\varepsilon}\}_{\alpha\in\alpha_0+2\varepsilon\mathbb{Z}}$ and $\{\psi_\alpha : G_\alpha \to F_{\alpha+\varepsilon}\}_{\alpha\in\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ be defined as in Definition 4.3. For any $\alpha \in \mathbb{R}$, let $n \in \mathbb{Z}$ be such that $\alpha_0 + 2(n-1)\varepsilon < \alpha \leq \alpha_0 + 2n\varepsilon$, and define $\phi'_\alpha : F_\alpha \to G_{\alpha+3\varepsilon}$ as follows:

$$\phi'_\alpha = g_{\alpha_0+(2n+1)\varepsilon}^{\alpha+3\varepsilon} \circ \phi_{\alpha_0+2n\varepsilon} \circ f_\alpha^{\alpha_0+2n\varepsilon}.$$

In other words, ϕ'_α corresponds to the path $F_\alpha \to F_{\alpha_0+2n\varepsilon} \to G_{\alpha_0+(2n+1)\varepsilon} \to G_{\alpha+3\varepsilon}$ in the diagram of Eq. (3). Symmetrically, let $\psi'_\alpha : G_\alpha \to F_{\alpha+3\varepsilon}$ correspond to the path $G_\alpha \to G_{\alpha_0+(2m+1)\varepsilon} \to F_{\alpha_0+(2m+2)\varepsilon} \to F_{\alpha+3\varepsilon}$, where $m \in \mathbb{Z}$ is such that $\alpha_0 + (2m-1)\varepsilon < \alpha \leq \alpha_0 + (2m+1)\varepsilon$. The families of homomorphisms $\{\phi'_\alpha\}_{\alpha\in\mathbb{R}}$ and $\{\psi'_\alpha\}_{\alpha\in\mathbb{R}}$ strongly 3ε-interleave $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$, provided that the diagrams of Eq. (4) – with ε replaced by 3ε – commute, which is an easy consequence of the commutativity of the diagram of Eq. (3). Indeed, given $\alpha \leq \alpha' \in \mathbb{R}$ and $n \leq n' \in \mathbb{Z}$ such that $\alpha_0 + (2n-1)\varepsilon < \alpha \leq \alpha_0 + (2n+1)\varepsilon$ and $\alpha_0 + (2n'-1)\varepsilon < \alpha' \leq \alpha_0 + (2n'+1)\varepsilon$, we have:

$$\begin{aligned}
&\psi'_{\alpha'} \circ g_\alpha^{\alpha'} \circ \phi'_{\alpha-3\varepsilon} \\
&= \left(f_{\alpha_0+(2n'+2)\varepsilon}^{\alpha'+3\varepsilon} \circ \psi_{\alpha_0+(2n'+1)\varepsilon} \circ g_{\alpha'}^{\alpha_0+(2n'+1)\varepsilon}\right) \circ g_\alpha^{\alpha'} \circ \left(g_{\alpha_0+(2n-1)\varepsilon}^{\alpha} \circ \phi_{\alpha_0+(2n-2)\varepsilon} \circ f_{\alpha-3\varepsilon}^{\alpha_0+(2n-2)\varepsilon}\right) \\
&= f_{\alpha_0+(2n'+2)\varepsilon}^{\alpha'+3\varepsilon} \circ \psi_{\alpha_0+(2n'+1)\varepsilon} \circ \left(g_{\alpha'}^{\alpha_0+(2n'+1)\varepsilon} \circ g_\alpha^{\alpha'} \circ g_{\alpha_0+(2n-1)\varepsilon}^{\alpha}\right) \circ \phi_{\alpha_0+(2n-2)\varepsilon} \circ f_{\alpha-3\varepsilon}^{\alpha_0+(2n-2)\varepsilon} \\
&= f_{\alpha_0+(2n'+2)\varepsilon}^{\alpha'+3\varepsilon} \circ \psi_{\alpha_0+(2n'+1)\varepsilon} \circ g_{\alpha_0+(2n-1)\varepsilon}^{\alpha_0+(2n'+1)\varepsilon} \circ \phi_{\alpha_0+(2n-2)\varepsilon} \circ f_{\alpha-3\varepsilon}^{\alpha_0+(2n-2)\varepsilon},
\end{aligned}$$

which by commutativity of (3) is equal to $f_{\alpha-3\varepsilon}^{\alpha'+3\varepsilon}$. This proves that the upper-left diagram of Eq. (4) – with ε replaced by 3ε – commutes. The other cases are handled similarly.

As for the tightness of the 3ε bound, it is shown in the following example: given any $\eta \in (0, \varepsilon)$, let $f^\eta, g^\eta : \{0\} \to \mathbb{R}$ be defined by $f^\eta(0) = \eta$ and $g^\eta(0) = 3\varepsilon$. Then, the 0th persistence modules of f^η, g^η are weakly ε-interleaved (take $\alpha_0 = 0$), yet they are only strongly $(3\varepsilon - \eta)$-interleaved. □

4.2 Persistence diagrams of weakly interleaved persistence modules

Theorem 4.6 (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 $\mathrm{d}_B^\infty(D\mathcal{F}_A, D\mathcal{G}_B) \leq 3\varepsilon$, and this bound is tight.*

Proof. Let $\alpha_0 \in \mathbb{R}$ be as in Definition 4.3 (i). Consider the discrete persistence module $\mathcal{H}_{\alpha_0+\varepsilon\mathbb{Z}}$ defined by:

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

By commutativity of the diagram of Definition 4.3 (ii), $\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}$ 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. Therefore, by Corollary 3.6 we have:

$$\mathrm{d}_B^\infty(D\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}},\ D\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}) \leq 2\varepsilon. \tag{5}$$

