



CGL

Computational Geometric Learning

The structure and stability of persistence modules

Frédéric Chazal

Vin de Silva

Marc Glisse

Steve Oudot

CGL Technical Report No.: 35

Part of deliverable: WP-WP1/
Site: INRIA Saclay
Month: 24

Project co-funded by the European Commission within FP7 (2010–2013)
under contract nr. IST-25582

THE STRUCTURE AND STABILITY OF PERSISTENCE MODULES

FRÉDÉRIC CHAZAL, VIN DE SILVA, MARC GLISSE, AND STEVE OUDOT

ABSTRACT. We give a self-contained treatment of the theory of persistence modules indexed over the real line. We give new proofs of the standard results. Persistence diagrams are constructed using measure theory. Linear algebra lemmas are simplified using a new notation for calculations on quiver representations. We show that the stringent finiteness conditions required by traditional methods are not necessary to prove the existence and stability of the persistence diagram. We introduce weaker hypotheses for taming persistence modules, which are met in practice and are strong enough for the theory still to work. The constructions and proofs enabled by our framework are, we claim, cleaner and simpler.

CONTENTS

Introduction	3
Multisets	5
1. Persistence Modules	6
1.1. Persistence modules over a real parameter	6
1.2. Different index sets	8
1.3. Module categories	9
1.4. Interval modules	9
1.5. Interval decomposition	11
1.6. The persistence diagram of a decomposable module	14
1.7. Quiver calculations	15
2. Rectangle Measures	17
2.1. The persistence measure	18
2.2. The persistence measure (continued)	20
2.3. Abstract r-measures	22

2.4.	Equivalence of measures and diagrams	23
2.5.	Non-finite measures	28
2.6.	The diagram at infinity	29
2.7.	Diagrams of persistence modules.	31
2.8.	Tameness conditions	33
2.9.	Finite approximations	38
3.	Interleaving	39
3.1.	Shifted homomorphisms	39
3.2.	Interleaving	40
3.3.	Interleaving (continued)	41
3.4.	The interpolation lemma	43
4.	The Isometry Theorem	48
4.1.	The interleaving distance	48
4.2.	The bottleneck distance	50
4.3.	The bottleneck distance (continued)	55
4.4.	The isometry theorem	57
4.5.	The converse stability theorem	57
4.6.	The stability theorem	63
4.7.	The measure stability theorem	65
4.8.	The measure stability theorem (continued)	69
5.	Examples	72
5.1.	Partial interleavings	72
5.2.	Extended persistence	75
	Acknowledgements	77
	References	77

INTRODUCTION

We intend this paper to be a self-contained treatment of the theory of persistence modules over the real line. We give the best proofs we know of the most important results. Each theorem is located at the appropriate level of abstraction (we believe).

Many authors have studied persistence modules in recent years, and many of the theorems presented here are not original in themselves. The originality lies in the methods we use, which give easier proofs and sharper results. Our main innovations are these:

- We use measure theory to construct persistence diagrams. The existence of a diagram is equivalent to the existence of a certain kind of measure on rectangles in the plane.
- We introduce ‘decorated’ real numbers to remove ambiguity about interval endpoints. Decorations are also what make the measure theory work.
- We define several kinds of ‘tameness’ for a persistence module. These occur naturally in practice. The *most* restrictive of these, finite type, is what is normally seen in the literature. We show how to work effectively with the less restrictive hypotheses.
- We introduce a special notation for calculations on quiver representations. This considerably simplifies the linear algebra (for instance, in proving the ‘box lemma’).

Our goal in introducing these ideas is to enable other authors to define persistence diagrams cleanly, in a wide variety of situations, without imposing unnecessary restrictions (such as assuming a function to be Morse). For instance, it will be seen in forthcoming work that the approach here can be used to define the levelset zigzag persistence of [5] quite broadly.

This paper owes much to [7] (and its published journal version [6]), which established the existence and stability of persistence diagrams for modules whose persistence maps are of finite rank. In the present work we call these modules ‘q-tame’.

Traditionally, continuous persistence diagrams have been treated in one of two ways. Most commonly, one makes the aggressive assumption that the situation being studied has only finitely many ‘critical values’. Alternatively, as carried out in [7] for q-tame modules, the diagram is constructed using a careful limiting process through ever-finer discretisations of the parameter. The former strategy may be appropriate when working with real-world data, where every persistence module really is finite in every way, but it excludes very common theoretical situations. The latter strategy was devised to overcome these restrictions, but unfortunately the limiting arguments turn out to be quite complicated. Our new approach gives the best of both worlds; we are able to work with broader classes of persistence modules, and we can reason about their diagrams in a clean way using arguments of a finite nature.

The other debt to [7] is the recasting of stability as a statement about interleaved persistence modules, the so-called ‘algebraic stability theorem’. In this paper we re-recast the result as a statement about 1-parameter families of measures. This allows us to prove stability results for even quite badly behaved persistence modules.

Overview. The paper is organised as follows.

In section 1, we set up the basic properties of persistence modules. These can be defined over any partially ordered set; we are primarily interested in persistence modules over the real line. In the best case a persistence module can be expressed as a direct sum of ‘interval modules’, which can be thought of as the atomic building blocks of the theory. Not all persistence modules decompose in this way, so we spend much of this paper developing techniques that work without this assumption. These techniques depend on a thorough understanding of finitely-indexed persistence modules known as ‘ A_n -quiver representations’ [15, 12]. We introduce a special notation for performing calculations on these quiver representations. This ‘quiver calculus’ is used throughout the paper.¹

Section 2 addresses the question of how to define the diagram of a persistence module. This is easy for modules which decompose into intervals. To handle the general case, we establish an equivalence between diagrams and a certain kind of measure defined on rectangles in the plane. This means that whenever a persistence diagram is sought, it is enough to construct the corresponding persistence measure; and theorems about a diagram can be replaced by simpler-to-prove theorems about its measure. The diagram exists wherever the measure takes finite values. This leads to several different notions of ‘tameness’. There are large classes of examples of naturally occurring persistence modules which are tame enough for their diagrams to be defined everywhere or almost everywhere.

In order to make the measure theory work, we use real numbers that are ‘decorated’ with a superscript $+$ or $-$. We use decorations also to indicate whether a real interval is open, closed or half-open. For persistence modules which decompose into interval modules, these two uses match up perfectly.

There is a ‘snapping principle’ by which our abstractly defined diagrams are seen to agree with the diagrams that are produced by the standard algorithms [13, 20], for example when working with finite simplicial complexes derived from real data.

In section 3, we study interleavings. An interleaving is an approximate isomorphism between two persistence modules. They occur naturally in applications when the input data are known only up to some bounded error. After presenting the basic properties, we give a clean proof of the technical lemma (from [7]) that two interleaved modules can be interpolated by a 1-parameter family.

Section 4 is devoted to the isometry theorem, which asserts that the interleaving distance between two persistence modules is equal to the bottleneck distance between their persistence diagrams. The two inequalities that comprise this result are treated separately. One direction is the celebrated stability theorem of [10]. The more recent converse inequality appears in [16]. We formulate the stability theorem as a statement about measures and their diagrams. The proof of this more abstract result closely follows the original proof in [10].

Our version of the isometry theorem supposes that the persistence modules are ‘q-tame’. We consider this to be the natural realm of the theorem. We also prove a more general version

¹Readers who wish to adopt our notation are invited to contact us for the L^AT_EX macros.

of the stability theorem which allows us to compare diagrams of persistence modules with no assumptions on their tameness: wherever the two diagrams are defined, they must be close to each other.

Given the length of this paper, the reader may wonder if the framework developed here is truly simpler than existing approaches. In section 5, we give examples of how to use the results and ideas in this paper.

Related Work. The early history of persistence is concerned with the quantity

$$r_s^t = \text{rank}(H(X_s) \rightarrow H(X_t))$$

for an object X represented at two different scales s, t , and where H is homology. This appeared in the early 1990s in the work of Frosini [14], with different notation and under the name ‘size function’. Independently, a few years later, Robins [17] introduced the term ‘persistent Betti numbers’ for quantities of the form $r_\epsilon^{\epsilon+\rho}$, and noted their stability with respect to Hausdorff distance.

The modern theory of persistence is built on three central pillars.

- The persistence diagram, and an algorithm for computing it, were introduced by Edelsbrunner, Letscher and Zomorodian [13].
- Zomorodian and Carlsson [20] defined persistence modules, indexed by the natural numbers and viewed as graded modules over the polynomial ring $\mathbf{k}[t]$.
- Cohen-Steiner, Edelsbrunner and Harer [10] proved the stability theorem.

All three papers make strong finiteness assumptions about the starting data. In [7], the results are generalised to persistence modules parametrised over the real line, under the assumption that $r_s^t < \infty$ for $s < t$. The present paper extends that work, with many new concepts and proofs.

We draw attention to two recent papers which share our goal of understanding continuous-parameter persistence modules. Lesnick [16] gives an extensive algebraic treatment of modules over one or more real parameters. The converse stability inequality (and hence the isometry theorem) appears for the first time there. Bubenik and Scott [3] develop the category-theoretical view of persistence modules.

Both [16] and [3] appeared during the writing of this paper. There is a fair amount of overlap between the three papers, most of it reached independently, but not quite all: from Lesnick [16] we learned of the results of Webb [19], which resolved a sticking-point for us.

MULTISSETS

Persistence diagrams are multisets rather than sets. For our purposes, a multiset is a pair $A = (S, m)$ where S is a set and

$$m : S \rightarrow \{1, 2, 3, \dots\} \cup \{\infty\}$$

is the multiplicity function, which tells us how many times each element of S occurs in \mathbf{A} .

Here are our conventions regarding multisets.

- The cardinality of $\mathbf{A} = (S, m)$ is defined to be

$$\text{card } \mathbf{A} = \sum_{s \in S} m(s)$$

which takes values in $\{0, 1, 2, \dots\} \cup \{\infty\}$. We do not distinguish between different infinite cardinals.

- We never form the intersection of two multisets, but we will sometimes restrict a multiset \mathbf{A} to a set B :

$$\mathbf{A}|_B = (S \cap B, m|_{S \cap B})$$

We may write this as $\mathbf{A} \cap B$ when $\mathbf{A}|_B$ is typographically inconvenient.

- A pair (B, m) where

$$m : B \rightarrow \{0, 1, 2, \dots\} \cup \{\infty\}$$

is implicitly regarded as defining a multiset $\mathbf{A} = (S, m|_S)$ where $S = B - m^{-1}(0)$ is the support of m .

- If $\mathbf{A} = (S, m)$ is a multiset and $f : S \rightarrow B$ where B is a set, then the notation

$$\{f(a) \mid a \in \mathbf{A}\}$$

is interpreted as the multiset in B with multiplicity function

$$m'(b) = \sum_{f^{-1}(b)} m(s)$$

Except in definitions like these, we seldom refer explicitly to S .

1. PERSISTENCE MODULES

All vector spaces are taken to be over an arbitrary field \mathbf{k} , fixed throughout the paper.

1.1. Persistence modules over a real parameter. A persistence module \mathbb{V} over the real numbers \mathbf{R} is defined to be an indexed family of vector spaces

$$(V_t \mid t \in \mathbf{R}),$$

and a doubly-indexed family of linear maps

$$(v_s^t : V_s \rightarrow V_t \mid s \leq t)$$

which satisfy the composition law

$$v_s^t \circ v_r^s = v_r^t$$

whenever $r \leq s \leq t$, and where v_t^t is the identity map on V_t .

Remark. Equivalently, a persistence module is a functor from the real line (viewed as a category with a unique morphism $s \rightarrow t$ whenever $s \leq t$) to the category of vector spaces. The uniqueness of the morphism $s \rightarrow t$ corresponds to the fact that all possible compositions

$$v_{s_{n-1}}^t \circ v_{s_{n-2}}^{s_{n-1}} \circ \cdots \circ v_{s_1}^{s_2} \circ v_s^{s_1}$$

from V_s to V_t are equal to each other, and in particular to v_s^t .

Here is the standard class of examples from applied topology. Let X be a topological space and let $f : X \rightarrow \mathbf{R}$ be a function (not necessarily continuous). Consider the sublevel sets:

$$X^t = (X, f)^t = \{x \in X \mid f(x) \leq t\}$$

The inclusion maps

$$(i_s^t : X^s \rightarrow X^t \mid s \leq t)$$

trivially satisfy the composition law

$$i_s^t \circ i_r^s = i_r^t$$

whenever $r \leq s \leq t$, and i_t^t is the identity on X^t . Collectively this information is called the **sublevelset filtration** of (X, f) and may be denoted \mathbb{X}_{sub} or $\mathbb{X}_{\text{sub}}^f$.

We obtain a persistence module by applying any functor from topological spaces to vector spaces. For example, let $H = H_p(-; \mathbf{k})$ be the functor ‘ p -dimensional singular homology with coefficients in \mathbf{k} ’. We define a persistence module \mathbb{V} by setting

$$V_t = H(X^t),$$

and

$$v_s^t = H(i_s^t) : H(X^s) \rightarrow H(X^t)$$

(the maps on homology induced by the inclusion maps). More simply:

$$\mathbb{V} = H(\mathbb{X}_{\text{sub}})$$

In the applied topology literature, there are many examples (X, f) whose persistent homology is of interest. Very often X is a finite simplicial complex and each X^t is a subcomplex. It follows that the vector spaces $H(X^t)$ are finite-dimensional; and as t increases there are finitely many ‘critical values’ at which the complex changes, growing by one or more new cells. Suppose these critical values are

$$a_1 < a_2 < \cdots < a_n.$$

Then all the information in the persistence module is contained in the finite diagram

$$H(X^{a_1}) \rightarrow H(X^{a_2}) \rightarrow \cdots \rightarrow H(X^{a_n})$$

of finite-dimensional vector spaces and linear maps. In this situation,

- the isomorphism type of $H(\mathbb{X}_{\text{sub}})$ admits a compact description [13, 20];
- there is a fast algorithm for computing this description [13, 20];
- the description is continuous (indeed 1-Lipschitz) in f [10].

This description is the famous **persistence diagram**, or **barcode**.

In practical applications all the persistence modules that we encounter are finite, for the trivial reason that a computer only stores finite data. However, there are good grounds for extending the results of [13, 20, 10] to a more general setting. For example, theoretical guarantees are commonly formulated in terms of an idealised model, such as a continuous space to which the data form an approximation. Finiteness becomes unnatural and difficult to enforce in these ideal models, but one still wants the main results to be true.

Following [7], we say that a persistence module \mathbb{V} is **tame** if

$$r_s^t = \text{rank}(v_s^t) < \infty \quad \text{whenever } s < t.$$

Since the word ‘tame’ is overloaded with too many meanings in the persistence literature, we offer **q-tame** as an alternative (for reasons that will be explained later).

It is shown in [7] that persistence diagrams can be constructed for q-tame persistence modules, and that these diagrams are stable with respect to certain natural metrics. We reproduce these results here, using different methods for many of the arguments. We complete the picture by showing that the map from tame persistence modules to persistence diagrams is an isometry. This isometry theorem is due independently to Lesnick [16].

We believe that q-tame persistence modules are the ‘right’ class of objects to work with. This is for two complementary reasons: (i) we can prove almost everything we want to prove about tame modules and their persistence diagrams; and (ii) they occur in practice. For example, a continuous function on a finite simplicial complex has tame sublevelset persistent homology (Theorem 2.22). See [8] for many other examples.

1.2. Different index sets. We can define persistence modules over any partially ordered set \mathbf{T} , formally in the same way as for \mathbf{R} , by specifying indexed families

$$(V_t \mid t \in \mathbf{T}) \quad \text{and} \quad (v_s^t \mid s, t \in \mathbf{T}, s \leq t)$$

of vector spaces and linear maps, for which $v_r^t = v_s^t \circ v_r^s$ whenever $r \leq s \leq t$, and where v_t^t is the identity on V_t . The resulting collection of data is called a **T-persistence module** or a persistence module **over T**.

If \mathbb{V} is a **T**-persistence module and $\mathbf{S} \subset \mathbf{T}$, then we get an **S**-persistence module by considering only those spaces and maps with indices in **S**. This is called the **restriction** of \mathbb{V} to **S**, and may be written $\mathbb{V}_{\mathbf{S}}$.

Most commonly, we work with locally finite subsets $\mathbf{T} \subset \mathbf{R}$, that is, those subsets with no accumulation points in \mathbf{R} . We collect information about an **R**-persistence module by considering its restriction to finite or locally finite subsets. This works well because persistence modules over $\{1, 2, \dots, n\}$ or over the integers \mathbf{Z} are well understood.

In section 3, we will consider some other posets.

1.3. Module categories. A homomorphism Φ between two \mathbf{T} -persistence modules \mathbb{U}, \mathbb{V} is a collection of linear maps

$$(\phi_t : U_t \rightarrow V_t \mid t \in \mathbf{T})$$

such that the diagram

$$\begin{array}{ccc} U_s & \xrightarrow{u_s^t} & U_t \\ \phi_s \downarrow & & \downarrow \phi_t \\ V_s & \xrightarrow{v_s^t} & V_t \end{array}$$

commutes for all $s \leq t$. Composition is defined in the obvious way, as are identity homomorphisms. This makes the collection of persistence modules into a category.

The category contains kernel, image, and cokernel objects for every map Φ , and there is a zero object.

Write

$$\begin{aligned} \text{Hom}(\mathbb{U}, \mathbb{V}) &= \{\text{homomorphisms } \mathbb{U} \rightarrow \mathbb{V}\}, \\ \text{End}(\mathbb{V}) &= \{\text{homomorphisms } \mathbb{V} \rightarrow \mathbb{V}\}. \end{aligned}$$

Note that $\text{End}(\mathbb{V})$ is a \mathbf{k} -algebra.

Later we will consider homomorphisms that shift the index. We will introduce these when we consider the interleaving relation between persistence modules.

1.4. Interval modules. The basic building blocks of persistence are the interval modules. One seeks to understand a persistence module by decomposing it into intervals. This is not always possible, but it is possible often enough for our purposes.

Let $\mathbf{T} \subseteq \mathbf{R}$, and let $J \subseteq \mathbf{T}$ be an interval.² Then \mathbb{I}^J is defined to be the \mathbf{T} -persistence module with spaces

$$I_t = \begin{cases} \mathbf{k} & \text{if } t \in J \\ 0 & \text{otherwise} \end{cases}$$

and maps

$$i_s^t = \begin{cases} 1 & \text{if } s, t \in J \\ 0 & \text{otherwise} \end{cases}$$

In informal language, \mathbb{I}^J represents a ‘feature’ which ‘persists’ over the interval J but is absent elsewhere. We write $\mathbb{I}_{\mathbf{T}}^J$ when we wish to specify the index set unambiguously.

We now establish notation for describing interval modules.

If \mathbf{T} is a locally finite subset of \mathbf{R} , then any bounded interval contains its endpoints. We write these as closed intervals, and unbounded intervals as open or half-open intervals. For example, in the case of \mathbf{Z} -persistence modules there are four kinds of interval module:

$$\mathbb{I}^{[m,n]} \quad \mathbb{I}^{(-\infty,n]} \quad \mathbb{I}^{[m,+\infty)} \quad \mathbb{I}^{(-\infty,+\infty)}$$

²By this we mean that if $r < s < t$ are elements of T with $r, t \in J$, then $s \in J$.

For visual clarity, we sometimes lower the superscripts, writing

$$\mathbb{I}[m, n] \quad \mathbb{I}(-\infty, n] \quad \mathbb{I}[m, +\infty) \quad \mathbb{I}(-\infty, +\infty)$$

instead.

For persistence modules over \mathbf{R} , it becomes important to distinguish intervals which have the same endpoints but different topology (open, closed, half-open). For this purpose we introduce **decorated real numbers**, which are written as ordinary real numbers but with a superscript $+$ (plus) or $-$ (minus). For finite intervals we adopt the following dictionary:

$$\begin{aligned} (p^-, q^-) &\text{ means } [p, q) \\ (p^-, q^+) &\text{ means } [p, q] \\ (p^+, q^-) &\text{ means } (p, q) \\ (p^+, q^+) &\text{ means } (p, q] \end{aligned}$$

We require $p < q$, except for the special case (r^-, r^+) , which represents the 1-point interval $[r, r]$.

We also include the symbols $-\infty$ and $+\infty$, for infinite intervals. Since real intervals are always open at infinity these implicitly carry the superscripts $-\infty^+$ and $+\infty^-$, but we usually omit the superscripts. The interval notation extends in the obvious way, so

$$(-\infty, q^+) \text{ means } (-\infty, q]$$

for instance.

When we wish to refer to a decorated real number but we don't know what the decoration is, we use an asterisk. Thus p^* means p^+ or p^- . The notation for an arbitrary interval is (p^*, q^*) , where $p^* < q^*$ in the obvious ordering on decorated reals.

Here are some visual conventions for representing interval modules over \mathbf{R} .

We work in the half-plane

$$\mathcal{H} = \{(p, q) \mid p \leq q\}$$

of points in \mathbf{R}^2 which lie on or above the diagonal. Any finite interval module $\mathbb{I}(p^*, q^*)$ may be represented in several different ways (see Figure 1):

- as an interval in the real line;
- as $\text{rank}(i_s^t)$, viewed as a function $\mathcal{H} \rightarrow \{0, 1\}$;
- as a point (p, q) in \mathcal{H} , with a tick to specify the decoration.

Here are the four tick directions explicitly:

$$\begin{aligned} (p^-, q^+) &= \bullet \swarrow & (p^+, q^+) &= \bullet \nearrow \\ (p^-, q^-) &= \bullet \searrow & (p^+, q^-) &= \bullet \nwarrow \end{aligned}$$

The convention is that the tick points into the quadrant suggested by the decorations.

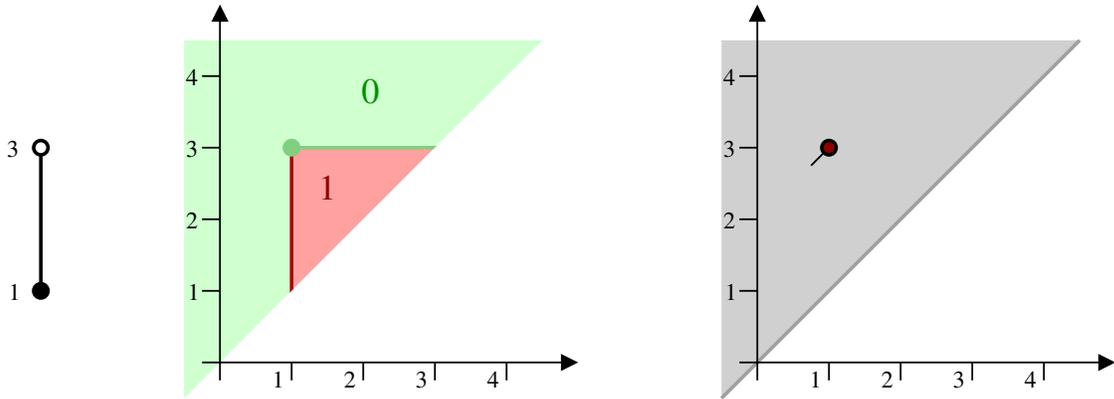


FIGURE 1. The interval, rank function, and decorated point representations of the interval module $\mathbb{I}[1, 3) = \mathbb{I}(1^-, 3^-)$.

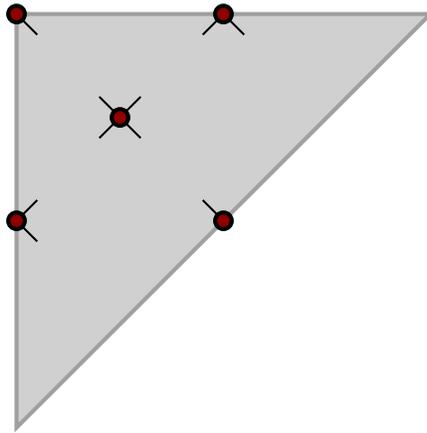


FIGURE 2. The extended half-plane $\overline{\mathcal{H}}$ with examples of each type of interval, drawn as points with ticks.

A fourth option is to draw the point (p, q) without indicating the decoration. This ‘forgetful’ representation is the classical convention for persistence diagrams, and is adequate for most purposes. However, the extra precision provided by decoration is important for the correspondence between diagrams and measures.

To represent an infinite interval as a (decorated or undecorated) point, we work in the extended half-plane

$$\overline{\mathcal{H}} = \mathcal{H} \cup \{-\infty\} \times \mathbf{R} \cup \mathbf{R} \times \{+\infty\} \cup \{(-\infty, +\infty)\}$$

which can be drawn schematically as a triangle. See Figure 2.

1.5. Interval decomposition. The **direct sum** $\mathbb{W} = \mathbb{U} \oplus \mathbb{V}$ of two persistence modules \mathbb{U}, \mathbb{V} is defined as follows:

$$W_t = U_t \oplus V_t, \quad w_s^t = u_s^t \oplus v_s^t$$

This generalises immediately to arbitrary (finite or infinite) direct sums.

A persistence module \mathbb{W} is **indecomposable** if the only decompositions $\mathbb{W} = \mathbb{U} \oplus \mathbb{V}$ are the trivial decompositions $\mathbb{W} \oplus 0$ and $0 \oplus \mathbb{W}$.

Direct sums play both a *synthetic* role and an *analytic* role in our theory. On the one hand, given an indexed family of intervals $(J_\ell \mid \ell \in L)$ we can synthesise a persistence module

$$\mathbb{V} = \bigoplus_{\ell \in L} \mathbb{I}^{J_\ell}$$

whose isomorphism type depends only on the multiset $\{J_\ell \mid \ell \in L\}$. In light of the direct-sum decomposition, we can think of \mathbb{V} as having an independent feature for each $\ell \in L$, supported over the interval J_ℓ . On the other hand, we can attempt to analyse a given persistence module \mathbb{V} by decomposing it into interval modules.

We now present the necessary theory for this. A ‘building block’ in a module category can be characterised by having a comparatively simple endomorphism ring. Interval modules have the simplest possible:

Proposition 1.1. *Let $\mathbb{I} = \mathbb{I}_{\mathbf{T}}^J$ be an interval module over $\mathbf{T} \subseteq \mathbf{R}$; then $\text{End}(\mathbb{I}) = \mathbf{k}$.*

Proof. Any endomorphism of \mathbb{I} acts on each nonzero $I_t = \mathbf{k}$ by scalar multiplication. By the commutation law, it is the same scalar for each t . \square

Proposition 1.2. *Interval modules are indecomposable.*

Proof. Given a decomposition $\mathbb{I} = \mathbb{U} \oplus \mathbb{V}$, the projection maps onto \mathbb{U} and \mathbb{V} are idempotents in the endomorphism ring of \mathbb{I} . But the only idempotents of $\text{End}(\mathbb{I}) = \mathbf{k}$ are 0 and 1. \square

Theorem 1.3 (Krull–Remak–Schmidt–Azumaya). *Suppose that a persistence module \mathbb{V} over $\mathbf{T} \subseteq \mathbf{R}$ can be expressed as a direct sum of interval modules in two different ways:*

$$\mathbb{V} = \bigoplus_{\ell \in L} \mathbb{I}^{J_\ell} = \bigoplus_{m \in M} \mathbb{I}^{K_m}$$

Then there is a bijection $\sigma : L \rightarrow M$ such that $J_\ell = K_{\sigma(\ell)}$ for all ℓ .

Proof. This is from Azumaya [2] (Theorem 1), plus the trivial observation that $\mathbb{I}^J \cong \mathbb{I}^K$ implies $J = K$. The theorem requires a certain condition on the endomorphism ring of each possible interval module: if $\alpha, \beta \in \text{End}(\mathbb{I})$ are non-isomorphisms then $\alpha + \beta$ is a non-isomorphism. Since each $\text{End}(\mathbb{I}) = \mathbf{k}$, the only non-isomorphism is the zero map and the condition is satisfied. \square

In other words, provided we can decompose a given persistence module \mathbb{V} as a direct sum of interval modules, then the multiset of intervals is an isomorphism invariant of \mathbb{V} . But when does such a decomposition exist?

Theorem 1.4 (Gabriel, Auslander, Ringel–Tachikawa, Webb). *Let \mathbb{V} be a persistence module over $\mathbf{T} \subseteq \mathbf{R}$. In each of the following situations, \mathbb{V} can be decomposed as a direct sum of interval modules:*

- (1) \mathbf{T} is a finite set.
- (2) \mathbf{T} is a locally finite subset of \mathbf{R} and each V_t is countable-dimensional.

On the other hand, (3) there exists a persistence module over \mathbf{Z} (indeed, over the nonpositive integers) which does not admit an interval decomposition.

We are grateful to William Crawley-Boevey for clarifying the situation for us, and for introducing us to the module with no interval decomposition. We acknowledge Lesnick [16] for drawing our attention to Webb [19].

Proof. (1) The decomposition of a diagram

$$V_1 \rightarrow V_2 \rightarrow \cdots \rightarrow V_n$$

into interval summands, when each $\dim(V_i)$ is finite, is one of the simpler instances of Gabriel's theorem [15]; see [20] or [4] for a concrete explanation. The extension to infinite-dimensional modules follows abstractly from a theorem of Auslander [1] and, independently, Ringel and Tachikawa [18]. Alternatively, note that the argument given in [4] does not really require finite-dimensionality.

(2) We may assume that $\mathbf{T} = \mathbf{Z}$ because any locally finite subset of \mathbf{R} is equivalent as an ordered set to a subset of \mathbf{Z} . Then this is a theorem of Webb [19].

(3) Webb [19] gives the following example, indexed over the nonpositive integers:

$$\begin{aligned} V_0 &= \{\text{sequences } (x_1, x_2, x_3, \dots) \text{ of real numbers}\} \\ V_{-n} &= \{\text{sequences with } x_1 = \cdots = x_n = 0\} \quad \text{for } n \geq 1 \end{aligned}$$

The maps v_{-m}^{-n} are the inclusions $V_{-m} \subset V_{-n}$ ($m \geq n$).

Suppose \mathbb{V} has an interval decomposition. Since each map v_{-n-1}^{-n} is injective, all of the intervals must be of the form $[-n, 0]$ or $(-\infty, 0]$. Since $\dim(V_{-n}/V_{-n-1}) = 1$, each interval $[-n, 0]$ occurs with multiplicity 1. Since $\bigcap V_{-n} = \{0\}$, the interval $(-\infty, 0]$ does not occur at all.

This would imply that $\mathbb{V} \cong \bigoplus_{n \geq 0} \mathbb{I}^{[-n, 0]}$. This contradicts the fact that $\dim(V_0)$ is uncountable. Therefore \mathbb{V} does not admit an interval decomposition after all. \square

For a persistence module which decomposes into intervals, the way is now clear to define its persistence diagram. Simply record which intervals occur in the decomposition (with multiplicity): see section 1.6. Theorem 1.3 tells us that this is an isomorphism invariant.

However, we have seen that arbitrary persistence modules over \mathbf{R} are not guaranteed an interval decomposition. Here are three ways around the problem:

- Work in restricted settings to ensure that the structure of \mathbb{V} depends only on finitely many index values $t \in \mathbf{R}$. For example, if X is a compact manifold and f is a Morse

function, then $H(\mathbb{X}_{\text{sub}})$ is determined by the finite sequence

$$H(X^{a_1}) \rightarrow H(X^{a_2}) \rightarrow \dots \rightarrow H(X^{a_n})$$

where a_1, a_2, \dots, a_n are the critical values of f . This is the traditional approach. In this setting, the word ‘tame’ is often used to demarcate pairs (X, f) for which $H(\mathbb{X}_{\text{sub}})$ is determined by a finite diagram of finite-dimensional vector spaces.

- Sample the persistence module \mathbb{V} over a finite grid. Consider limits as the grid converges to the whole real line. This is the approach taken in [7], where it was shown that the q-tame hypothesis is sufficient to guarantee good limiting behaviour.
- Show that the persistence intervals (in the decomposable case) can be inferred from the behaviour of \mathbb{V} on short finite index sets. Apply this indirect definition to define the persistence diagram in the non-decomposable case. This is the method of ‘rectangle measures’ developed in this paper.

Although Theorem 1.4 doesn’t cover it, it remains possible that there is a decomposition theorem for q-tame modules (in fact, it seems somewhat likely). Such a theorem would remove the impetus for some of the arguments in this paper.

1.6. The persistence diagram of a decomposable module. If a persistence module \mathbb{V} indexed over \mathbf{R} can be decomposed

$$\mathbb{V} \cong \bigoplus_{\ell \in L} \mathbb{I}(p_\ell^*, q_\ell^*),$$

then we define the **decorated persistence diagram** to be the multiset

$$\text{Dgm}(\mathbb{V}) = \text{Int}(\mathbb{V}) = \{(p_\ell^*, q_\ell^*) \mid \ell \in L\}$$

and the **undecorated persistence diagram** to be the multiset

$$\text{dgm}(\mathbb{V}) = \text{int}(\mathbb{V}) = \{(p_\ell, q_\ell) \mid \ell \in L\} - \Delta$$

where $\Delta = \{(r, r) \mid r \in \mathbf{R}\}$ is the diagonal in the plane.

Remark. In section 2.7 we will give a quite different definition of the persistence diagram of \mathbb{V} , based on the persistence measure rather than the interval decomposition. When occasionally we must distinguish between the two notions, we use the alternate names Int, int rather than Dgm, dgm for the diagrams defined here.

Theorem 1.3 implies that $\text{Dgm}(\mathbb{V})$ and $\text{dgm}(\mathbb{V})$ do not depend on the decomposition of \mathbb{V} .