In addition $\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}$ (resp. $\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$) is a discretization of $\mathcal{F}_A$ (resp. $\mathcal{G}_B$), therefore by Corollary 3.10 we have:

$$\mathrm{d}_B^\infty(D\mathcal{F}_A,\ D\mathcal{F}_{\alpha_0+2\varepsilon\mathbb{Z}}) \leq 2\varepsilon \qquad \mathrm{d}_B^\infty(D\mathcal{G}_B,\ D\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}) \leq 2\varepsilon \tag{6}$$

Combining Eqs. (5) and (6) with the triangle inequality, we obtain:

$$d_B^\infty(D\mathcal{F}_A,\ D\mathcal{G}_B) \leq 6\varepsilon.$$

In order to improve 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 $\mathrm{pix}_{\alpha_0+2\varepsilon\mathbb{Z}}$ 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 $\mathrm{pix}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ 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 $D\mathcal{G}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$). Thus, the map $m = m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1$ is a multi-bijection between $D\mathcal{F}_A$ and $D\mathcal{G}_B$. Consider the possible images of a point $p \in D\mathcal{F}_A$ through this multi-bijection (see Figure 6):

- $m_1(p)$ is at a vertex (u, v) of the grid $\Gamma_{\alpha_0+2\varepsilon\mathbb{Z}}$;
- $m_2^{-1} \circ m_1(p)$ 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 $(u-\varepsilon, v-\varepsilon)$;
- 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 $(u+\varepsilon, v+\varepsilon)$;
- since m_4 is the restriction of $\mathrm{pix}_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$ 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 $\Gamma_{\alpha_0+\varepsilon+2\varepsilon\mathbb{Z}}$.

All in all, $m(p)$ is located in the box $(u-3\varepsilon, u+\varepsilon] \times (v-3\varepsilon, v+\varepsilon]$. Since p is located in $(u-2\varepsilon, u] \times (v-2\varepsilon, v]$, we conclude that $\|p - m(p)\|_\infty < 3\varepsilon$.

To show the tightness of this upper bound, consider the example given at the end of the proof of Lemma 4.5: the two functions $f^\eta, g^\eta : \{0\} \to \mathbb{R}$ have weakly ε-interleaved 0th persistence modules, whose persistence diagrams are respectively $\{(\eta, +\infty)\} \cup \Delta$ and $\{(3\varepsilon, +\infty)\} \cup \Delta$, whose bottleneck distance is $3\varepsilon - \eta$. $\square$

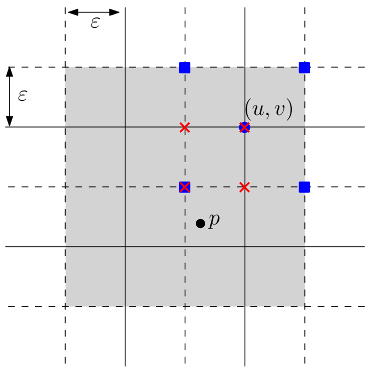

Figure 6: Possible images of p through the multi-bijection $m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1$: $m_1(p)$ is the blue disk with coordinates (u, v); $m_2^{-1} \circ m_1(p)$ lies among the four red crosses; $m_3 \circ m_2^{-1} \circ m_1(p)$ is among the four blue squares, therefore $m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1(p)$ is located in the pink box.

4.3 Persistence diagrams of strongly interleaved persistence modules

We now turn our focus to strongly ε-interleaved tame persistence modules. At a high level, our analysis follows the same scheme as in [12]: first, we bound the Hausdorff distance between the persistence diagrams of such modules; then, we move from Hausdorff to bottleneck distance using an interpolation argument.

4.3.1 Bound on the Hausdorff distance

The key argument is presented in the following lemma from [12], of which we provide a new proof based solely on pixelization arguments:

Lemma 4.7 (Box Lemma) *Let $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ be two tame persistence modules that are strongly ε-interleaved. Given any $\alpha < \beta < \gamma < \delta$, let $\square$ be the box $(\alpha, \beta] \times (\gamma, \delta] \subset \bar{\mathbb{R}}^2$, and $\square_\varepsilon$ 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 $\square_\varepsilon$.*

Proof. If $\beta + \varepsilon > \gamma - \varepsilon$, then $\square_\varepsilon$ 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 ϕ_α, ψ_α introduced in Definition 4.4):

$$\begin{tikzcd}
 & F_\alpha \arrow[r] & F_\beta \arrow[rrr] \arrow[dr] & & & F_\gamma \arrow[r] & F_\delta \arrow[dr] & \\
G_{\alpha-\varepsilon} \arrow[ur] \arrow[rrr] & & & G_{\beta+\varepsilon} \arrow[r] & G_{\gamma-\varepsilon} \arrow[ur] \arrow[rrr] & & & G_{\delta+\varepsilon}
\end{tikzcd}$$

It follows that $\mathcal{F}_A$ and $\mathcal{G}_B$ are two discretizations of the (discrete) mixed persistence module $\mathcal{H}_{A\cup B}$ defined by the path $G_{\alpha-\varepsilon} \to F_\alpha \to F_\beta \to G_{\beta+\varepsilon} \to G_{\gamma-\varepsilon} \to F_\gamma \to F_\delta \to G_{\delta+\varepsilon}$ in the above diagram. Let $m_2 : D\mathcal{H}_{A\cup B} \to D\mathcal{F}_A$ and $m_3 : D\mathcal{H}_{A\cup B} \to D\mathcal{G}_B$ be the associated multi-bijections between their persistence diagrams, as of Theorem 3.5. In addition, $\mathcal{F}_A$ (resp. $\mathcal{G}_B$) is a discretization of $\mathcal{F}_\mathbb{R}$ (resp. $\mathcal{G}_\mathbb{R}$), and let $m_1 : D\mathcal{F}_\mathbb{R} \to D\mathcal{F}_A$ (resp. $m_4 : D\mathcal{G}_\mathbb{R} \to D\mathcal{G}_B$) be the associated multi-bijection between their persistence diagrams, as of Theorem 3.9. We will now track the trajectory of each point $p \in D\mathcal{F}_\mathbb{R} \cap \square$ through the mapping $m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1$ (see Figure 7):

- $m_1(p)$ is the upper-right corner (β, δ) of the box $\square$;
- since α, β and γ, δ are consecutive in $A \cup B$, (β, δ) remains fixed through m_2^{-1};
- the image of (β, δ) through m_3 is the upper-right corner $(\beta + \varepsilon, \delta + \varepsilon)$ of the box $\square_\varepsilon$, therefore $m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1(p)$ lies in $\square_\varepsilon$.