Notice that Dgm is a multiset of decorated points in $\overline{\mathcal{H}}$, whereas dgm is a multiset of undecorated points in the interior of $\overline{\mathcal{H}}$. Here ‘interior’ means that we exclude the diagonal but keep the points at infinity. The information retained by dgm is precisely the information we care about later, when we discuss bottleneck distances. See section 4.

Let us see how these definitions play out in an example of the traditional kind.

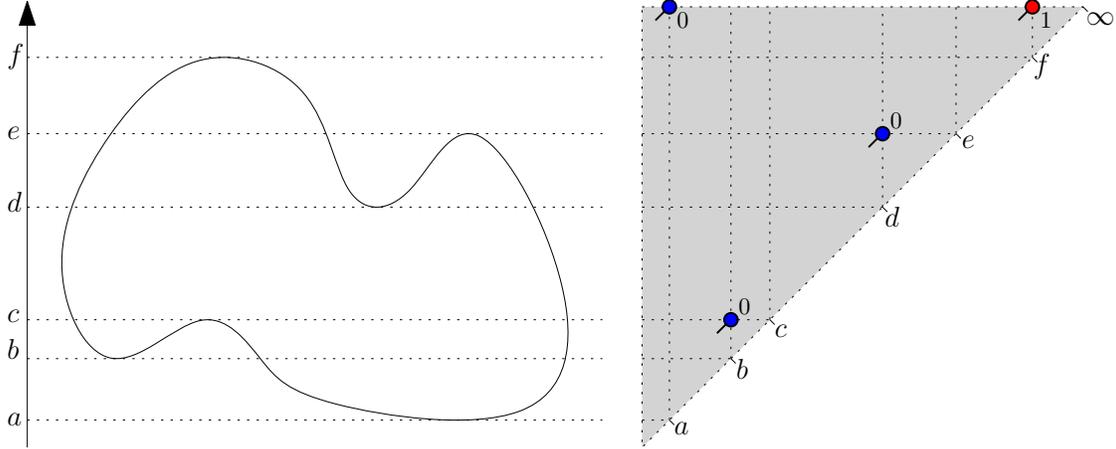


FIGURE 3. A traditional example in persistence theory: (left) X is a smoothly embedded curve in the plane, and f is its y -coordinate or ‘height’ function; (right) the decorated persistence diagram of $H(\mathbb{X}_{\text{sub}})$. There are three intervals in H_0 and one interval in H_1 .

Example 1.5. Consider the curve in \mathbf{R}^2 shown in Figure 3, filtered by the height function. The sublevelset persistent homology decomposes into half-open intervals as follows:

$$\begin{aligned} H_0(\mathbb{X}_{\text{sub}}) &\cong \mathbb{I}(a^-, +\infty) \oplus \mathbb{I}(b^-, c^-) \oplus \mathbb{I}(d^-, e^-) \\ H_1(\mathbb{X}_{\text{sub}}) &\cong \mathbb{I}(f^-, +\infty) \end{aligned}$$

If we let $[x]$ denote the chain corresponding to the critical point with critical value x , then the three summands of H_0 can be taken to be generated by $[a]$, $[b] - [a]$, and $[d] - [a]$ respectively. For instance, the independent 0-cycle $[b] - [a]$, which is born at time b , becomes a boundary at time c . Thus it gives rise to the half-open interval $[b, c) = (b^-, c^-)$.

Remark. For a Morse function on a compact manifold with critical values (a_i) , the intervals are always half-open, of type $[a_i, a_j) = (a_i^-, a_j^-)$. See section 2.9.

1.7. Quiver calculations. We now set up the notation and algebraic tools for handling persistence modules over a finite index set.

A persistence module \mathbb{V} indexed over a finite subset

$$\mathbf{T} : \quad a_1 < a_2 < \cdots < a_n$$

of the real line can be thought of as a diagram of n vector spaces and $n - 1$ linear maps:

$$\mathbb{V} : \quad V_{a_1} \rightarrow V_{a_2} \rightarrow \cdots \rightarrow V_{a_n}$$

Such a diagram is a **representation** of the following **quiver**:

$$\bullet \longrightarrow \bullet \longrightarrow \cdots \longrightarrow \bullet$$

We have seen (Theorem 1.4) that \mathbb{V} decomposes as a finite sum of interval modules $\mathbb{I}[a_i, a_j]$. When n is small, we can represent these interval modules pictorially. The following example illustrates how.

Example 1.6. Let $a < b < c$. There are six interval modules over $\{a, b, c\}$, namely:

$$\begin{aligned} \mathbb{I}[a, a] &= \bullet_a \text{---} \circ_b \text{---} \circ_c & \mathbb{I}[a, b] &= \bullet_a \text{---} \bullet_b \text{---} \circ_c & \mathbb{I}[a, c] &= \bullet_a \text{---} \bullet_b \text{---} \bullet_c \\ \mathbb{I}[b, b] &= \circ_a \text{---} \bullet_b \text{---} \circ_c & \mathbb{I}[b, c] &= \circ_a \text{---} \bullet_b \text{---} \bullet_c \\ \mathbb{I}[c, c] &= \circ_a \text{---} \circ_b \text{---} \bullet_c \end{aligned}$$

In the notation we use filled circles \bullet to indicate where the module has rank 1, and clear circles \circ to indicate where the module has rank 0. The connecting maps have full rank.

Now let \mathbb{V} be a persistence module indexed over \mathbf{R} . For any finite set of indices

$$\mathbf{T} : \quad a_1 < a_2 < \cdots < a_n$$

and any interval $[a_i, a_j] \subseteq \mathbf{T}$, we define the multiplicity of $[a_i, a_j]$ in $\mathbb{V}_{\mathbf{T}}$ to be the number of copies of $\mathbb{I}[a_i, a_j]$ to occur in the interval decomposition of $\mathbb{V}_{\mathbf{T}}$. This takes values in the set $\{0, 1, 2, \dots, \infty\}$. (We do not distinguish different infinite cardinals.)

It is useful to have notation for these multiplicities. Again, we define by example.

Example 1.7. We write

$$\langle [b, c] \mid \mathbb{V}_{a,b,c} \rangle \quad \text{or} \quad \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \mid \mathbb{V} \rangle$$

for the multiplicity of $\circ_a \text{---} \bullet_b \text{---} \bullet_c$ in the 3-term module

$$\mathbb{V}_{a,b,c} = (V_a \rightarrow V_b \rightarrow V_c).$$

When \mathbb{V} is clear from the context, we may simply write

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \rangle.$$

The abbreviation $\langle [b, c] \rangle$ is not permitted since it is ambiguous. For example, $\langle [b, c] \mid V_{b,c} \rangle$ and $\langle [b, c] \mid V_{a,b,c} \rangle$ are not generally the same. See Proposition 1.10 and Example 1.11.

Example 1.8. The invariants of a single linear map $V_a \xrightarrow{v} V_b$ are:

$$\begin{aligned} \text{rank}(v) &= \langle \bullet_a \text{---} \bullet_b \mid \mathbb{V} \rangle \\ \text{nullity}(v) &= \langle \bullet_a \text{---} \circ_b \mid \mathbb{V} \rangle \\ \text{conullity}(v) &= \langle \circ_a \text{---} \bullet_b \mid \mathbb{V} \rangle \end{aligned}$$

(The conullity is the dimension of the cokernel.)

Proposition 1.9 (direct sums). *Suppose a persistence module \mathbb{V} can be written as a direct sum*

$$\mathbb{V} = \bigoplus_{\ell \in L} \mathbb{V}^{\ell}$$

Then

$$\langle [a_i, a_j] \mid \mathbb{V}_{\mathbf{T}} \rangle = \sum_{\ell \in L} \langle [a_i, a_j] \mid \mathbb{V}_{\mathbf{T}}^{\ell} \rangle$$

for any index set $\mathbf{T} = \{a_1, a_2, \dots, a_n\}$ and interval $[a_i, a_j] \subseteq \mathbf{T}$.

Proof. Each summand $\mathbb{V}_{\mathbf{T}}^{\ell}$ can be decomposed separately into interval modules. Putting these together we get an interval decomposition of $\mathbb{V}_{\mathbf{T}}$. The number of summands of a given type in $\mathbb{V}_{\mathbf{T}}$ is then equal to the total number of summands of that type in all of the $\mathbb{V}_{\mathbf{T}}^{\ell}$. \square

Often we wish to compare multiplicities of intervals in different finite restrictions of \mathbb{V} . The principle is very simple:

Proposition 1.10 (restriction principle). *Let \mathbf{S}, \mathbf{T} be finite index sets with $\mathbf{S} \subset \mathbf{T}$. Then*

$$\langle \mathbb{I} \mid \mathbb{V}_{\mathbf{S}} \rangle = \sum_{\mathbb{J}} \langle \mathbb{J} \mid \mathbb{V}_{\mathbf{T}} \rangle$$

where the sum is over those intervals $\mathbb{J} \subseteq \mathbf{T}$ which restrict over \mathbf{S} to \mathbb{I} .

Proof. Take an arbitrary interval decomposition of $\mathbb{V}_{\mathbf{T}}$. This induces an interval decomposition of $\mathbb{V}_{\mathbf{S}}$. Summands of $\mathbb{V}_{\mathbf{S}}$ of type \mathbb{I} arise precisely from those summands of $\mathbb{V}_{\mathbf{T}}$ of types \mathbb{J} as above. \square

Example 1.11. Let $a < b < c$. Consider new indices p and q , arranged $a < p < b < q < c$. Then

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \rangle = \langle \circ_a \text{---} \bullet_b \text{---} \bullet_q \text{---} \bullet_c \rangle$$

whereas

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \rangle = \langle \circ_a \text{---} \circ_p \text{---} \bullet_b \text{---} \bullet_c \rangle + \langle \circ_a \text{---} \bullet_p \text{---} \bullet_b \text{---} \bullet_c \rangle.$$

The extra term occurs when the inserted new index occurs between a clear node and a filled node, because then there are two possible intervals which restrict to the original interval.

We will make frequent use of the restriction principle. Here is a simple illustration, to serve as a template for similar arguments that we will encounter later on.

Example 1.12. Consider the standard fact that $\text{rank}(V_b \rightarrow V_c) \geq \text{rank}(V_a \rightarrow V_d)$ when $a \leq b \leq c \leq d$. The proof using quiver notation runs as follows:

$$\begin{aligned} \text{rank}(V_b \rightarrow V_c) &= \langle \text{---} \bullet_b \text{---} \bullet_c \text{---} \rangle \\ &= \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle + \text{three other terms} \\ &\geq \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle \\ &= \langle \bullet_a \text{---} \text{---} \bullet_d \rangle \\ &= \text{rank}(V_a \rightarrow V_d) \end{aligned}$$

The ‘three other terms’ are

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle, \quad \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle, \quad \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle$$

as indicated by the restriction principle.

2. RECTANGLE MEASURES

For a decomposable \mathbf{R} -persistence module

$$\mathbb{V} \cong \bigoplus_{\ell \in L} \mathbb{I}(p_{\ell}^*, q_{\ell}^*),$$

we have defined the decorated persistence diagram to be the multiset

$$\text{Dgm}(\mathbb{V}) = \{(p_{\ell}^*, q_{\ell}^*) \mid \ell \in L\},$$

and the undecorated persistence diagram to be the multiset

$$\text{dgm}(\mathbb{V}) = \{(p_\ell, q_\ell) \mid \ell \in L\}.$$

If we don't know that \mathbb{V} is decomposable, then we have to proceed differently. The rough idea is that if we know how many points of Dgm are contained in each rectangle in the upper-half space, then we know Dgm itself. For persistence modules, counting points in rectangles turns out to be easy.

The language of measure theory is well suited to this argument. We will show that each persistence module defines an integer-valued measure on rectangles. If the module is tame then this measure is finite-valued, and therefore (Theorem 2.8) it is concentrated at a discrete set of points: this is the persistence diagram.

In the decomposable case, we will see that this agrees with the definition above. And when the module is not known to be decomposable, we can proceed regardless.

The persistence measures that we construct are not true measures on subsets of \mathbf{R}^2 : they are additive in the sense of tiling rather than in the usual sense of disjoint set union. The discrepancy arises when we split a rectangle into two: what happens to the points on the common edge? To which rectangle do they belong? In resolving this, one is naturally led to the notion of decorated points. This fits perfectly with our earlier use of decorations to distinguish open and closed interval ends.

2.1. The persistence measure. Let \mathbb{V} be a persistence module. The persistence measure of \mathbb{V} is the function

$$\mu_{\mathbb{V}}(R) = \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{V} \rangle$$

defined on rectangles $R = [a, b] \times [c, d]$ in the plane with $a < b \leq c < d$.

For a decomposable persistence module, there is a clear relationship between $\mu_{\mathbb{V}}$ and the interval summands of \mathbb{V} . Let us first consider the case of an interval module.

Proposition 2.1. *Let $\mathbb{V} = \mathbb{I}^J$ where $J = (p^*, q^*)$ is a real interval. Let $R = [a, b] \times [c, d]$ where $a < b \leq c < d$. Then*

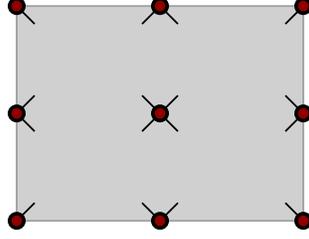
$$\mu_{\mathbb{V}}(R) = \begin{cases} 1 & \text{if } [b, c] \subseteq J \subseteq (a, d) \\ 0 & \text{otherwise} \end{cases}$$

Proof. It is clear that \mathbb{I}^J restricted to $\{a, b, c, d\}$ is an interval or is zero. Thus, $\mu_{\mathbb{V}}(R) \leq 1$. Moreover $\mu_{\mathbb{V}}(R) = 1$ precisely when

$$\mathbb{I}_{a,b,c,d}^J = \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d,$$

which happens if and only if $b, c \in J$ and $a, d \notin J$. This is equivalent to the condition $[b, c] \subseteq J \subseteq (a, d)$. \square

Proposition 2.1 has a graphical interpretation. Represent the interval $J \subseteq \mathbf{R}$ as a decorated point in the extended plane. The following picture indicates exactly which decorated points (p^*, q^*) are detected by $\mu_{\mathbb{V}}(R)$:



If (p, q) is in the interior of R then (p^*, q^*) is always detected regardless of the decoration. If (p, q) is on the boundary, then (p^*, q^*) is detected if the tick is directed inwards.

We formalise this by defining a membership relation between decorated points and rectangles.

Definition. Let $R = [a, b] \times [c, d]$ where $a < b \leq c < d$, and consider a decorated point (p^*, q^*) with $p^* < q^*$. We write $(p^*, q^*) \in R$ if

- (i) the interval $J = [p^*, q^*]$ satisfies $[b, c] \subseteq J \subseteq (a, d)$;

or equivalently

- (ii) the point (p, q) and its decoration tick are contained in the closed rectangle R .

We use the notation

$$R^\blacksquare = \{(p^*, q^*) \in R\}$$

when we wish to explicitly name the set of decorated points contained in R . However, we prefer to write $(p^*, q^*) \in R$ instead of $(p^*, q^*) \in R^\blacksquare$ when we are simply indicating the relationship between a decorated point and a rectangle.

With this understanding we state the following counting result.

Corollary 2.2. *Suppose \mathbb{V} is a decomposable persistence module over \mathbf{R} :*

$$\mathbb{V} = \bigoplus_{\ell \in L} \mathbb{I}[p_\ell^*, q_\ell^*]$$

Then:

$$(*) \quad \mu_{\mathbb{V}}(R) = \text{card}(\text{Dgm}(\mathbb{V})|_R)$$

Proof. This follows immediately from Proposition 2.1 and Proposition 1.9 (direct sums). \square

We can now articulate our strategy for defining the persistence diagram without assuming that the module \mathbb{V} is decomposable:

- construct the persistence measure $\mu_{\mathbb{V}}$;
- let $\text{Dgm}(\mathbb{V})$ be a multiset in the half-plane such that $(*)$ holds for all rectangles R .

To make this work, we need to know that such a multiset exists and is unique. This is the content of Theorem 2.8, under the hypothesis that $\mu_{\mathbb{V}}$ is finite and additive. The result is a sort of ‘Riesz’ representation theorem for measures on rectangles.

When \mathbb{V} is decomposable, Corollary 2.2 confirms that our new definition agrees with the old.

2.2. The persistence measure (continued). We call $\mu_{\mathbb{V}}$ a measure because it is additive with respect to splitting a rectangle into two rectangles. We prove this shortly. First, we give a new proof of an ‘alternating sum’ formula for $\mu_{\mathbb{V}}(R)$ that appears in [10].

Proposition 2.3. *Let \mathbb{V} be a persistence module, and let $a < b \leq c < d$. If the spaces V_a, V_b, V_c, V_d are finite-dimensional, or less stringently if $r_b^c < \infty$, then*

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{V} \rangle = r_b^c - r_a^c - r_b^d + r_a^d.$$

(Here as before $r_s^t = \text{rank}(v_s^t : V_s \rightarrow V_t)$.)

Proof. Decompose the 4-term module $\mathbb{V}_{a,b,c,d}$ into intervals. The left-hand side counts intervals of type $[b, c]$. By the restriction principle, the four terms on the right-hand side evaluate as follows:

$$\begin{aligned} r_b^c &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle + \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle + \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle + \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle \\ r_a^c &= \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle + \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle \\ r_b^d &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle + \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle \\ r_a^d &= \langle \bullet_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_d \rangle \end{aligned}$$

These expressions are all finite: the hypothesis $r_b^c < \infty$ implies that the other three ranks are finite too (Example 1.12). We can legitimately take the alternating sum, whereupon all terms cancel except for the $\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle$. \square

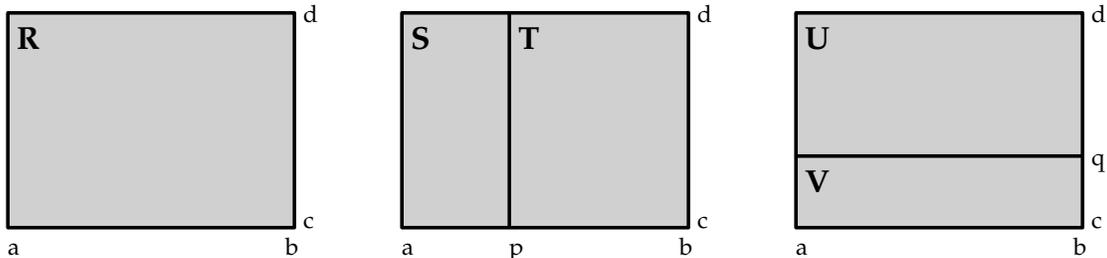
We give three proofs of additivity. The first is completely general, whereas the other two work under restricted settings but are illuminating in their own way.

Proposition 2.4. $\mu_{\mathbb{V}}$ is additive under vertical and horizontal splitting, meaning that

$$\begin{aligned} \mu_{\mathbb{V}}([a, b] \times [c, d]) &= \mu_{\mathbb{V}}([a, p] \times [c, d]) + \mu_{\mathbb{V}}([p, b] \times [c, d]) \\ \mu_{\mathbb{V}}([a, b] \times [c, d]) &= \mu_{\mathbb{V}}([a, b] \times [c, q]) + \mu_{\mathbb{V}}([a, b] \times [q, d]) \end{aligned}$$

whenever $a < p < b \leq c < q < d$.

This additivity property is illustrated by the following figure



where the claim is that $\mu_{\mathbb{V}}(R) = \mu_{\mathbb{V}}(S) + \mu_{\mathbb{V}}(T) = \mu_{\mathbb{V}}(U) + \mu_{\mathbb{V}}(V)$.

First proof. Let $a < p < b \leq c < q < d$. Then we calculate

$$\begin{aligned}
\mu_{\mathbb{V}}([a, b] \times [c, d]) &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle \\
&= \langle \circ_a \text{---} \bullet_p \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle + \langle \circ_a \text{---} \circ_p \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle \\
&= \langle \circ_a \text{---} \bullet_p \text{---} \bullet_c \text{---} \circ_d \rangle + \langle \text{---} \circ_p \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle \\
&= \mu_{\mathbb{V}}([a, p] \times [c, d]) + \mu_{\mathbb{V}}([p, b] \times [c, d])
\end{aligned}$$

for additivity with respect to a horizontal split, and

$$\begin{aligned}
\mu_{\mathbb{V}}([a, b] \times [c, d]) &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \text{---} \circ_d \rangle \\
&= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_q \text{---} \circ_d \rangle + \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \bullet_q \text{---} \circ_d \rangle \\
&= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_q \text{---} \text{---} \rangle + \langle \circ_a \text{---} \bullet_b \text{---} \text{---} \bullet_q \text{---} \circ_d \rangle \\
&= \mu_{\mathbb{V}}([a, b] \times [c, q]) + \mu_{\mathbb{V}}([a, b] \times [q, d])
\end{aligned}$$

for additivity with respect to a vertical split. \square

Second proof, assuming $r_b^c < \infty$. The alternating sum formula (Proposition 2.3) gives

$$r_b^c - r_a^c - r_b^d + r_a^d = (r_p^c - r_a^c - r_p^d + r_a^d) + (r_b^c - r_p^c - r_b^d + r_p^d)$$

and

$$r_b^c - r_a^c - r_b^d + r_a^d = (r_b^c - r_a^c - r_b^q + r_a^q) + (r_b^q - r_a^q - r_b^d + r_a^d)$$

as required. Note that $r_b^c < \infty$ implies that $r_p^c, r_b^q < \infty$, so the formula is valid for all the rectangles involved. \square

This second proof is particularly transparent when drawn geometrically in the plane: the + and - signs at the corners of the rectangles cancel in a pleasant way.

Third proof, assuming \mathbb{V} is decomposable. By Corollary 2.2, the measure of a rectangle is equal to the number of interval summands whose corresponding decorated points lie in the rectangle. Additivity now follows from the elementary observation that a decorated point in R belongs to exactly one of S and T , and to exactly one of U and V . \square

We finish this section with two further descriptions of $\mu_{\mathbb{V}}([a, b] \times [c, d])$.

Proposition 2.5. *We have the following formulae:*

$$\begin{aligned}
\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{V} \rangle &= \dim \left[\frac{\text{im}(v_b^c) \cap \ker(v_c^d)}{\text{im}(v_a^c) \cap \ker(v_c^d)} \right] \\
&= \dim \left[\frac{\ker(v_b^d)}{\ker(v_b^c) + \text{im}(v_a^b) \cap \ker(v_b^d)} \right]
\end{aligned}$$

Proof. This is covered, for instance, in the localisation discussion in section 5.1 of [4]. The two formulae are obtained by localising at c, b respectively. \square

Proposition 2.5 expresses the measure of a rectangle as the dimension of a vector space constructed functorially from \mathbb{V} . (Ostensibly there are two vector spaces, one for each formula, but the map v_b^c induces a natural isomorphism between them.) The functoriality has its uses, but in other regards this characterisation is quite hard to use. For instance, additivity is not at all obvious in this formulation.

2.3. Abstract r-measures. We now consider rectangle measures more abstractly. Persistence measures are of course our primary example, but the general formulation allows for many other situations.

For ease of exposition, we initially work in the plane \mathbf{R}^2 rather than the extended plane $\overline{\mathbf{R}}^2$. The picture is completed in section 2.6 when we discuss the points at infinity.

Definition. Let \mathcal{D} be a subset of \mathbf{R}^2 . Define

$$\text{Rect}(\mathcal{D}) = \{[a, b] \times [c, d] \subset \mathcal{D} \mid a < b \text{ and } c < d\}$$

(the set of closed rectangles contained in \mathcal{D}). A **rectangle measure** or **r-measure** on \mathcal{D} is a function

$$\mu : \text{Rect}(\mathcal{D}) \rightarrow \{0, 1, 2, \dots\} \cup \{\infty\}$$

which is additive under vertical and horizontal splitting (as in Proposition 2.4).

Proposition 2.6. *Let μ be an r-measure on $\mathcal{D} \subseteq \mathbf{R}^2$. Then μ is:*

(Finitely additive) If $R \in \text{Rect}(\mathcal{D})$ can be written as a union $R = R_1 \cup \dots \cup R_k$ of rectangles with disjoint interiors, then $\mu(R) = \mu(R_1) + \dots + \mu(R_k)$.

(Monotone) If $R \subseteq S$ then $\mu(R) \leq \mu(S)$.

Proof. (Finitely additive) Let $R = [a, b] \times [c, d]$. By induction and the vertical splitting property, it follows that finite additivity holds for decompositions of the form

$$R = \bigcup_i R_i$$

where $R_i = [a_i, a_{i+1}] \times [c, d]$ with $a = a_1 < a_2 < \dots < a_m = b$.

By induction and the horizontal splitting property, it then follows that finite additivity holds for ‘product’ decompositions

$$R = [a, b] \times [c, d] = \bigcup_{i,j} R_{ij}$$

where $R_{ij} = [a_i, a_{i+1}] \times [c_j, c_{j+1}]$ with $a = a_1 < a_2 < \dots < a_m = b$ and $c = c_1 < c_2 < \dots < c_n = d$.

For an arbitrary decomposition $R = R_1 \cup \dots \cup R_k$, the result follows by considering a product decomposition of R by which each R_i is itself product-decomposed.

(Monotone) Decompose S into a collection of interior disjoint rectangles R, R_1, \dots, R_{k-1} , one of which is R . (This can be done with at most 9 rectangles using a product decomposition.)

Then

$$\begin{aligned}\mu(S) &= \mu(R) + \mu(R_1) + \cdots + \mu(R_{k-1}) \\ &\geq \mu(R)\end{aligned}$$

by finite additivity and the fact that $\mu \geq 0$. □

Here is one more plausible-and-true statement about abstract r-measures.

Proposition 2.7 (Subadditivity). *Let μ be an r-measure on $\mathcal{D} \subseteq \mathbf{R}^2$. If a rectangle $R \in \text{Rect}(\mathcal{D})$ is contained in a finite union*

$$R \subseteq R_1 \cup \cdots \cup R_k$$

of rectangles $R_i \in \text{Rect}(\mathcal{D})$, then

$$\mu(R) \leq \mu(R_1) + \cdots + \mu(R_k).$$

Proof. Let

$$a_1 < a_2 < \cdots < a_m$$

include all the x -coordinates of the corners of all the rectangles, and let

$$c_1 < c_2 < \cdots < c_n$$

include all the y -coordinates. Each rectangle is then tiled as a union of pieces

$$[a_i, a_{i+1}] \times [c_j, c_{j+1}]$$

with disjoint interiors, and the measure of the rectangle is the sum of the measures of its tiles, by additivity. Since each tile belonging to R must also belong to one or more of the R_i , the inequality follows. □

2.4. Equivalence of measures and diagrams. We wish to establish a correspondence between r-measures and decorated diagrams. The task of defining a continuous persistence diagram can then be replaced by the simpler task of defining an r-measure. For this to work, the measure has to be finite.

A useful notion is the **r-interior** of a region $\mathcal{D} \subseteq \mathbf{R}^2$, defined

$$\mathcal{D}^\blacksquare = \{(p^*, q^*) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p^*, q^*) \in R\}.$$

This is the set of decorated points that are ‘seen’ by the rectangles in \mathcal{D} . The decorated diagram will be a multiset in \mathcal{D}^\blacksquare . Clearly an r-measure cannot tell us what happens outside \mathcal{D}^\blacksquare .

The **interior** of \mathcal{D} in the classical sense is written \mathcal{D}° . It may be defined

$$\mathcal{D}^\circ = \{(p, q) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p, q) \in R^\circ\},$$

supposing that we already agree that the interior of a closed rectangle $R = [a, b] \times [c, d]$ is the open rectangle $R^\circ = (a, b) \times (c, d)$. The undecorated diagram will be a multiset in \mathcal{D}° .

Theorem 2.8 (The equivalence theorem). *Let $\mathcal{D} \subseteq \mathbf{R}^2$. There is a bijective correspondence between:*

- *Finite r-measures μ on \mathcal{D} . Here ‘finite’ means that $\mu(R) < \infty$ for every $R \in \text{Rect}(\mathcal{D})$.*
- *Locally finite multisets \mathbf{A} in \mathcal{D}^\blacksquare . Here ‘locally finite’ means that $\text{card}(\mathbf{A}|_R) < \infty$ for every $R \in \text{Rect}(\mathcal{D})$.*

The measure μ corresponding to a multiset \mathbf{A} is related to it by the formula

$$(2.9) \quad \mu(R) = \text{card}(\mathbf{A}|_R)$$

for every $R \in \text{Rect}(\mathcal{D})$.

Remark. We can write equation (2.9) equivalently as

$$(2.10) \quad \mu(R) = \sum_{(p^*, q^*) \in R} m(p^*, q^*),$$

where

$$m : \mathcal{D}^\blacksquare \rightarrow \{0, 1, 2, \dots\}$$

is the multiplicity function for \mathbf{A} .

The theorem leads immediately to the following definitions. Let μ be a finite r-measure on a region $\mathcal{D} \subset \mathbf{R}^2$.

(i) The **decorated diagram** of μ is the unique locally finite multiset $\text{Dgm}(\mu)$ in \mathcal{D}^\blacksquare such that

$$\mu(R) = \text{card}(\text{Dgm}(\mu)|_R)$$

for every $R \in \text{Rect}(\mathcal{D})$.

(ii) The **undecorated diagram** of μ is the locally finite multiset in \mathcal{D}°

$$\text{dgm}(\mu) = \{(p, q) \mid (p^*, q^*) \in \text{Dgm}(\mu)\} \cap \mathcal{D}^\circ$$

obtained by forgetting the decorations on the points and restricting to the interior.

Remark. Note that **dgm** is locally finite in \mathcal{D}° , but not necessarily locally finite in \mathbf{R}^2 —it may have accumulation points on the boundary of \mathcal{D} .

Proof of Theorem 2.8. One direction of the correspondence is easy. If \mathbf{A} is a multiset on \mathcal{D}^\blacksquare then the function $\mu(R)$ on rectangles defined by equation (2.9) is indeed an r-measure. It is finite if \mathbf{A} is locally finite. To verify additivity, suppose that a rectangle R is split vertically or horizontally into two rectangles R_1, R_2 . Notice that every decorated point $(p^*, q^*) \in R$ belongs to exactly one of R_1, R_2 . It follows that

$$\mu(R) = \text{card}(\mathbf{A}|_R) = \text{card}(\mathbf{A}|_{R_1}) + \text{card}(\mathbf{A}|_{R_2}) = \mu(R_1) + \mu(R_2),$$

as required.

The reverse direction takes more work. Given an r-measure μ we will (1) construct a multiset \mathbf{A} in \mathcal{D}^\blacksquare , (2) show that μ and \mathbf{A} are related by equation (2.9), and (3) show that \mathbf{A} is unique.

In practice we will work with the multiplicity function m and equation (2.10), rather than referring to \mathbf{A} directly.

Step 1. Let μ be a finite r-measure on \mathcal{D} . For (p^*, q^*) in \mathcal{D}^\blacksquare , define

$$(2.11) \quad \mathfrak{m}(p^*, q^*) = \min \{ \mu(R) \mid R \in \text{Rect}(\mathcal{D}), (p^*, q^*) \in R \}.$$

Note that the minimum is attained because the set is nonempty and μ takes values in the natural numbers.

Here is an alternative characterisation. Rather than minimising over all rectangles, we can take the limit through a decreasing sequence of rectangles:

Lemma 2.12. *Let (ξ_i) and (η_i) be non-increasing sequences of positive real numbers which tend to zero as $i \rightarrow \infty$. Then*

$$\mathfrak{m}(p^+, q^+) = \lim_{i \rightarrow \infty} \mu([p, p + \xi_i] \times [q, q + \eta_i]),$$

and similarly

$$\mathfrak{m}(p^+, q^-) = \lim_{i \rightarrow \infty} \mu([p, p + \xi_i] \times [q - \eta_i, q]),$$

$$\mathfrak{m}(p^-, q^+) = \lim_{i \rightarrow \infty} \mu([p - \xi_i, p] \times [q, q + \eta_i]),$$

$$\mathfrak{m}(p^-, q^-) = \lim_{i \rightarrow \infty} \mu([p - \xi_i, p] \times [q - \eta_i, q]).$$

Proof. The key observation is that the sequence of rectangles $R_i = [p, p + \xi_i] \times [q, q + \eta_i]$ is cofinal in the set of rectangles R containing (p^+, q^+) . In other words, for any such R we have $R_i \subseteq R$ for all sufficiently large i .

By monotonicity, the sequence of nonnegative integers $\mu(R_i)$ is non-increasing, and hence eventually stabilises to a limit. Then

$$\mathfrak{m}(p^+, q^+) \leq \min_i \mu(R_i) = \lim_{i \rightarrow \infty} \mu(R_i) \leq \mu(R)$$

for any R containing (p^+, q^+) . Taking the minimum over all R , the right-hand side becomes $\mathfrak{m}(p^+, q^+)$ and hence by squeezing

$$\mathfrak{m}(p^+, q^+) = \lim_{i \rightarrow \infty} \mu(R_i).$$

The other three cases of the lemma are similar. □

We return to the main proof.

Step 2. Having defined $\mathfrak{m}(p^*, q^*)$, we now show that this is the ‘correct’ definition, meaning that equation (2.10) is satisfied. We have seen already that \mathfrak{m} corresponds to an r-measure

$$(2.13) \quad \nu(R) = \sum_{(p^*, q^*) \in R} \mathfrak{m}(p^*, q^*),$$

and it remains to show (for this step) that $\nu = \mu$. We prove this by induction on $k = \mu(R)$.

Base case. $\mu(R) = 0$. Then for every $(p^*, q^*) \in R$ we have

$$0 \leq \mathfrak{m}(p^*, q^*) \leq \mu(R) = 0$$

so $\nu(R) = 0$.