Since the above is true for all points $p \in D\mathcal{F}_\mathbb{R} \cap \square$, and since $m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1 : D\mathcal{F}_\mathbb{R} \to D\mathcal{G}_\mathbb{R}$ is a multi-bijection, we deduce that the total multiplicity of $D\mathcal{F}_\mathbb{R} \cap \square$ is at most the total multiplicity of $D\mathcal{G}_\mathbb{R} \cap \square_\varepsilon$, thereby concluding the proof of the lemma. $\square$

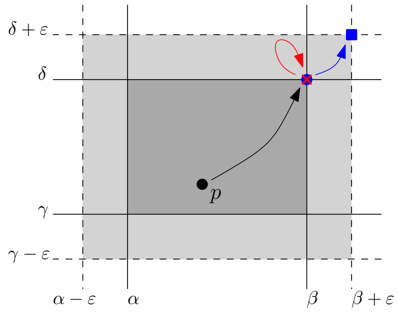

Figure 7: Possible images of $p \in D\mathcal{F}_{\mathbb{R}} \cap \Box$ through the multi-bijection $m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1$: $m_1(p)$ is the blue disk with coordinates (β, δ); $m_2^{-1} \circ m_1(p)$ is the red cross coinciding with the blue disk; $m_3 \circ m_2^{-1} \circ m_1(p)$ is the blue square; $m_4^{-1} \circ m_3 \circ m_2^{-1} \circ m_1(p)$ is located in the light gray box $\Box_\varepsilon$.

The Box Lemma provides the following upper bound on the Hausdorff distance between strongly interleaved persistence modules:

Theorem 4.8 *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 $\mathrm{d}_{\mathcal{H}}^{\infty}(D\mathcal{F}_{\mathbb{R}}, D\mathcal{G}_{\mathbb{R}}) \leq \varepsilon$.*

Proof. Let $p = (\alpha, \beta)$ be a point of $D\mathcal{F}_{\mathbb{R}}$. If $p \in \Delta$, then p belongs to $D\mathcal{G}_{\mathbb{R}}$ as well. Else, for any value $\delta > 0$, the box $(\alpha - \delta, \alpha] \times (\beta - \delta, \beta]$ contains at least one of point of $D\mathcal{F}_{\mathbb{R}}$, namely p. It follows that $(\alpha-\delta-\varepsilon, \alpha+\varepsilon] \times (\beta-\delta-\varepsilon, \beta+\varepsilon]$ contains at least one point of $D\mathcal{G}_{\mathbb{R}}$, by the Box Lemma 4.7. Since this is true for any $\delta > 0$, the l^∞-distance of p to $D\mathcal{G}_{\mathbb{R}}$ must be at most ε. The case of a point $p \in D\mathcal{G}_{\mathbb{R}}$ is symmetric. $\Box$

4.3.2 Bound on the bottleneck distance

The main result of this section is an improvement to the bound of Theorem 4.6 in the case where the persistence modules are strongly ε-interleaved:

Theorem 4.9 (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 $\mathrm{d}_B^{\infty}(D\mathcal{F}_{\mathbb{R}}, D\mathcal{G}_{\mathbb{R}}) \leq \varepsilon$.*

The rest of the section is devoted to the proof of the theorem. As in [12], our proof relies on an interpolation argument. However, differently from [12], we do not interpolate at functional level, but rather at algebraic level directly, since in our context persistence modules are the only available data. In addition to being more general, this strategy is technically interesting since our interpolation between tame persistence modules remains tame throughout, whereas a naive function interpolation yields persistence modules that may no longer be tame.

Lemma 4.10 *Let $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ be two strongly ε-interleaved persistence modules. For all $s \in [0, \varepsilon]$, there exists a persistence module $\widetilde{\mathcal{H}}_{\mathbb{R}}^s$ that is strongly s-interleaved with $\mathcal{F}_{\mathbb{R}}$ and strongly $(\varepsilon - s)$-interleaved with $\mathcal{G}_{\mathbb{R}}$. Furthermore, if $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{G}_{\mathbb{R}}$ are tame, then so is $\widetilde{\mathcal{H}}_{\mathbb{R}}^s$.*

Proof. For clarity, we let $\varepsilon_1 = s$ and $\varepsilon_2 = \varepsilon - s$. We first define the *translated sum* $\mathcal{H}_\mathbb{R}$ by $H_\alpha = F_{\alpha-\varepsilon_1} \oplus G_{\alpha-\varepsilon_2}$.

Denote $f_\alpha^\beta : F_\alpha \to F_\beta$ (resp. $g_\alpha^\beta : G_\alpha \to G_\beta$) the homomorphisms that make $\mathcal{F}$ (resp $\mathcal{G}$) a persistence module. The applications $h_\alpha^\beta : H_\alpha \to H_\beta$ defined for all $x \in F_{\alpha-\varepsilon_1}$ and $y \in F_{\beta-\varepsilon_2}$ 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.

Notice that if $\mathcal{F}$ and $\mathcal{G}$ are tame, $\mathcal{H}$ is tame as well.

Call $\phi_\alpha^{\mathcal{F},\mathcal{G}} : F_\alpha \to G_{\alpha+\varepsilon}$ and $\phi_\alpha^{\mathcal{G},\mathcal{F}} : G_\alpha \to F_{\alpha+\varepsilon}$ the homomorphisms in the interleave of $\mathcal{F}$ and $\mathcal{G}$. We define $\phi_\alpha^{\mathcal{F},\mathcal{H}} : F_\alpha \to H_{\alpha+\varepsilon_1}$ by $\phi_\alpha^{\mathcal{F},\mathcal{H}}(x) = (x,0)$ and $\phi_\alpha^{\mathcal{G},\mathcal{H}} : G_\alpha \to H_{\alpha+\varepsilon_2}$ by $\phi_\alpha^{\mathcal{G},\mathcal{H}}(x) = (0,x)$. We also define $\phi_\alpha^{\mathcal{H},\mathcal{F}} : H_\alpha \to F_{\alpha+\varepsilon_1}$ by $\phi_\alpha^{\mathcal{H},\mathcal{F}}((x,y)) = f_{\alpha-\varepsilon_1}^{\alpha+\varepsilon_1}(x) + \phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{F}}(y)$ and $\phi_\alpha^{\mathcal{H},\mathcal{G}} : H_\alpha \to G_{\alpha+\varepsilon_2}$ by $\phi_\alpha^{\mathcal{H},\mathcal{G}}((x,y)) = \phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{G}}(x) + g_{\alpha-\varepsilon_2}^{\alpha+\varepsilon_2}(y)$.

These homomorphisms satisfy some easy properties. In particular, $\phi_{\alpha+\varepsilon_1}^{\mathcal{H},\mathcal{F}} \circ \phi_\alpha^{\mathcal{F},\mathcal{H}} = f_\alpha^{\alpha+2\varepsilon_1}$, $\phi_{\alpha+\varepsilon_2}^{\mathcal{H},\mathcal{G}} \circ \phi_\alpha^{\mathcal{G},\mathcal{H}} = g_\alpha^{\alpha+2\varepsilon_2}$, $\phi_{\alpha+\varepsilon_1}^{\mathcal{H},\mathcal{G}} \circ \phi_\alpha^{\mathcal{F},\mathcal{H}} = \phi_\alpha^{\mathcal{F},\mathcal{G}}$ 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 $\phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}}$:

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

We then also check it on the image of $\phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{H}}$:

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

Since by definition H_α is generated by the images of $\phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}}$ and $\phi_{\alpha-\varepsilon_2}^{\mathcal{G},\mathcal{H}}$, 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, $g_{\alpha+\varepsilon_2}^{\beta+\varepsilon_2} \circ \phi_\alpha^{\mathcal{H},\mathcal{G}} = \phi_\beta^{\mathcal{H},\mathcal{G}} \circ h_\alpha^\beta$.

$$\begin{tikzcd}
F_{\alpha-\varepsilon_1-2\varepsilon_2} \arrow[rrr] \arrow[dr] & & & F_{\alpha-\varepsilon_1} \arrow[rr] \arrow[dr] & & F_{\alpha+\varepsilon_1} \arrow[dr] & \\
& H_{\alpha-2\varepsilon_2} \arrow[drr] & H_{\alpha-2\varepsilon_1} \arrow[ur] & & H_\alpha \arrow[ur] \arrow[rr] & & H_{\alpha+2\varepsilon_1} \\
& G_{\alpha-2\varepsilon_1-\varepsilon_2} \arrow[ur] \arrow[rr] & & G_{\alpha-\varepsilon_2} \arrow[ur] \arrow[rr] & & G_{\alpha+2\varepsilon_1-\varepsilon_2} \arrow[ur] &
\end{tikzcd}$$

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 $\phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}}$: $\phi_{\alpha+\varepsilon_1}^{\mathcal{F},\mathcal{H}} \circ \phi_\alpha^{\mathcal{H},\mathcal{F}} \circ \phi_{\alpha-\varepsilon_1}^{\mathcal{F},\mathcal{H}} =$

$\phi^{\mathcal{F},\mathcal{H}}_{\alpha+\varepsilon_1} \circ f^{\alpha+\varepsilon_1}_{\alpha-\varepsilon_1} = h^{\alpha+2\varepsilon_1}_{\alpha} \circ \phi^{\mathcal{F},\mathcal{H}}_{\alpha-\varepsilon_1}$, the relation is satisfied. On the image of $\phi^{\mathcal{G},\mathcal{H}}_{\alpha-\varepsilon_2}$: $(\phi^{\mathcal{F},\mathcal{H}}_{\alpha+\varepsilon_1} \circ \phi^{\mathcal{H},\mathcal{F}}_{\alpha} - h^{\alpha+2\varepsilon_1}_{\alpha}) \circ \phi^{\mathcal{G},\mathcal{H}}_{\alpha-\varepsilon_2} = \phi^{\mathcal{F},\mathcal{H}}_{\alpha+\varepsilon_1} \circ \phi^{\mathcal{H},\mathcal{F}}_{\alpha} \circ \phi^{\mathcal{G},\mathcal{H}}_{\alpha-\varepsilon_2} - h^{\alpha+2\varepsilon_1}_{\alpha} \circ \phi^{\mathcal{G},\mathcal{H}}_{\alpha-\varepsilon_2} = \phi^{\mathcal{F},\mathcal{H}}_{\alpha+\varepsilon_1} \circ \phi^{\mathcal{G},\mathcal{F}}_{\alpha-\varepsilon_2} - \phi^{\mathcal{G},\mathcal{H}}_{\alpha+2\varepsilon_1-\varepsilon_2} \circ g^{\alpha+2\varepsilon_1-\varepsilon_2}_{\alpha-\varepsilon_2}$ so $(\phi^{\mathcal{F},\mathcal{H}}_{\alpha+\varepsilon_1} \circ \phi^{\mathcal{H},\mathcal{F}}_{\alpha} - h^{\alpha+2\varepsilon_1}_{\alpha})((x,y)) = (\phi^{\mathcal{G},\mathcal{F}}_{\alpha-\varepsilon_2}(y), -g^{\alpha+2\varepsilon_1-\varepsilon_2}_{\alpha-\varepsilon_2}(y))$ which is not necessarily null.