Inductive step. Suppose $\mu(R) = \nu(R)$ for every rectangle R with $\mu(R) < k$. Consider a rectangle R_0 with $\mu(R_0) = k$. We must show that $\nu(R_0) = k$.

Split the rectangle into four equal quadrants S_1, S_2, S_3, S_4 . Certainly

$$\begin{aligned}\mu(R_0) &= \mu(S_1) + \mu(S_2) + \mu(S_3) + \mu(S_4) \\ \nu(R_0) &= \nu(S_1) + \nu(S_2) + \nu(S_3) + \nu(S_4)\end{aligned}$$

by finite additivity (Proposition 2.6). If every quadrant satisfies $\mu(S_i) < k$, then by induction we deduce that $\mu(R_0) = \nu(R_0)$. Otherwise, one of the quadrants has $\mu = k$ and the other three quadrants satisfy $\mu = 0$ (and hence $\nu = 0$). Let R_1 be the distinguished quadrant, so $\mu(R_1) = k$. It is now enough to show that $\nu(R_1) = k$.

We repeat the argument. Subdivide R_i into four equal quadrants. Either all four quadrants satisfy the inductive hypothesis $\mu < k$, in which case we are done. Otherwise we find a quadrant R_{i+1} with $\mu(R_{i+1}) = k$, and we are reduced to showing that $\nu(R_{i+1}) = k$.

In the worst case—the remaining unresolved case—this iteration never terminates and we obtain a sequence of closed rectangles

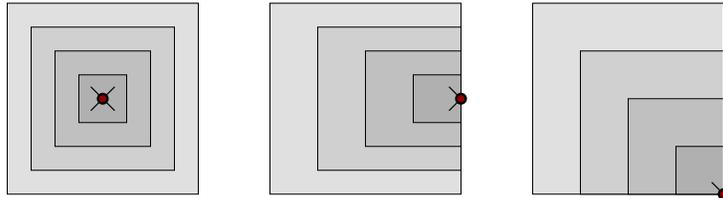
$$R_0 \supset R_1 \supset R_2 \supset \dots$$

each being a quadrant of the previous one, with $\mu(R_i) = k$. Since the diameters of the rectangles tend to zero, their intersection $\bigcap_i R_i$ contains a single point (r, s) .

We are now in a position to show that $\nu(R_0) = k$, by evaluating the sum explicitly over all decorated points in R_0 .

First of all, consider decorated points that eventually leave the sequence (R_i) . Specifically, suppose that $(p^*, q^*) \in R_0$ but $(p^*, q^*) \notin R_{i-1} - R_i$ for some i . This means that (p^*, q^*) belongs to one of the three quadrants of R_{i-1} for which $\mu = 0$. It follows immediately that $m(p^*, q^*) = 0$.

Thus, the only contribution to $\nu(R_0)$ comes from decorated points (p^*, q^*) which belong to every rectangle in the sequence (R_i) . Clearly these must be decorated versions (r^*, s^*) of the intersection point (r, s) . There are 4, 2 or 1 of them depending on how the nested sequence of rectangles converges to its limit. Here we illustrate the three cases:



Suppose first that (r, s) lies in the interior of every rectangle R_i , so that all four decorated points (r^+, s^+) , (r^+, s^-) , (r^-, s^+) , (r^-, s^-) belong to every R_i . Divide each R_i into 4 subrectangles R_i^{++} , R_i^{+-} , R_i^{-+} , R_i^{--} , which share a common corner at (r, s) so that each of the four decorated points (r^*, s^*) belongs to one of the subrectangles in the obvious notation.

By Lemma 2.12,

$$\begin{aligned} m(r^+, s^+) &= \lim_{i \rightarrow \infty} \mu(R_i^{++}) & m(r^+, s^-) &= \lim_{i \rightarrow \infty} \mu(R_i^{+-}) \\ m(r^-, s^+) &= \lim_{i \rightarrow \infty} \mu(R_i^{-+}) & m(r^-, s^-) &= \lim_{i \rightarrow \infty} \mu(R_i^{--}) \end{aligned}$$

and moreover each of these decreasing integer sequences eventually stabilises at its limiting value. Thus, for sufficiently large i ,

$$\begin{aligned} \nu(R_0) &= m(r^+, s^+) + m(r^+, s^-) + m(r^-, s^+) + m(r^-, s^-) \\ &= \mu(R_i^{++}) + \mu(R_i^{+-}) + \mu(R_i^{-+}) + \mu(R_i^{--}) \\ &= \mu(R_i) \\ &= k \end{aligned}$$

as required.

A similar argument (with fewer terms) can be made in the cases where only 2 or 1 of the decorated points (r^*, s^*) belong to every R_i . For instance, if (r, s) lies on the interior of the right-hand edge of the rectangles (R_i) for all sufficiently large i , we split each rectangle into two parts R_i^{-+} and R_i^{--} and obtain

$$\nu(R_0) = m(r^-, s^+) + m(r^-, s^-) = \mu(R_i^{-+}) + \mu(R_i^{--}) = \mu(R_i) = k$$

in the same way. In this case (r^+, s^+) and (r^+, s^-) eventually leave (or were never in) the sequence (R_i) and therefore do not contribute to $\nu(R_0)$. We omit the details of the remaining cases, which are equally straightforward.

This completes the inductive step. Thus $\mu(R) = \nu(R)$ for every $R \in \text{Rect}(\mathcal{D})$.

Step 3. Suppose $m'(p^*, q^*)$ is some other multiplicity function on \mathcal{D}^\blacksquare whose associated r-measure

$$\nu'(R) = \sum_{(p^*, q^*) \in R} m'(p^*, q^*)$$

satisfies $\mu = \nu'$. We must show that $m = m'$.

Consider an arbitrary decorated point $(p^*, q^*) \in \mathcal{D}^\blacksquare$. Let R be a rectangle which contains (p^*, q^*) at its corner. Since

$$\nu(R) = \nu'(R) = \mu(R) < \infty,$$

there are only finitely many other decorated points $(r^*, s^*) \in R$ with positive multiplicity in m or m' . By making R smaller, we can therefore assume that (p^*, q^*) is the only decorated point in R with positive multiplicity in either measure. Then

$$m(p^*, q^*) = \nu(R) = \mu(R) = \nu'(R) = m'(p^*, q^*).$$

Since (p^*, q^*) was arbitrary it follows that $m = m'$.

This completes the proof of Theorem 2.8. □

2.5. **Non-finite measures.** If a measure is not everywhere finite, we restrict our attention to the parts of the plane where it is finite. Define the **finite r-interior** of an r-measure μ to be the set of decorated points

$$\mathcal{F}^\blacksquare(\mu) = \{(p^*, q^*) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p^*, q^*) \in R \text{ and } \mu(R) < \infty\}.$$

The **finite interior** is

$$\mathcal{F}^\circ(\mu) = \{(p, q) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p, q) \in R^\circ \text{ and } \mu(R) < \infty\}.$$

This is an open subset of the plane, being a union of open rectangles. It is easy to see that $(p, q) \in \mathcal{F}^\circ(\mu)$ if and only if $(p^*, q^*) \in \mathcal{F}^\blacksquare(\mu)$ for all possible decorations.

We can apply Theorem 2.8 to each rectangle R of finite measure to obtain a decorated diagram in R^\blacksquare . These diagrams will agree, by uniqueness, on the common intersection of any two such rectangles. Therefore we can combine these local definitions to get a well-defined multiset $\text{Dgm}(\mu)$ in the whole of $\mathcal{F}^\blacksquare(\mu)$.

This multiset has the property that $\mu(R) = \text{card}(\text{Dgm}(\mu)|_R)$ for any rectangle $R \in \text{Rect}(\mathcal{D})$ with finite measure. There is one subtlety to address.

Proposition 2.14. *Let $R \in \text{Rect}(\mathcal{D})$. If $R^\blacksquare \subseteq \mathcal{F}^\blacksquare(\mu)$ then $\mu(R) < \infty$.*

Proof. We show that each $(p, q) \in R$ is contained in the relative interior of a rectangle $S \subseteq R$ of finite measure. Then R , being compact, is the union of finitely many of these rectangles, and by subadditivity (Proposition 2.7) it must have finite measure.

If (p, q) lies in the interior of R , then each of the four decorated points (p^*, q^*) belongs to $\mathcal{F}^\blacksquare(\mu)$ so we can find four finite-measure rectangles containing them. The union of these rectangles contains a neighbourhood of (p, q) , and we can take $S \subseteq R$ to be a rectangle contained in this union with (p, q) in its interior. It has finite measure, by subadditivity.

If (p, q) lies on the interior of an edge, we take two finite-measure rectangles containing a relative neighbourhood of (p, q) ; and if (p, q) is a corner point we take just one rectangle. \square

Corollary 2.15. *Let μ be an r-measure on $\mathcal{D} \subseteq \mathbf{R}^2$. Then there is a uniquely defined locally finite multiset $\text{Dgm}(\mu)$ in $\mathcal{F}^\blacksquare(\mu)$ such that*

$$\mu(R) = \text{card}(\text{Dgm}(\mu)|_R)$$

for every $R \in \text{Rect}(\mathcal{D})$ with $R^\blacksquare \subseteq \mathcal{F}^\blacksquare(\mu)$. \square

Similarly, or as an easy consequence of Proposition 2.14, we have:

Proposition 2.16. *Let $R \in \text{Rect}(\mathcal{D})$. If $R \subseteq \mathcal{F}^\circ(\mu)$ then $\mu(R) < \infty$.* \square

Now we can define the diagrams of a general r-measure.

- The **decorated diagram** of an r-measure μ is the pair $(\text{Dgm}(\mu), \mathcal{F}^\blacksquare(\mu))$, where $\text{Dgm}(\mu)$ is the multiset in $\mathcal{F}^\blacksquare(\mu)$ described in the corollary.

- The **undecorated diagram** is the pair $(\text{dgm}(\mu), \mathcal{F}^\circ(\mu))$, where

$$\text{dgm}(\mu) = \{(p, q) \mid (p^*, q^*) \in \text{Dgm}(\mu)\} \cap \mathcal{F}^\circ(\mu)$$

is the locally finite³ multiset in $\mathcal{F}^\circ(\mu)$ obtained by forgetting the decorations in $\text{Dgm}(\mu)$ and restricting to the finite interior.

Remark. The point of Proposition 2.14 is that there is no ‘hidden information’ in μ beyond what is recorded in the diagram. The measure of any rectangle R with $R^\blacksquare \subseteq \mathcal{F}^\blacksquare(\mu)$ is recovered by counting decorated points of $\text{Dgm}(\mu)$, and any other rectangle has infinite measure.

We can bring these definitions into accord with the previous one (for finite r-measures) by agreeing that $\text{Dgm}(\mu)$ and $\text{dgm}(\mu)$ are abbreviations for $(\text{Dgm}(\mu), \mathcal{D}^\blacksquare)$ and $(\text{dgm}(\mu), \mathcal{D}^\circ)$, when \mathcal{D} is clear from the context.

It is sometimes useful to adopt the following **extension convention**. An r-measure defined on a subset $\mathcal{D} \subset \mathbf{R}^2$ can be interpreted as an r-measure on the whole plane \mathbf{R}^2 , by agreeing that $\mu(R) = \infty$ for any rectangle that meets $\mathbf{R}^2 - \mathcal{D}$. The extension has the same diagram as the original r-measure.

2.6. The diagram at infinity. We now discuss r-measures in the extended plane $\overline{\mathbf{R}}^2$. This requires infinite rectangles and decorated points at infinity.

- For $\mathcal{D} \subseteq \overline{\mathbf{R}}^2$, the set $\text{Rect}(\mathcal{D})$ consists of all rectangles $R \subseteq \mathcal{D}$ of the form

$$R = [a, b] \times [c, d]$$

where now $-\infty \leq a < b \leq +\infty$ and $-\infty \leq c < d \leq +\infty$.

- Decorated points in $\overline{\mathbf{R}}^2$ are pairs (p^*, q^*) where p^* and q^* are decorated real numbers or $-\infty^+$ or $+\infty^-$. The symbols $-\infty^+, +\infty^-$ may be abbreviated to $-\infty, +\infty$.

With these preliminaries, several other concepts are formally unchanged. The r-interior and the interior of \mathcal{D} are

$$\begin{aligned} \mathcal{D}^\blacksquare &= \{(p^*, q^*) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p^*, q^*) \in R\}, \\ \mathcal{D}^\circ &= \{(p, q) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p, q) \in R^\circ\}. \end{aligned}$$

One subtlety is that we take R° to be the relative interior of R as a subspace of $\overline{\mathbf{R}}^2$. For instance, if $R = [-\infty, b] \times [c, d]$ where b, c, d are finite, then $R^\circ = [-\infty, b) \times (c, d]$.

The concept of an r-measure goes through as before: it is a function

$$\mu : \text{Rect}(\mathcal{D}) \rightarrow \{0, 1, 2, \dots\} \cup \{\infty\}$$

which is additive with respect to horizontal and vertical splittings. An r-measure has a finite r-interior and a finite interior:

$$\begin{aligned} \mathcal{F}^\blacksquare(\mu) &= \{(p^*, q^*) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p^*, q^*) \in R \text{ and } \mu(R) < \infty\}, \\ \mathcal{F}^\circ(\mu) &= \{(p, q) \mid \text{there exists } R \in \text{Rect}(\mathcal{D}) \text{ such that } (p, q) \in R^\circ \text{ and } \mu(R) < \infty\}. \end{aligned}$$

³As before, dgm is locally finite in \mathcal{F}° , but may have accumulation points on the boundary of \mathcal{F}° .

We sometimes write $\mathcal{S}(\mu) = \overline{\mathbf{R}^2} - \mathcal{F}^\circ(\mu)$, the **singular support** of μ .

Claim. Corollary 2.15 is valid for r-measures on $\mathcal{D} \subseteq \overline{\mathbf{R}^2}$.

Proof. The statements (and indeed the proofs) of Theorem 2.8 and Corollary 2.15 are invariant under reparametrisations of the plane of the form

$$x' = f(x), \quad y' = g(y),$$

where f, g are homeomorphic embeddings. We can view $\overline{\mathbf{R}^2}$ as a rectangle in \mathbf{R}^2 via a transformation of this type; for instance

$$x' = \arctan(x), \quad y' = \arctan(y)$$

identifies $\overline{\mathbf{R}^2}$ with the rectangle $[-\pi/2, \pi/2] \times [-\pi/2, \pi/2]$ in \mathbf{R}^2 . The theorem and its corollary now apply directly. \square

The decorated diagram $(\text{Dgm}(\mu), \mathcal{F}^\blacksquare(\mu))$ and the undecorated diagram $(\text{dgm}(\mu), \mathcal{F}^\circ(\mu))$ are defined as before.

Let \mathbb{V} be a persistence module. The persistence measure $\mu_{\mathbb{V}}$, previously defined by

$$\mu_{\mathbb{V}}([a, b] \times [c, d]) = \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{V} \rangle$$

for $a < b \leq c < d$, extends easily to infinite rectangles. We allow for the possibility that $a = -\infty$ or $d = +\infty$ by setting

$$V_{-\infty} = 0, \quad V_{+\infty} = 0.$$

In these cases the alternating sum formula of Proposition 2.3 becomes:

$$\begin{aligned} \mu_{\mathbb{V}}([-\infty, b] \times [c, +\infty]) &= \langle \bullet_b \text{---} \bullet_c \mid \mathbb{V} \rangle &&= r_b^c \\ \mu_{\mathbb{V}}([a, b] \times [c, +\infty]) &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \mid \mathbb{V} \rangle &&= r_b^c - r_a^c \quad (\text{if } r_a^c < \infty) \\ \mu_{\mathbb{V}}([-\infty, b] \times [c, d]) &= \langle \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{V} \rangle &&= r_b^c - r_b^d \quad (\text{if } r_b^d < \infty) \end{aligned}$$

The first of these corresponds to the ‘ k -triangle lemma’ of [10].

So $\mu_{\mathbb{V}}$ can be thought of as an r-measure on the extended half-plane

$$\overline{\mathcal{H}} = \{(p, q) \mid -\infty \leq p \leq q \leq +\infty\}$$

and its diagram $\text{Dgm}(\mu_{\mathbb{V}})$ is defined on the subset of the r-interior

$$\overline{\mathcal{H}}^\blacksquare = \{(p^*, q^*) \mid -\infty^+ \leq p^* < q^* \leq +\infty^-\}$$

on which $\mu_{\mathbb{V}}$ is finite.