We thus define

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

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

and let $\widetilde{\mathcal{H}}$ be the persistence module defined by $\widetilde{H}_\alpha = H_\alpha/(\mathbb{F}_\alpha + \mathbb{G}_\alpha)$. If $\mathcal{H}$ is tame, $\widetilde{\mathcal{H}}$ is tame as well. Call $\pi_\alpha : H_\alpha \to \widetilde{H}_\alpha$ the quotient application.

Let y be an element of $\mathbb{F}_\alpha$. Then $y = (f^{\alpha-\varepsilon_1}_{\alpha-\varepsilon_1-2\varepsilon_2}(x), -\phi^{\mathcal{F},\mathcal{G}}_{\alpha-\varepsilon_1-2\varepsilon_2}(x))$ for some $x \in F_{\alpha-\varepsilon_1-2\varepsilon_2}$. $h^\beta_\alpha(y) = (f^{\beta-\varepsilon_1}_{\alpha-\varepsilon_1}(f^{\alpha-\varepsilon_1}_{\alpha-\varepsilon_1-2\varepsilon_2}(x)), g^{\beta-\varepsilon_2}_{\alpha-\varepsilon_2}(-\phi^{\mathcal{F},\mathcal{G}}_{\alpha-\varepsilon_1-2\varepsilon_2}(x)))$. $h^\beta_\alpha(y) = (f^{\beta-\varepsilon_1}_{\beta-\varepsilon_1-2\varepsilon_2}(z), -\phi^{\mathcal{F},\mathcal{G}}_{\beta-\varepsilon_1-2\varepsilon_2}(z))$ where $z = f^{\beta-\varepsilon_1-2\varepsilon_2}_{\alpha-\varepsilon_1-2\varepsilon_2}(x)$, *i.e.* $h^\beta_\alpha(y) \in \mathbb{F}_\beta$.

Thus $h^\beta_\alpha(\mathbb{F}_\alpha) \subseteq \mathbb{F}_\beta$, $h^\beta_\alpha(\mathbb{G}_\alpha) \subseteq \mathbb{G}_\beta$ (same proof) and h^β_α induces a homomorphism $\widetilde{h}^\beta_\alpha : \widetilde{H}_\alpha \to \widetilde{H}_\beta$ such that $\widetilde{h}^\beta_\alpha \circ \pi_\alpha = \pi_\beta \circ h^\beta_\alpha$ and the $\widetilde{h}^\beta_\alpha$ satisfy the persistence module property: $\widetilde{h}^\gamma_\beta \circ \widetilde{h}^\beta_\alpha \circ \pi_\alpha = \widetilde{h}^\gamma_\beta \circ \pi_\beta \circ h^\beta_\alpha = \pi_\gamma \circ h^\gamma_\beta \circ h^\beta_\alpha = \pi_\gamma \circ h^\gamma_\alpha = \widetilde{h}^\gamma_\alpha \circ \pi_\alpha$ so $\widetilde{h}^\gamma_\beta \circ \widetilde{h}^\beta_\alpha = \widetilde{h}^\gamma_\alpha$.

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

A key property[4] of $\mathbb{F}_\alpha$ and $\mathbb{G}_\alpha$ is that they are included in the kernels of $\phi^{\mathcal{H},\mathcal{F}}_\alpha$ and $\phi^{\mathcal{H},\mathcal{G}}_\alpha$. Indeed, $\phi^{\mathcal{H},\mathcal{F}}_\alpha((f^{\alpha-\varepsilon_1}_{\alpha-\varepsilon_1-2\varepsilon_2}(x), -\phi^{\mathcal{F},\mathcal{G}}_{\alpha-\varepsilon_1-2\varepsilon_2}(x))) = f^{\alpha+\varepsilon_1}_{\alpha-\varepsilon_1} \circ f^{\alpha-\varepsilon_1}_{\alpha-\varepsilon_1-2\varepsilon_2}(x) - \phi^{\mathcal{G},\mathcal{F}}_{\alpha-\varepsilon_2} \circ \phi^{\mathcal{F},\mathcal{G}}_{\alpha-\varepsilon_1-2\varepsilon_2}(x) = (f^{\alpha+\varepsilon_1}_{\alpha-\varepsilon_1-2\varepsilon_2} - \phi^{\mathcal{G},\mathcal{F}}_{\alpha-\varepsilon_2} \circ \phi^{\mathcal{F},\mathcal{G}}_{\alpha-\varepsilon_1-2\varepsilon_2})(x) = 0$ on one hand, $\phi^{\mathcal{H},\mathcal{F}}_\alpha((\phi^{\mathcal{G},\mathcal{F}}_{\alpha-2\varepsilon_1-\varepsilon_2}(x), -g^{\alpha-\varepsilon_2}_{\alpha-2\varepsilon_1-\varepsilon_2}(x))) = f^{\alpha+\varepsilon_1}_{\alpha-\varepsilon_1} \circ \phi^{\mathcal{G},\mathcal{F}}_{\alpha-2\varepsilon_1-\varepsilon_2}(x) - \phi^{\mathcal{G},\mathcal{F}}_{\alpha-\varepsilon_2} \circ g^{\alpha-\varepsilon_2}_{\alpha-2\varepsilon_1-\varepsilon_2}(x) = 0$ on the other hand, and the proof is identical for $\phi^{\mathcal{H},\mathcal{G}}_\alpha$. This implies that $\phi^{\mathcal{H},\mathcal{F}}_\alpha$ and $\phi^{\mathcal{H},\mathcal{G}}_\alpha$ induce homomorphisms $\widetilde{\phi}^{\mathcal{H},\mathcal{F}}_\alpha : \widetilde{H}_\alpha \to F_{\alpha+\varepsilon_1}$ and $\widetilde{\phi}^{\mathcal{H},\mathcal{G}}_\alpha : \widetilde{H}_\alpha \to G_{\alpha+\varepsilon_2}$ such that $\phi^{\mathcal{H},\mathcal{F}}_\alpha = \widetilde{\phi}^{\mathcal{H},\mathcal{F}}_\alpha \circ \pi_\alpha$ and $\phi^{\mathcal{H},\mathcal{G}}_\alpha = \widetilde{\phi}^{\mathcal{H},\mathcal{G}}_\alpha \circ \pi_\alpha$.