Notice that there is a 1–1 correspondence between nonempty real intervals and elements of $\overline{\mathcal{H}}^\blacksquare$. The interval corresponding to (p^*, q^*) is written (p^*, q^*) as before. Proposition 2.1 and Corollary 2.2 extend straightforwardly to this setting. In particular:

Corollary 2.17. *If \mathbb{V} is decomposable into interval modules then $\mu_{\mathbb{V}}(R)$ counts the interval summands corresponding to decorated points which lie in R .* \square

We finish this section by defining the ‘measures at infinity’ induced by an r-measure μ . The extended plane has 4 lines at infinity

$$(-\infty, \mathbf{R}), \quad (+\infty, \mathbf{R}), \quad (\mathbf{R}, -\infty), \quad (\mathbf{R}, +\infty)$$

and 4 points at infinity

$$(-\infty, -\infty), \quad (+\infty, -\infty), \quad (-\infty, +\infty), \quad (+\infty, +\infty)$$

and, depending on where μ is finite, there are measures defined on each of these. On the lines at infinity these are ‘interval measures’ (the 1-dimensional analogue of r-measures), and at the points at infinity these are simply numbers.

We write out the three cases of direct relevance to persistence modules. The other cases are similar.

- **the line** $(-\infty, \mathbf{R})$:

$$\mu(-\infty, [c, d]) = \lim_{b \rightarrow -\infty} \mu([-\infty, b] \times [c, d]) = \min_b \mu([-\infty, b] \times [c, d])$$

for any interval $[c, d] \subseteq \mathbf{R}$.

- **the line** $(\mathbf{R}, +\infty)$:

$$\mu([a, b], +\infty) = \lim_{c \rightarrow +\infty} \mu([a, b] \times [c, +\infty]) = \min_c \mu([a, b] \times [c, +\infty])$$

for any interval $[a, b] \subseteq \mathbf{R}$.

- **the point** $(-\infty, +\infty)$:

$$\mu(-\infty, +\infty) = \lim_{e \rightarrow +\infty} \mu([-\infty, -e] \times [e, +\infty]) = \min_e \mu([-\infty, -e] \times [e, +\infty])$$

Monotonicity of μ guarantees that each limit exists, except when none of the rectangles in the expression belongs to $\text{Rect}(\mathcal{D})$. When that happens we adopt the extension convention (see section 2.5), and regard these rectangles as having infinite measure and the limit as being ∞ .

In order for the limit to be finite, we need at least one rectangle in the limiting expression to have finite measure. By considering the decorated diagram in that rectangle, it becomes clear what the ‘measure at infinity’ measures: it counts the decorated points of $\text{Dgm}(\mu)$ that lie in the segment, or corner, at infinity.

For example, when $\mu(-\infty, [c, d])$ is finite, it counts the decorated points of $\text{Dgm}(\mu)$ of the form $(-\infty^+, q^*)$ where $q^* \in [c, d]$. When it is infinite, it means that some such $(-\infty^+, q^*)$ does not belong to $\mathcal{F}^\blacksquare(\mu)$, either because it does not belong to \mathcal{D}^\blacksquare or because all rectangles containing it have infinite mass.

2.7. Diagrams of persistence modules. We now have two competing definitions of the decorated diagram of a persistence module.

- If $\mathbb{V} = \bigoplus_{\ell \in L} \mathbb{I}(p_\ell^*, q_\ell^*)$ is decomposable into intervals, then define

$$\text{Dgm}(\mathbb{V}) = \text{Int}(\mathbb{V}) = \{(p_\ell^*, q_\ell^*) \mid \ell \in L\}.$$

This is a multiset in $\overline{\mathcal{H}}^\blacksquare$.

- Let $\mu_{\mathbb{V}}$ be the persistence measure of \mathbb{V} . Then

$$\text{Dgm}(\mathbb{V}) = \text{Dgm}(\mu_{\mathbb{V}}).$$

This is a multiset in $\mathcal{F}^\blacksquare(\mu)$, the finite r-interior of $\mu = \mu_{\mathbb{V}}$.

The overloading of the term ‘ $\text{Dgm}(\mathbb{V})$ ’ is partially excused by the following fact.

Proposition 2.18. *If \mathbb{V} is decomposable into intervals, then $\text{Int}(\mathbb{V})$ agrees with $\text{Dgm}(\mu_{\mathbb{V}})$ where the latter is defined, that is, on $\mathcal{F}^\blacksquare(\mu)$.*

Proof. By Corollary 2.17 we have

$$\text{card}(\text{Int}(\mathbb{V})|_R) = \mu_{\mathbb{V}}(R)$$

for all rectangles. On the other hand, we have

$$\text{card}(\text{Dgm}(\mu_{\mathbb{V}})|_R) = \mu_{\mathbb{V}}(R)$$

for all rectangles with $\mu_{\mathbb{V}}(R) < \infty$. By uniqueness, it follows that $\text{Int}(\mathbb{V})$ and $\text{Dgm}(\mu_{\mathbb{V}})$ must be the same multiset when restricted to $\mathcal{F}^\blacksquare(\mu_{\mathbb{V}})$. \square

Neither definition strictly outperforms the other, as the following examples show.

Example 2.19. Let

$$\mathbb{V} = \bigoplus_{\ell \in L} \mathbb{I}(p_\ell^*, q_\ell^*)$$

where the undecorated pairs (p_ℓ, q_ℓ) form a dense subset of the half-plane $\overline{\mathcal{H}}$. Then $\text{Int}(\mathbb{V})$ is defined; but $\mu_{\mathbb{V}}(R) = \infty$ for every rectangle, so $\mathcal{F}^\blacksquare(\mu_{\mathbb{V}})$ is the empty set and $\text{Dgm}(\mu_{\mathbb{V}})$ is nowhere defined.

Example 2.20. Adapting the example of Webb [19] to the real line:

$$W_t = 0 \quad \text{for } t > 0$$

$$W_0 = \{\text{sequences } (x_1, x_2, x_3, \dots) \text{ of real numbers}\}$$

$$W_t = \{\text{sequences with } x_n = 0 \text{ for all } n \leq |t|\} \quad \text{for } t < 0$$

As before, this is not decomposable into intervals. On the other hand, it is easy to see that

$$\langle \circ_a \text{---} \bullet_b \mid \mathbb{W} \rangle = \text{conullity}(W_a \rightarrow W_b) < \infty$$

except when $a = -\infty$, and

$$\langle \bullet_c \text{---} \circ_d \mid \mathbb{W} \rangle = \text{nullity}(W_c \rightarrow W_d) < \infty$$

except when $c \leq 0 < d$. Each of these terms dominates

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{W} \rangle,$$

which is therefore finite for all rectangles which do not contain $(-\infty, 0^+)$. It is not difficult to complete the argument that

$$\mathcal{F}^\blacksquare(\mu) = \overline{\mathcal{H}}^\blacksquare - \{(-\infty, 0^+)\}$$

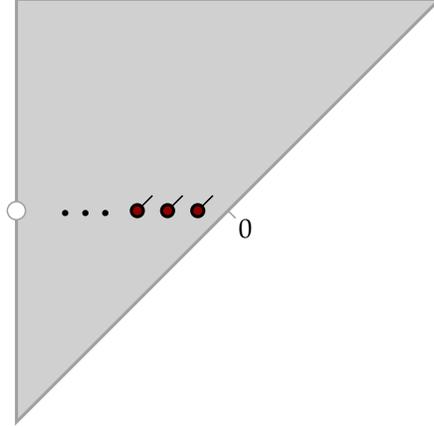


FIGURE 4. The persistence diagram $\text{Dgm}(\mathbb{W}) = \text{Dgm}(\mu_{\mathbb{W}})$ for the example of Webb. This is defined everywhere in the extended half-plane except $(-\infty, 0^+)$.

and to determine the persistence diagram

$$\text{Dgm}(\mu_{\mathbb{W}}) = \{(-n^+, 0^+) \mid n = 1, 2, 3, \dots\}$$

over $\mathcal{F}^{\blacksquare}(\mu)$. See Figure 4. In other words, although this example is not decomposable into intervals, its persistence measure is well-behaved away from the singular support, which is the point $(-\infty, 0^+)$.

2.8. Tameness conditions. We now describe several different levels of ‘tameness’ for a persistence module, starting with the tamest.

A persistence module \mathbb{V} is of **finite type**, if it is a finite direct sum of interval modules.

It is **locally finite**, if it is a direct sum of interval modules such that every $t \in \mathbf{R}$ has a neighbourhood which meets only finitely many of the intervals. Since closed bounded intervals are compact, this is equivalent to the stronger assertion that every bounded subset of \mathbf{R} meets only finitely many of the intervals.

Proposition 2.21. *A persistence module \mathbb{V} is locally finite if and only if (i) each V_t is finite-dimensional and (ii) there is a locally finite set $\mathbf{S} \subset \mathbf{R}$ such that v_b^c is an isomorphism for every pair $b < c$ with $[b, c] \cap \mathbf{S} = \emptyset$.*

In less formal language, condition (ii) asserts that \mathbb{V} is constant over each interval of the open set $\mathbf{R} - \mathbf{S}$.

Proof. Suppose conditions (i) and (ii) hold. Select a locally finite countable set

$$\mathbf{T} : \quad \dots < t_{-n} < \dots < t_{-1} < t_0 < t_1 < \dots < t_n < \dots$$

unbounded in both directions, where the even-numbered points include the entire set \mathbf{S} . By the result of Webb used in Theorem 1.4, we can decompose $\mathbb{V}_{\mathbf{T}}$ into intervals. Since \mathbb{V} is constant over each (t_{2n}, t_{2n+2}) , this extends uniquely to a decomposition of \mathbb{V} into interval

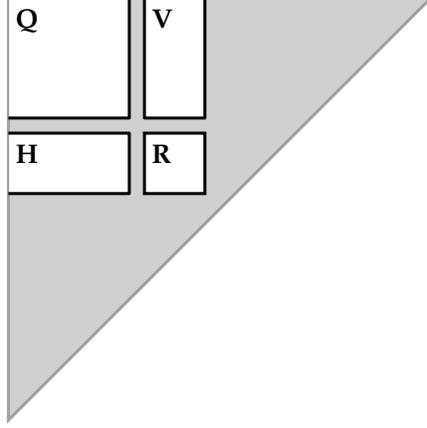


FIGURE 5. A quadrant, horizontal strip, vertical strip, and finite rectangle in $\overline{\mathcal{H}}$.

modules. Any bounded subset of \mathbf{R} meets only finitely many of the t_{2n} and (t_{2n}, t_{2n+2}) . Applying condition (i) we find that it meets only finitely many of the intervals.

The converse direction is immediate: take \mathbf{S} to be the set of endpoints of the intervals in the decomposition of \mathbb{V} . \square

For a finite or locally finite persistent module, it is easy to see that $\mathcal{F}^\blacksquare(\mu) = \overline{\mathcal{H}}$, so both persistence diagrams are defined on all of $\overline{\mathcal{H}}$, including on the diagonal and at infinity. Diagonal points must be decorated (p^-, p^+) , since only these belong to rectangles in $\overline{\mathcal{H}}$.

We introduce four more kinds of tameness. The assumptions here are about finiteness of $\mu_{\mathbb{V}}$ over different types of rectangle (see Figure 5). Each condition guarantees the existence of the persistence diagram over a certain subset of the extended half-plane. The finite part of the plane (except the diagonal) is always included; it is at infinity that the four conditions differ.

(i) We say that \mathbb{V} is **q-tame**, if $\mu_{\mathbb{V}}(Q) < \infty$ for every quadrant Q not touching the diagonal. In other words

$$\langle \bullet_b \text{---} \bullet_c \mid \mathbb{V} \rangle < \infty$$

(that is, $r_b^c < \infty$) for all $b < c$. The persistence diagram $\text{Dgm}(\mu_{\mathbb{V}})$ is defined over the set:

$$\{(p^*, q^*) \mid -\infty \leq p < q \leq +\infty\} = \blacksquare$$

(ii) We say that \mathbb{V} is **h-tame**, if $\mu_{\mathbb{V}}(H) < \infty$ for every horizontally infinite strip H not touching the diagonal. In other words,

$$\langle \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{V} \rangle < \infty$$

for all $b < c < d$. The persistence diagram $\text{Dgm}(\mu_{\mathbb{V}})$ is defined over the set:

$$\{(p^*, q^*) \mid -\infty \leq p < q < +\infty\} = \blacksquare$$

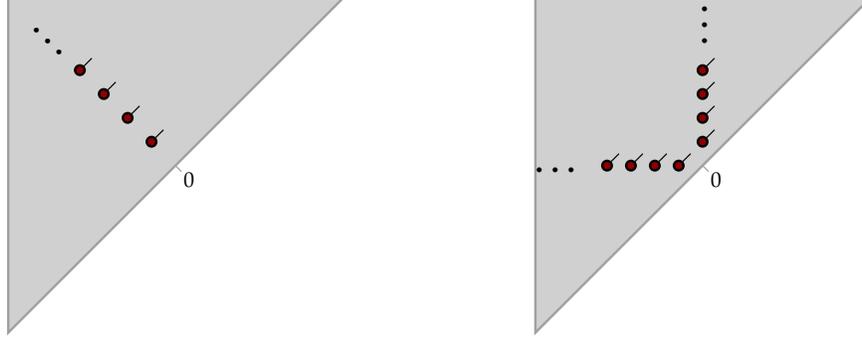


FIGURE 6. Diagrams of persistence modules which are: (left) h-tame and v-tame but not q-tame; (right) r-tame but not h-tame or v-tame.

(iii) We say that \mathbb{V} is **v-tame**, if $\mu_{\mathbb{V}}(V) < \infty$ for every vertically infinite strip V not touching the diagonal. In other words,

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \mid \mathbb{V} \rangle < \infty$$

for all $a < b < c$. The persistence diagram $\text{Dgm}(\mu_{\mathbb{V}})$ is defined over the set:

$$\{(p^*, q^*) \mid -\infty < p < q \leq +\infty\} = \blacktriangle$$

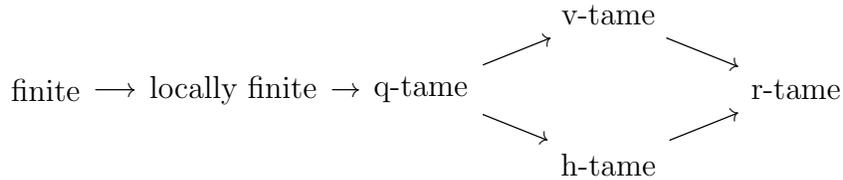
(iv) We say that \mathbb{V} is **r-tame**, if $\mu_{\mathbb{V}}(R) < \infty$ for every finite rectangle R not touching the diagonal. In other words,

$$\langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{V} \rangle < \infty$$

for all $a < b < c < d$. The persistence diagram $\text{Dgm}(\mu_{\mathbb{V}})$ is defined over the set:

$$\{(p^*, q^*) \mid -\infty < p < q < +\infty\} = \blacktriangle$$

Here is the diagram of inclusions between the different classes of tame persistence modules:



One can show that all the inclusions are strict, and also

$$\text{q-tame} \subsetneq \text{h-tame} \cap \text{v-tame}, \quad \text{h-tame} \cup \text{v-tame} \subsetneq \text{r-tame}.$$

Examples for these last two assertions are suggested by the diagrams in Figure 6.

We have not determined whether

$$\text{h-tame} \oplus \text{v-tame} = \text{r-tame}.$$

If an r-tame module is decomposable then it can certainly be written as the direct sum of an h-tame module and a v-tame module, by partitioning the intervals appropriately. The general situation seems more subtle.

Later we will show that ‘q-tame’ can be thought of as the completion of ‘locally finite’: a persistence module is q-tame if and only if it can be approximated arbitrarily well by locally finite modules. See Theorem 4.19.

Many naturally occurring persistence modules are q-tame. Here is a typical example.

Theorem 2.22. *Let X be a finite polyhedron,⁴ and let $f : X \rightarrow \mathbf{R}$ be a continuous function. Then persistent homology $H(\mathbb{X}_{\text{sub}})$ of the sublevelset filtration of (X, f) is q-tame.*

Proof. For any $b < c$ we must show that

$$H(X^b) \rightarrow H(X^c)$$

has finite rank. Begin with any triangulation of X , and subdivide it repeatedly until no simplex meets both $f^{-1}(b)$ and $f^{-1}(c)$. If we define Y to be the union of the closed simplices which meet X^b , then we have

$$X^b \subseteq Y \subseteq X^c$$

and hence the factorisation

$$H(X^b) \rightarrow H(Y) \rightarrow H(X^c).$$

Since Y is a finite polyhedron, $H(Y)$ is finite dimensional, so $H(X^b) \rightarrow H(X^c)$ has finite rank. \square

Remark. Individual terms vector spaces $H(X^b)$ in the persistence module may be infinite-dimensional. For instance let X be a closed disk in the plane, and let f be a non-negative function whose zero set is a Hawaiian earring; then $H(X^0)$ is uncountable-dimensional. However, the theorem shows that such pathological behaviour is ephemeral; it does not persist over any nontrivial interval of parameter values.

Other examples are developed in [8]. It is shown, for instance, that the Vietoris–Rips and Čech complexes of a pre-compact metric space have q-tame persistent homology.