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

The family $(\widetilde{\mathcal{H}}^s_{\mathbb{R}})_{s\in[0,\varepsilon]}$ of persistence modules introduced in Lemma 4.10 *interpolates* between the two initial persistence modules in the following sense:

Lemma 4.11 *Let $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ be two strongly ε-interleaved persistence modules, and let $(\widetilde{\mathcal{H}}^s_{\mathbb{R}})_{s\in[0,\varepsilon]}$ be the family of persistence modules introduced in Lemma 4.10. We have $\widetilde{\mathcal{H}}^0_\mathbb{R} \simeq \mathcal{F}_\mathbb{R}$, $\widetilde{\mathcal{H}}^\varepsilon_\mathbb{R} \simeq \mathcal{G}_\mathbb{R}$, and for all $s, s' \in [0,\varepsilon]$, $\widetilde{\mathcal{H}}^s_\mathbb{R}$ and $\widetilde{\mathcal{H}}^{s'}_\mathbb{R}$ are strongly $|s-s'|$-interleaved.*

[4]Indeed if we define $\widetilde{H}_\alpha$ as $H_\alpha/(Ker(\phi^{\mathcal{H},\mathcal{F}}_\alpha) \cap Ker(\phi^{\mathcal{H},\mathcal{F}}_\alpha))$ the construction works as well.

Proof. To avoid confusion, for all $\varepsilon_1+\varepsilon_2=\varepsilon$ we denote by $\widetilde{\mathcal{H}}=\mathcal{F}\otimes_{\varepsilon_1,\varepsilon_2}\mathcal{G}$ 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 $\psi_\alpha:K_\alpha\to J_\alpha$ the function defined by $\psi_\alpha(((x,y),z))=(x,g_{\alpha-2\varepsilon_2-\varepsilon_3}^{\alpha-\varepsilon_3}(y)+z)$.

Define

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

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

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

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

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

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

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

$\mathbb{H}_\alpha^J$ 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 $\widetilde{J}_\alpha=J_\alpha/(\mathbb{F}_{\alpha-\varepsilon_2}^H\times 0+\mathbb{G}_{\alpha-\varepsilon_2}^H\times 0+\mathbb{F}_\alpha^J+\mathbb{G}_\alpha^J)$.

$\mathbb{G}_\alpha^J$ 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 $\widetilde{\psi}_\alpha$ between $\widetilde{J}_\alpha$ and $\widetilde{K}_\alpha$. It is easy to check that the persistence homomorphisms $J_\alpha\to J_\beta$ and $K_\alpha\to K_\beta$ 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 $\widetilde{\mathcal{H}}_\mathbb{R}^0$ are strongly 0-interleaved, which means that $\widetilde{\phi}^{\mathcal{F},\mathcal{H}}$ defines an isomorphism between $\mathcal{F}_\mathbb{R}$ and $\widetilde{\mathcal{H}}_\mathbb{R}^0$ with inverse $\widetilde{\phi}^{\mathcal{H},\mathcal{F}}$. The same holds for $\mathcal{G}_\mathbb{R}$ and $\widetilde{\mathcal{H}}_\mathbb{R}^\varepsilon$.

$$\begin{tikzcd}
 & F_{\alpha-\varepsilon_1-\varepsilon_2-2\varepsilon_3} \arrow[rrr] \arrow[dr] & & & F_{\alpha-\varepsilon_1-\varepsilon_2} \arrow[dr] & & \\
 & & H_{\alpha-\varepsilon_2-2\varepsilon_3} \arrow[rrr] \arrow[ddrr] & & & H_{\alpha-\varepsilon_2} \arrow[dr] & \\
 & & & & & & J_\alpha \simeq K_\alpha \\
G_{\alpha-2\varepsilon_2-3\varepsilon_3} \arrow[uurr] \arrow[rrr, bend left] & G_{\alpha-2\varepsilon_1-2\varepsilon_2-\varepsilon_3} \arrow[uuurrr] \arrow[rr] & & G_{\alpha-2\varepsilon_2-\varepsilon_3} \arrow[uurr] \arrow[r] & G_{\alpha-\varepsilon_3} \arrow[urr] & &
\end{tikzcd}$$

□

We now have the necessary ingredients at hand to apply the interpolation argument of [12]. However, we need to handle some additional technical details that stem from the fact that our persistence diagrams are allowed to contain infinitely many points off the diagonal Δ, some of which may diverge or accumulate towards Δ.

Let $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ be two tame persistence modules that are strongly ε-interleaved, and let $(\widetilde{\mathcal{H}}^s_\mathbb{R})_{s\in[0,\varepsilon]}$ be the interpolating family of persistence modules defined in Lemma 4.10. Choose two arbitrary numbers $\alpha < \beta$, as well as an arbitrarily small value $\eta > 0$. According to Corollary 3.4, any discretization $\widetilde{\mathcal{H}}^s_A$ of $\widetilde{\mathcal{H}}^s_\mathbb{R}$ has the property that the support of $D\widetilde{\mathcal{H}}^s_A$ contains finitely many points in the vertical band $[\alpha, \beta] \times \bar{\mathbb{R}}$ minus the diagonal Δ. Then, Theorem 3.9 (applied with a sufficiently small discretization step) guarantees that the support of $D\widetilde{\mathcal{H}}^s_\mathbb{R}$ contains finitely many points in the area $B^{\alpha,\beta}_\eta = ([\alpha, \beta] \times \bar{\mathbb{R}}) \cap \Delta^\eta_+ = \{(u, v) \in \bar{\mathbb{R}}^2 \mid \alpha \leq u \leq \beta \text{ and } v \geq u + 2\eta\}$. Let

$$\delta^{\alpha,\beta}_\eta(s) = \frac{1}{4}\min\{\|p - q\|_\infty \mid p, q \in B^{\alpha,\beta}_\eta \text{ and } p \neq q\} > 0. \tag{7}$$

For any $s, s' \in [0, \varepsilon]$, $\widetilde{\mathcal{H}}^{s'}_\mathbb{R}$ is said to be *η-close* to $\widetilde{\mathcal{H}}^s$ if $|s - s'| < \delta^{\alpha,\beta}_\eta(s)$. 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 $B^{\alpha,\beta}_\eta$:

Lemma 4.12 (Easy Bijection Lemma) *Let $s, s' \in [0, \varepsilon]$ be such that $\widetilde{\mathcal{H}}^{s'}_\mathbb{R}$ is η-close to $\widetilde{\mathcal{H}}^s_\mathbb{R}$. Then, there exists a multi-bijection m between $D\widetilde{\mathcal{H}}^s_\mathbb{R}$ and $D\widetilde{\mathcal{H}}^{s'}_\mathbb{R}$ such that:*

(i) $\|p - m(p)\|_\infty \leq |s' - s|$ *for all* $p \in D\widetilde{\mathcal{H}}^s_\mathbb{R} \cap B^{\alpha+4|s-s'|,\, \beta-4|s-s'|}_{\eta+4|s-s'|}$,

(ii) $\|p - m(p)\|_\infty \leq 3|s' - s|$ *for any other point* p *of* $D\widetilde{\mathcal{H}}^s_\mathbb{R}$.

Proof. By Lemmas 4.10 and 4.11, $\widetilde{\mathcal{H}}^s_\mathbb{R}$ and $\widetilde{\mathcal{H}}^{s'}_\mathbb{R}$ are tame and strongly $|s - s'|$-interleaved. Therefore, the Weak Stability Theorem 4.6 implies that there exists a multi-bijection $m : D\widetilde{\mathcal{H}}^s_\mathbb{R} \to D\widetilde{\mathcal{H}}^{s'}_\mathbb{R}$ that moves the points of $D\widetilde{\mathcal{H}}^s_\mathbb{R}$ by at most $3|s - s'|$ in the l^∞-distance, thus proving (ii).

Consider now an arbitrary point $p \in D\widetilde{\mathcal{H}}^s_\mathbb{R} \cap B^{\alpha+4|s-s'|,\, \beta-4|s-s'|}_{\eta+4|s-s'|}$. By Eq. (7), for all $q \in D\widetilde{\mathcal{H}}^s_\mathbb{R} \cap B^{\alpha,\beta}_\eta$ we have $\|p - q\|_\infty \geq 4\delta^{\alpha,\beta}_\eta(s)$, which is greater than $4|s - s'|$ since $\widetilde{\mathcal{H}}^{s'}_\mathbb{R}$ is η-close to $\widetilde{\mathcal{H}}^s_\mathbb{R}$. 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 $D\widetilde{\mathcal{H}}^s_{\mathbb{R}} \setminus \{p\}$ 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 $D\widetilde{\mathcal{H}}^s_{\mathbb{R}} \setminus \{p\}$ is greater than $|s-s'|$. Now, by Theorem 4.8, the l^∞-distance of $m(p)$ to $D\widetilde{\mathcal{H}}^s_{\mathbb{R}}$ is at most $|s-s'|$, which implies that $\|p - m(p)\|_\infty \leq |s-s'|$. $\square$

Consider now an arbitrary positive function $r : [0, \varepsilon] \to \mathbb{R}$ that is bounded from above by $\delta^{\alpha,\beta}_\eta$, that is: $\forall s \in [0,\varepsilon],\ 0 < r(s) \leq \delta^{\alpha,\beta}_\eta(s)$. 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 $\widetilde{\mathcal{H}}^{s_i}_{\mathbb{R}}$ is η-close to $\widetilde{\mathcal{H}}^{s_{i+1}}_{\mathbb{R}}$, or the other way around. It follows that there exists a multi-bijection $m_i : D\widetilde{\mathcal{H}}^{s_i}_{\mathbb{R}} \to D\widetilde{\mathcal{H}}^{s_{i+1}}_{\mathbb{R}}$ satisfying assertions (i) and (ii) of the Easy Bijection Lemma 4.12 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 $m_0 : D\mathcal{F}_{\mathbb{R}} \to D\widetilde{\mathcal{H}}^{s_1}_{\mathbb{R}}$ and $m_k : D\widetilde{\mathcal{H}}^{s_k}_{\mathbb{R}} \to D\mathcal{G}_{\mathbb{R}}$ satisfying conditions (i) and (ii) above. Let

$$m = m_k \circ m_{k-1} \circ \cdots \circ m_1 \circ m_0 \tag{8}$$

be the induced multi-bijection between $D\mathcal{F}_{\mathbb{R}}$ and $D\mathcal{G}_{\mathbb{R}}$. Combining assertion (ii) of Lemma 4.12 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 4.13 *For any point $p \in (D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}) \cap B^{\alpha+8\varepsilon,\ \beta-8\varepsilon}_{\eta+7\sup r+\varepsilon}$, p and its pair q satisfy $\|p - q\|_\infty \leq \varepsilon$.*