With a little more work we can show:

Theorem 2.23. *Let X be a locally compact polyhedron, and let $f : X \rightarrow \mathbf{R}$ be a proper continuous function. Then persistent homology $H(\mathbb{X}_{\text{sub}})$ of the sublevelset filtration of (X, f) is h-tame and v-tame (and r-tame).*

Proof. h-tameness. Let $b < c < d$. We must show that

$$(h^*) \quad \langle \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid H(\mathbb{X}_{\text{sub}}) \rangle < \infty.$$

Begin with a triangulation of X . Only finitely many simplices meet the compact set $f^{-1}(b)$, so again after a finite number of subdivisions no simplex meets both $f^{-1}(b)$ and $f^{-1}(c)$.

Now let Y be the union of the closed simplices which meet X^b , and let Z be the union of the closed simplices which meet X^d . This gives a diagram of inclusions

$$X^b \subseteq Y \subseteq X^c \subseteq X^d \subseteq Z.$$

⁴By ‘polyhedron’ we mean the realisation of a simplicial complex as a topological space.

Note that the polyhedron Z differs from its subpolyhedron Y by the addition of only finitely many simplices, since each such simplex must meet the compact set $f^{-1}[b, d]$. Thus the relative homology $H(Z, Y)$ is finite-dimensional.

We now work with the induced homology diagram

$$H(X^b) \rightarrow H(Y) \rightarrow H(X^c) \rightarrow H(X^d) \rightarrow H(Z).$$

In the obvious notation,

$$\begin{aligned} \langle \bullet_b \text{---} \bullet_c \text{---} \circ_d \text{---} \rangle &= \langle \bullet_b \text{---} \bullet_y \text{---} \bullet_c \text{---} \circ_d \text{---} \circ_z \rangle \\ &\leq \langle \text{---} \bullet_y \text{---} \text{---} \circ_z \rangle \\ &= \dim[\ker(H(Y) \rightarrow H(Z))]. \end{aligned}$$

By the homology long exact sequence for the pair (Z, Y) , we have

$$\ker(H(Y) \rightarrow H(Z)) = \text{im}(H(Z, Y) \rightarrow H(Y))$$

which is finite-dimensional. This confirms (h-*).

v-tameness. Let $a < b < c$. We must show that

$$(v^*) \quad \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \mid H(\mathbb{X}_{\text{sub}}) \rangle < \infty$$

Using a similar argument to the above, we construct a diagram of inclusions

$$Y \subseteq X^a \subseteq X^b \subseteq Z \subseteq X^c$$

where Y, Z are polyhedra with $H(Z, Y)$ finite-dimensional. Working with the homology diagram

$$H(Y) \rightarrow H(X^a) \rightarrow H(X^b) \rightarrow H(Z) \rightarrow H(X^c),$$

we estimate

$$\begin{aligned} \langle \text{---} \circ_a \text{---} \bullet_b \text{---} \bullet_c \rangle &= \langle \circ_y \text{---} \circ_a \text{---} \bullet_b \text{---} \bullet_z \text{---} \bullet_c \rangle \\ &\leq \langle \circ_y \text{---} \text{---} \bullet_z \text{---} \rangle \\ &= \dim[\text{coker}(H(Y) \rightarrow H(Z))] \end{aligned}$$

By the homology long exact sequence of the pair (Z, Y) , we have

$$\text{coker}(H(Y) \rightarrow H(Z)) \cong \text{im}(H(Z) \rightarrow H(Z, Y))$$

which is finite-dimensional. This confirms (v-*). □

Remark. The theorem essentially asserts that the only possible bad behaviour of $H(\mathbb{X}_{\text{sub}})$ occurs at $(-\infty, +\infty)$. It is easy to construct examples which are definitely not q-tame. The simplest example is $X = \mathbf{Z}$ (the integers), with $f(n) = n$. The 0-homology of any sublevelset is infinite dimensional, and all inclusions have infinite rank.

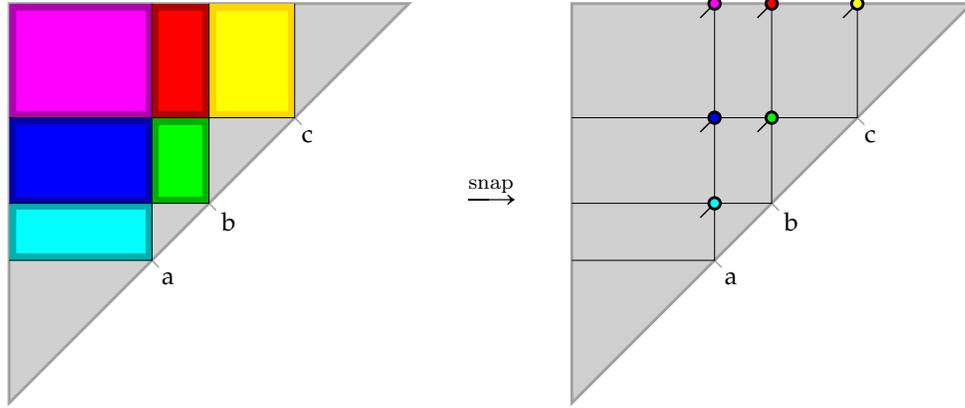


FIGURE 7. A persistence module \mathbb{V} , discretised at a, b, c . The persistence diagram $\text{Dgm}(\mathbb{V}_{a,b,c})$ is localised at six grid vertices. The multiplicity of each vertex is equal to the number of decorated points of $\text{Dgm}(\mathbb{V})$ in the rectangle immediately below and to the left of it. Decorated points of $\text{Dgm}(\mathbb{V})$ in the remaining triangular regions do not show up in $\text{Dgm}(\mathbb{V}_{a,b,c})$.

2.9. **Finite approximations.** Away from the finite r -interior, the r -measure gives a decidedly limited view of the structure of a persistence module. For example:

- It is not possible to distinguish between the many nonisomorphic persistence modules \mathbb{V} for which $\mu_{\mathbb{V}}$ is infinite on every rectangle.
- If the persistence diagram of \mathbb{V} contains a sequence of points (p_n^*, q_n^*) with p_n, q_n converging to r from below and above, respectively, then there is no way to determine the multiplicity of (r^-, r^+) from the measure alone.

On the other hand, from $\mu_{\mathbb{V}}$ we do recover all information obtainable by restricting \mathbb{V} to finite subsets $\mathbf{T} \subset \mathbf{R}$. We may call this the ‘finitely observable’ part of \mathbb{V} . Specifically, for any finite index set

$$\mathbf{T} : a_1 < a_2 < \cdots < a_n$$

we have

$$(\dagger) \quad \langle [a_i, a_j] \mid \mathbb{V}_{\mathbf{T}} \rangle = \langle \circ_{a_{i-1}^-} \bullet_{a_i} \text{---} \bullet_{a_j} \text{---} \circ_{a_{j+1}} \mid \mathbb{V} \rangle = \mu_{\mathbb{V}}([a_{i-1}, a_i] \times [a_j, a_{j+1}])$$

(where a_0, a_{n+1} are interpreted as $-\infty, +\infty$ respectively).

This can be interpreted as a sort of ‘snapping principle’ (see [6, 7] for the origin of this term). The left-hand side of (\dagger) gives the multiplicity of the interval $[a_i, a_{j+1})$ in the traditional definition of the barcode for $\mathbb{V}_{\mathbf{T}}$. For us, this is $[a_i^-, a_{j+1}^-)$. The right-hand side of (\dagger) counts the decorated points of $\text{Dgm}(\mu_{\mathbb{V}})$ in the rectangle $[a_{i-1}, a_i] \times [a_j, a_{j+1}]$. See Figure 7.

There are some well known situations in which the entire structure of \mathbb{V} is finitely observable. For example, suppose

- X is a compact manifold and f is a Morse function; or

- X is a compact polyhedron and f is piecewise linear.

In each case there is a finite set of critical points

$$\mathbf{T} : a_1 < a_2 < \cdots < a_n$$

such that the inclusion of sublevelsets $X^s \subset X^t$ is a homotopy equivalence provided that $(s, t]$ does not meet any of the critical points. It follows that the sublevelset persistent homology $\mathbb{V} = \mathbf{H}(\mathbb{X}_{\text{sub}})$ is constant on each of the intervals

$$(-\infty, a_1), [a_1, a_2), [a_2, a_3), \dots, [a_{n-1}, a_n), [a_n, +\infty).$$

The structure of \mathbb{V} is therefore determined by $\mathbb{V}_{\mathbf{T}}$, it is of finite type, and $\mathbf{Dgm}(\mathbb{V})$ is localised at the points

$$(a_i^-, a_j^-) \quad \text{for } 1 \leq i < j \leq n,$$

and

$$(a_i^-, +\infty) \quad \text{for } 1 \leq i \leq n.$$

3. INTERLEAVING

As with any category, two persistence modules \mathbb{U}, \mathbb{V} are said to be isomorphic if there are maps

$$\begin{aligned} \Phi &\in \mathbf{Hom}(\mathbb{U}, \mathbb{V}), \\ \Psi &\in \mathbf{Hom}(\mathbb{V}, \mathbb{U}), \end{aligned}$$

such that

$$\begin{aligned} \Psi\Phi &= 1_{\mathbb{U}}, \\ \Phi\Psi &= 1_{\mathbb{V}}. \end{aligned}$$

This relation is too strong for situations where the data leading to the construction of a persistence module is obtained with some uncertainty or noise. The natural response is to consider a weaker relation, **δ -interleaving**, where $\delta \geq 0$ quantifies the uncertainty.

In this section, we define the interleaving relation and study its elementary properties. We prove the nontrivial result (from [6]) that if two persistence modules are δ -interleaved, then they are connected in the space of persistence modules by a path of length δ . This ‘interpolation lemma’ is a crucial step in the proof of the stability theorem in section 4.

In this section, all persistence modules are indexed by \mathbf{R} unless explicitly stated otherwise.

3.1. Shifted homomorphisms. The first step is to consider homomorphisms which shift the value of the persistent index. Let \mathbb{U}, \mathbb{V} be persistence modules over \mathbf{R} , and let δ be any real number. A **homomorphism of degree δ** is a collection Φ of linear maps

$$\phi_t : U_t \rightarrow V_{t+\delta}$$

for all $t \in \mathbf{R}$, such that the diagram

$$\begin{array}{ccc} U_s & \xrightarrow{u_s^t} & U_t \\ \phi_s \downarrow & & \downarrow \phi_t \\ V_{s+\delta} & \xrightarrow{v_{s+\delta}^{t+\delta}} & V_{t+\delta} \end{array}$$

commutes whenever $s \leq t$.

We write

$$\begin{aligned} \text{Hom}^\delta(\mathbb{U}, \mathbb{V}) &= \{\text{homomorphisms } \mathbb{U} \rightarrow \mathbb{V} \text{ of degree } \delta\}, \\ \text{End}^\delta(\mathbb{V}) &= \{\text{homomorphisms } \mathbb{V} \rightarrow \mathbb{V} \text{ of degree } \delta\}. \end{aligned}$$

Composition gives a map

$$\text{Hom}^{\delta_2}(\mathbb{V}, \mathbb{W}) \times \text{Hom}^{\delta_1}(\mathbb{U}, \mathbb{V}) \rightarrow \text{Hom}^{\delta_1+\delta_2}(\mathbb{U}, \mathbb{W}).$$

For $\delta \geq 0$, the most important degree- δ endomorphism is the shift map

$$1_{\mathbb{V}}^\delta \in \text{End}^\delta(\mathbb{V}),$$

which is the collection of maps $(v_t^{t+\delta})$ from the persistence structure on \mathbb{V} . If Φ is a homomorphism $\mathbb{U} \rightarrow \mathbb{V}$ of any degree, then by definition

$$\Phi 1_{\mathbb{U}}^\delta = 1_{\mathbb{V}}^\delta \Phi$$

for all $\delta \geq 0$.

3.2. Interleaving. Let $\delta \geq 0$. Two persistence modules \mathbb{U}, \mathbb{V} are said to be δ -**interleaved** if there are maps

$$\begin{aligned} \Phi &\in \text{Hom}^\delta(\mathbb{U}, \mathbb{V}) \\ \Psi &\in \text{Hom}^\delta(\mathbb{V}, \mathbb{U}) \end{aligned}$$

such that

$$\begin{aligned} \Psi \Phi &= 1_{\mathbb{U}}^{2\delta} \\ \Phi \Psi &= 1_{\mathbb{V}}^{2\delta} \end{aligned}$$

More expansively (with many more indices), this means that there are maps

$$\phi_t : U_t \rightarrow V_{t+\delta} \quad \text{and} \quad \psi_t : V_t \rightarrow U_{t+\delta}$$

defined for all t , such that the following diagrams

$$(3.1) \quad \begin{array}{ccc} U_s & \xrightarrow{u_s^t} & U_t \\ \phi_s \downarrow & & \downarrow \phi_t \\ V_{s+\delta} & \xrightarrow{v_{s+\delta}^{t+\delta}} & V_{t+\delta} \end{array} \quad \begin{array}{ccc} U_{s-\delta} & \xrightarrow{u_{s-\delta}^{s+\delta}} & U_{s+\delta} \\ \phi_{s-\delta} \searrow & & \nearrow \psi_s \\ & V_s & \end{array}$$

$$\begin{array}{ccc} V_s & \xrightarrow{v_s^t} & V_t \\ \psi_s \downarrow & & \downarrow \psi_t \\ U_{s+\delta} & \xrightarrow{u_{s+\delta}^{t+\delta}} & U_{t+\delta} \end{array} \quad \begin{array}{ccc} V_{s-\delta} & \xrightarrow{v_{s-\delta}^{s+\delta}} & V_{s+\delta} \\ \psi_{s-\delta} \searrow & & \nearrow \phi_s \\ & U_s & \end{array}$$

commute for all eligible parameter values; that is, for all $s < t$. Where possible, we will be concise rather than expansive.

Here is the classic example of a pair of interleaved persistence modules.

Example 3.2. Let X be a topological space and let $f, g : X \rightarrow \mathbf{R}$. Suppose $\|f - g\|_\infty < \delta$. Then the persistence modules $H(\mathbb{X}_{\text{sub}}^f)$, $H(\mathbb{X}_{\text{sub}}^g)$ are δ -interleaved. Indeed, there are inclusions

$$\begin{aligned} (X, f)^t &\subseteq (X, g)^{t+\delta} \\ (X, g)^t &\subseteq (X, f)^{t+\delta} \end{aligned}$$

for all t , which induce maps

$$\begin{aligned} \Phi &: H(\mathbb{X}_{\text{sub}}^f) \rightarrow H(\mathbb{X}_{\text{sub}}^g) \\ \Psi &: H(\mathbb{X}_{\text{sub}}^g) \rightarrow H(\mathbb{X}_{\text{sub}}^f) \end{aligned}$$

of degree δ . Since all the maps are induced functorially from inclusion maps, the interleaving relations are automatically satisfied.

This is the situation for which the stability theorem of Cohen-Steiner, Edelsbrunner and Harer [10] was originally stated: if two functions f, g are close then the diagrams for their sublevelset persistent homology are close. Subsequently, stability was formulated as a theorem about the diagrams of interleaved persistence modules [6, 7]. In the present paper, we come to view stability as a theorem about r-measures.

3.3. Interleaving (continued). An interleaving between two persistence modules can itself be thought of as a persistence module over a certain partially ordered set (poset). We develop this idea now.

Consider the standard partial order on the plane:

$$(p_1, q_1) \leq (p_2, q_2) \quad \Leftrightarrow \quad p_1 \leq p_2 \text{ and } q_1 \leq q_2.$$

For any real number x , define the shifted diagonal

$$\Delta_x = \{(p, q) \mid q - p = 2x\} \subset \mathbf{R}^2.$$

As a poset, this is isomorphic to the real line. Specifically, we identify $t \in \mathbf{R}$ with the point $(t - x, t + x) \in \Delta_x$. Through this, we get a canonical identification between persistence modules over \mathbf{R} and persistence modules over Δ_x .

Proposition 3.3. *Let x, y be real numbers. Persistence modules \mathbb{U}, \mathbb{V} are $|y - x|$ -interleaved if and only if there is a persistence module \mathbb{W} over $\Delta_x \cup \Delta_y$ such that $\mathbb{W}|_{\Delta_x} = \mathbb{U}$ and $\mathbb{W}|_{\Delta_y} = \mathbb{V}$.*

Proof. We may assume $x < y$.

We claim that (i) the extra information carried by $(y - x)$ -interleaving maps Φ, Ψ is equivalent to (ii) the extra information carried by \mathbb{W} . Let us describe both, more carefully:

(i) In addition to \mathbb{U}, \mathbb{V} we have a system of maps $\Phi = (\phi_t)$, where

$$\phi_t : U_t \rightarrow V_{t+y-x},$$

and a system of maps $\Psi = (\psi_t)$, where

$$\psi_t : V_t \rightarrow U_{t+y-x}.$$

These are constrained by the relations (for all $\eta \geq 0$).

$$(\ddagger) \quad \Phi 1_{\mathbb{U}}^\eta = 1_{\mathbb{V}}^\eta \Phi, \quad \Psi 1_{\mathbb{V}}^\eta = 1_{\mathbb{U}}^\eta \Psi, \quad \Psi \Phi = 1_{\mathbb{U}}^{2y-2x}, \quad \Phi \Psi = 1_{\mathbb{V}}^{2y-2x}.$$

There are no other constraints.