Proof. Let $p \in (D\mathcal{F}_{\mathbb{R}} \cup D\mathcal{G}_{\mathbb{R}}) \cap B^{\alpha+8\varepsilon,\ \beta-8\varepsilon}_{\eta+7\sup r+\varepsilon}$. 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 $p = p_0, p_1, \cdots, p_k, p_{k+1} = m(p)$ be the images of p through the sequence of multi-bijections m_i introduced in Eq. (8). We will show that $\|p_i - p_{i+1}\|_\infty \leq |s_i - s_{i+1}|$ for all $i = 0, \cdots, k$, which by the triangle inequality implies that $\|p - m(p)\| \leq \varepsilon$. 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 $B^{\alpha+7\varepsilon,\ \beta-7\varepsilon}_{\eta+7\sup r}$, which is included in $B^{\alpha+7|s_l-s_{l+1}|,\ \beta-7|s_l-s_{l+1}|}_{\eta+7|s_l-s_{l+1}|}$ since, as we saw above, $|s_l - s_{l+1}|$ is bounded from above by $\min\{\varepsilon,\ \max\{r(s_l),\ r(s_{l+1})\}\} \leq \min\{\varepsilon,\ \sup r\}$. Now, by Lemma 4.12 (ii), we have $\|p_l - p_{l+1}\|_\infty \leq 3|s_l - s_{l+1}|$, which implies that p_{l+1} belongs to $B^{\alpha+4|s_l-s_{l+1}|,\ \beta-4|s_l-s_{l+1}|}_{\eta+4|s_l-s_{l+1}|}$. As a result, whether $\widetilde{\mathcal{H}}^{s_l}_{\mathbb{R}}$ be η-close to $\widetilde{\mathcal{H}}^{s_{l+1}}_{\mathbb{R}}$ or the other way around, Lemma 4.12 (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 $m : D\mathcal{F}_{\mathbb{R}} \to D\mathcal{G}_{\mathbb{R}}$ 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 4.13, 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 4.14 *For any $\nu > 0$, there exists a multi-bijection $D\mathcal{F}_{\mathbb{R}} \to D\mathcal{G}_{\mathbb{R}}$ 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 $m_n : D\mathcal{F}_{\mathbb{R}} \to D\mathcal{G}_{\mathbb{R}}$ the multi-bijection given by Corollary 4.14 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 4.9.

5 Conclusion

We have shown that the notion of persistence diagram can be extended, and its stability proven, beyond the framework of [12]. In particular, working at algebraic level directly, we have provided a mean of comparing the persistence diagrams of functions defined over different spaces, thus giving a partial answer to an open question of [13]. In order 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 over sampled Riemannian manifolds [9].

An important question is whether an equivalent of the structure theorem of [22] exists under our weaker notion of tameness, introduced in Definition 2.5. Is it true that two persistence modules with identical persistence diagrams are isomorphic, regardless of the fact that their diagrams may have infinitely (yet countably) many points off the diagonal?

Note that our main notion of proximity between persistence modules, introduced in Definition 4.4, satisfies the axioms of a distance in the following sense: it is symmetric, it satisfies the triangle inequality, and two persistence modules $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ are strongly 0-interleaved if and only if they are isomorphic. This latter condition, combined with the correspondence and structure theorems of [22], implies that $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ have identical persistence diagrams if and only if they are strongly 0-interleaved. It would be interesting to see whether an approximate version of this result exists, stating that $\mathcal{F}_\mathbb{R}$ and $\mathcal{G}_\mathbb{R}$ have ε-close persistence diagrams in the bottleneck distance if and only if they are strongly ε-interleaved. Such a result, if true, would draw a clear parallel between proximity of persistence modules and proximity of their persistence diagrams. Our main result (Theorem 4.9) proves one direction, but the other direction remains open.

Among the various possible extensions of the current persistence theory, multidimensional persistence is certainly one of the most actively explored. A natural question to ask is whether stability results exist for certain generalizations of the persistence diagrams to higher dimensions. To the best of our knowledge, results of this nature have been proposed in the following somewhat restrictive contexts:

- multidimensional size theory [3], where only 0-dimensional homology is concerned,
- multidimensional persistence on a class of continuous functions satisfying some stronger tameness condition that has the property of being closed with respect to the max operator [4].

A major difficulty stems from the fact that there is no complete discrete invariant for multidimensional persistence modules [5]. The authors of [3, 4] chose to use the discrete *rank invariant* proposed by [5], which is complete in 1-d but not in higher dimensions. Characterizing what should be a good choice of discrete invariant in higher dimensions remains open.

Acknowledgements

This work was carried out within the Associate Team *TGDA: Topological and Geometric Data Analysis* involving the Geometrica group at INRIA and the Geometric Computing group at Stanford University. It was supported by ANR project 05-JCJC-0246 (*GeoTopAl*), DARPA grant HR0011-05-1-0007, as well as NSF grants ITR-0205671, FRG-0354543, and CCF-0634803.

The authors wish to thank Gunnar Carlsson, Vin de Silva and Primoz Skraba for helpful discussions. They also thank Claudia Landi for providing relevant references in size theory.

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. *Interational 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, L. J. Guibas, and S. Y. Oudot. The stability of persistence diagrams revisited. Research Report 6568 (v.1), INRIA, July 2008. `http://hal.inria.fr/inria-00292566_v1/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, J. Harer, and A. Zomorodian. Hierarchical Morse complexes for piecewise linear 2-manifolds. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 70–79, 2001.

[18] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.

[19] 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.

[20] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2001.

[21] D. Morozov. *Homological Illusions of Persistence and Stability.* Ph.d. dissertation, Department of Computer Science, Duke University, August 2008.

[22] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.

Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 Orsay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399