(ii) In addition to \mathbb{U}, \mathbb{V} the persistence module \mathbb{W} carries maps between the two components Δ_x, Δ_y . These maps are constrained by the composition law

$$w_R^T = w_S^T \circ w_R^S$$

for all $R, S, T \in \Delta_x \cup \Delta_y$ with $R \leq S \leq T$.

First, observe that we recover the maps ϕ_t, ψ_t as vertical maps from Δ_x to Δ_y , and horizontal maps from Δ_y to Δ_x , respectively (see Figure 8):

$$\begin{aligned} U_t &= W_{t-x, t+x} \rightarrow W_{t-x, t+2y-x} = V_{t+y-x} \\ V_t &= W_{t-y, t+y} \rightarrow W_{t+y-2x, t+y} = U_{t+y-x} \end{aligned}$$

Next, observe that the composition law implies all of the relations (\ddagger) .

Finally, there is no additional information in \mathbb{W} , beyond the interleaving maps and relations. Indeed, all remaining maps w_S^T , where $S \leq T$, can all be factored in the form:

$$\begin{aligned} w_S^T &= v_{s+y-x}^t \circ \phi_s && \text{if } S \in \Delta_x \text{ and } T \in \Delta_y, \\ w_S^T &= u_{s+y-x}^t \circ \psi_s && \text{if } S \in \Delta_y \text{ and } T \in \Delta_x. \end{aligned}$$

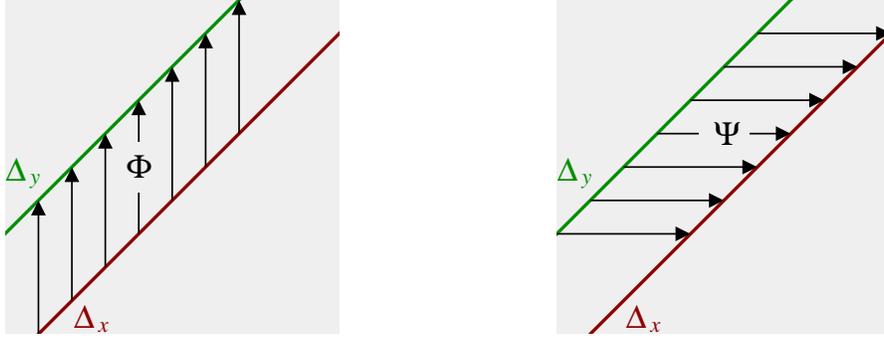


FIGURE 8. The maps Φ, Ψ recovered from the module \mathbb{W} over $\Delta_x \cup \Delta_y$.

Thus each map in \mathbb{W} is an instance of one of

$$\begin{aligned}
 1_{\mathbb{U}}^\eta & \quad \text{from } \Delta_x \text{ to } \Delta_x, \\
 1_{\mathbb{V}}^\eta & \quad \text{from } \Delta_y \text{ to } \Delta_y, \\
 1_{\mathbb{V}}^\eta \Phi & \quad \text{from } \Delta_x \text{ to } \Delta_y, \\
 1_{\mathbb{U}}^\eta \Psi & \quad \text{from } \Delta_y \text{ to } \Delta_x.
 \end{aligned}$$

It is a simple matter to verify that the composition law is satisfied for each composable pair of maps. For instance

$$(1_{\mathbb{V}}^\eta \Phi)(1_{\mathbb{U}}^\zeta \Psi) = 1_{\mathbb{V}}^\eta \Phi 1_{\mathbb{U}}^\zeta \Psi = 1_{\mathbb{V}}^\eta 1_{\mathbb{V}}^\zeta \Phi \Psi = 1_{\mathbb{V}}^{\eta+\zeta} 1_{\mathbb{V}}^{2y-2x} = 1_{\mathbb{V}}^{\eta+\zeta+2y-2x}.$$

This can be done using only the known relations, so there are no further constraints on the w_S^T . \square

Remark. This characterisation makes it clear (or, depends on the fact) that all composable combinations of the maps u, v, ϕ, ψ from a given domain to a given codomain must be equal: indeed, they must agree with the appropriate map w_S^T of \mathbb{W} .

3.4. The interpolation lemma. In this section we give a proof of the following theorem, which first appeared in [6]. The result strikes us as somewhat surprising.

Lemma 3.4 (Interpolation lemma). *Suppose \mathbb{U}, \mathbb{V} are a δ -interleaved pair of persistence modules. Then there exists a 1-parameter family of persistence modules $(\mathbb{U}_x \mid x \in [0, \delta])$ such that $\mathbb{U}_0, \mathbb{U}_\delta$ are equal to \mathbb{U}, \mathbb{V} respectively, and $\mathbb{U}_x, \mathbb{U}_y$ are $|y-x|$ -interleaved for all $x, y \in [0, \delta]$.*

We can make a sharper statement about what is proved. Given a specific pair of interleaving maps

$$\begin{aligned}
 \Phi & \in \text{Hom}^\delta(\mathbb{U}, \mathbb{V}) \\
 \Psi & \in \text{Hom}^\delta(\mathbb{V}, \mathbb{U})
 \end{aligned}$$

the construction explicitly provides, for each $x < y$, a pair of interleaving maps

$$\begin{aligned}\Phi_x^y &\in \text{Hom}^{y-x}(\mathbb{U}_x, \mathbb{U}_y) \\ \Psi_y^x &\in \text{Hom}^{y-x}(\mathbb{U}_y, \mathbb{U}_x)\end{aligned}$$

such that $\Phi_0^\delta, \Psi_\delta^0$ are equal to Φ, Ψ respectively, and moreover

$$\begin{aligned}\Phi_x^z &= \Phi_y^z \Phi_x^y \\ \Psi_z^x &= \Psi_y^x \Psi_z^y\end{aligned}$$

for all $x < y < z$.

In view of Proposition 3.3, this sharp form of the interpolation lemma can be restated as follows. (We have also replaced the interval $[0, \delta]$ with the more general $[x_0, x_1]$.)

Theorem 3.5. *Any persistence module \mathbb{W} over $\Delta_0 \cup \Delta_\delta$ extends to a persistence module $\overline{\mathbb{W}}$ over the diagonal strip*

$$\Delta_{[0,\delta]} = \{(p, q) \mid 0 \leq q - p \leq 2\delta\} \subset \mathbf{R}^2.$$

Remark. The extension is by no means unique.

Let us clarify how Theorem 3.5 implies Lemma 3.4. If \mathbb{U}, \mathbb{V} are δ -interleaved, then there exists a persistence module \mathbb{W} over $\Delta_0 \cup \Delta_\delta$ such that $\mathbb{W}|_{\Delta_0} = \mathbb{U}$ and $\mathbb{W}|_{\Delta_\delta} = \mathbb{V}$. By Theorem 3.5, this extends to $\overline{\mathbb{W}}$ over the strip $\Delta_{[0,\delta]}$. If we define a 1-parameter family $\mathbb{U}_x = \overline{\mathbb{W}}|_{\Delta_x}$, then $\mathbb{U}_x, \mathbb{U}_y$ are $|x - y|$ -interleaved for all $x, y \in [0, \delta]$.

The equivalence between Theorem 3.5 and the sharp form of Lemma 3.4 is left to the reader.

Proof of Theorem 3.5. In order to express the proof more symmetrically, it is convenient to replace the interval $[0, \delta]$ by the interval $[-1, 1]$. This can be done by a rescaling and translation of the plane. Accordingly, suppose we are given a persistence module \mathbb{W} over $\Delta_{-1} \cup \Delta_1$.

Our strategy is to construct two persistence modules over the strip $\Delta_{[-1,1]}$ and a module map between them. The image of this map is also a persistence module over the strip, and will be the required extension.

By Proposition 3.3, \mathbb{W} provides $\mathbb{U} = \mathbb{W}|_{\Delta_{-1}}$ and $\mathbb{V} = \mathbb{W}|_{\Delta_1}$, which we can view as persistence modules over \mathbf{R} ,

$$\begin{aligned}U_t &= W_{(t+1, t-1)}, \\ V_t &= W_{(t-1, t+1)},\end{aligned}$$

as well as interleaving maps $\Phi \in \text{Hom}^2(\mathbb{U}, \mathbb{V})$ and $\Psi \in \text{Hom}^2(\mathbb{V}, \mathbb{U})$ of degree 2.

From \mathbb{U}, \mathbb{V} we construct four persistence modules over \mathbf{R}^2 :

$$\begin{array}{llll}
\mathbb{A} = \mathbb{U}[p-1] & \text{defined by} & A_{(p,q)} = U_{p-1} & \text{and} & a_{(p,q)}^{(r,s)} = u_{p-1}^{r-1} \\
\mathbb{B} = \mathbb{V}[q-1] & \text{defined by} & B_{(p,q)} = V_{q-1} & \text{and} & b_{(p,q)}^{(r,s)} = v_{q-1}^{s-1} \\
\mathbb{C} = \mathbb{U}[q+1] & \text{defined by} & C_{(p,q)} = U_{q+1} & \text{and} & c_{(p,q)}^{(r,s)} = u_{q+1}^{s+1} \\
\mathbb{D} = \mathbb{V}[p+1] & \text{defined by} & D_{(p,q)} = V_{p+1} & \text{and} & d_{(p,q)}^{(r,s)} = v_{p+1}^{r+1}
\end{array}$$

Next, we construct four module maps.

$$\begin{array}{llll}
1_{\mathbb{U}} : \mathbb{A} \rightarrow \mathbb{C} & \text{defined at } (p, q) \text{ to be} & u_{p-1}^{q+1} : U_{p-1} \rightarrow U_{q+1} \\
\Phi : \mathbb{A} \rightarrow \mathbb{D} & \text{defined at } (p, q) \text{ to be} & \phi_{p-1} : U_{p-1} \rightarrow V_{p+1} \\
\Psi : \mathbb{B} \rightarrow \mathbb{C} & \text{defined at } (p, q) \text{ to be} & \psi_{q-1} : V_{q-1} \rightarrow U_{q+1} \\
1_{\mathbb{V}} : \mathbb{B} \rightarrow \mathbb{D} & \text{defined at } (p, q) \text{ to be} & v_{q-1}^{p+1} : V_{q-1} \rightarrow V_{p+1}
\end{array}$$

The maps Φ, Ψ are defined over the whole plane, whereas the map $1_{\mathbb{U}}$ is defined only where $p-1 \leq q+1$, and the map $1_{\mathbb{V}}$ is defined only where $q-1 \leq p+1$. It follows that all four maps are defined in the region where

$$-2 \leq q - p \leq 2,$$

which is precisely the strip $\Delta_{[-1,1]}$. Henceforth, we restrict to that strip.

It is easy to confirm that these are module maps. The required commutation relations involve composable combinations of the maps u, v, ϕ, ψ , which always agree by the remark following Proposition 3.3.

Define $\Omega \in \text{Hom}(\mathbb{A} \oplus \mathbb{B}, \mathbb{C} \oplus \mathbb{D})$ by the 2-by-2 matrix

$$\begin{bmatrix} 1_{\mathbb{U}} & \Psi \\ \Phi & 1_{\mathbb{V}} \end{bmatrix}$$

of module maps. Our claim is that $\overline{\mathbb{W}} = \text{im}(\Omega)$ is the required extension. An equivalent statement is that $\overline{\mathbb{W}} = \text{coim}(\Omega)$ is the required extension (using the coimage rather than the image).

Step 1: $\overline{\mathbb{W}}|_{\Delta_{-1}}$ is isomorphic to \mathbb{U} .

Proof. On $\Delta_{-1} = \{(t+1, t-1)\}$, we have

$$\begin{aligned}
(\mathbb{A} \oplus \mathbb{B})_t &= U_t \oplus V_{t-2}, \\
(\mathbb{C} \oplus \mathbb{D})_t &= U_t \oplus V_{t+2},
\end{aligned}$$

and the homomorphism $\Omega|_{\Delta_{-1}}$ is given by

$$\omega_t = \begin{bmatrix} u_t^t & \psi_{t-2} \\ \phi_t & v_{t-2}^{t+2} \end{bmatrix}.$$

The key point is that this factorises (in a block-matrix sense it ‘has rank 1’). The factorisation can be written in matrix form

$$\omega_t = \begin{bmatrix} u_t^t \\ \phi_t \end{bmatrix} \begin{bmatrix} u_t^t & \psi_{t-2} \end{bmatrix}$$

or as a diagram of \mathbf{R} -module maps

$$\mathbb{U} \oplus \mathbb{V}[t-2] \xrightarrow{\Omega_1} \mathbb{U} \xrightarrow{\Omega_2} \mathbb{U} \oplus \mathbb{V}[t+2]$$

where

$$\Omega_1(\mathbf{u} \oplus \mathbf{v}) = \mathbf{u} + \Psi(\mathbf{v}) \quad \text{and} \quad \Omega_2(\mathbf{u}) = \mathbf{u} \oplus \Phi(\mathbf{u}).$$

Since Ω_1 is surjective and Ω_2 is injective (at all indices t), it follows that

$$\text{coim}(\Omega) \cong \mathbb{U} \cong \text{im}(\Omega)$$

where the isomorphisms are given by Ω_1, Ω_2 respectively. □

Step 2. $\overline{\mathbb{W}}|_{\Delta_1}$ is isomorphic to \mathbb{V} .

Proof. On $\Delta_1 = \{(t-1, t+1)\}$, we have

$$\begin{aligned} (\mathbb{A} \oplus \mathbb{B})_t &= U_{t-2} \oplus V_t, \\ (\mathbb{C} \oplus \mathbb{D})_t &= U_{t+2} \oplus V_t, \end{aligned}$$

and

$$\omega_t = \begin{bmatrix} u_{t-2}^{t+2} & \psi_t \\ \phi_{t-2} & v_t^t \end{bmatrix}.$$

Again Ω factorises. The factorisation can be written in matrix form

$$\omega_t = \begin{bmatrix} \psi_t \\ v_t^t \end{bmatrix} \begin{bmatrix} \phi_{t-2} & v_t^t \end{bmatrix}$$

or as a diagram of \mathbf{R} -module maps

$$\mathbb{U}[t-2] \oplus \mathbb{V} \xrightarrow{\Omega_3} \mathbb{V} \xrightarrow{\Omega_4} \mathbb{U}[t+2] \oplus \mathbb{V}$$

where

$$\Omega_3(\mathbf{u} \oplus \mathbf{v}) = \Phi(\mathbf{u}) + \mathbf{v} \quad \text{and} \quad \Omega_4(\mathbf{v}) = \Psi(\mathbf{v}) \oplus \mathbf{v}.$$

Since Ω_3 is surjective and Ω_4 is injective, we have

$$\text{coim}(\Omega) \cong \mathbb{V} \cong \text{im}(\Omega)$$

where the isomorphisms are given by Ω_3, Ω_4 respectively. □

Step 3. Under the isomorphisms above, the cross maps of $\overline{\mathbb{W}}$ are precisely Φ and Ψ .

Proof. The cross maps for $\overline{\mathbb{W}}$ are induced by the cross maps for $\mathbb{A} \oplus \mathbb{B}$ and also by the cross maps for $\mathbb{C} \oplus \mathbb{D}$. (The result must be the same, because Ω is a module homomorphism.)

The following diagram shows the vertical maps for $\mathbb{A} \oplus \mathbb{B}$ (on the left) and $\mathbb{C} \oplus \mathbb{D}$ (on the right), as well as the factorisations of Steps 1 and 2.

$$\begin{array}{ccccc}
 \mathbb{U} \oplus \mathbb{V}[t-2] & \xrightarrow{\Omega_1} & \mathbb{U} & \xrightarrow{\Omega_2} & \mathbb{U} \oplus \mathbb{V}[t+2] \\
 \downarrow 1_{\mathbb{U}} \oplus 1_{\mathbb{V}}^4 & & \downarrow \Phi & & \downarrow 1_{\mathbb{U}}^4 \oplus 1_{\mathbb{V}} \\
 \mathbb{U} \oplus \mathbb{V}[t+2] & \xrightarrow{\Omega_3} & \mathbb{V}[t+2] & \xrightarrow{\Omega_4} & \mathbb{U}[t+4] \oplus \mathbb{V}[t+2]
 \end{array}$$

The diagram, with Φ in the middle, is easily verified to commute using the explicit forms of the maps Ω_i . It follows that Φ is the induced vertical cross map for $\overline{\mathbb{W}}$.

A similar argument using the diagram

$$\begin{array}{ccccc}
 \mathbb{U}[t-2] \oplus \mathbb{V} & \xrightarrow{\Omega_3} & \mathbb{V} & \xrightarrow{\Omega_4} & \mathbb{U}[t+2] \oplus \mathbb{V} \\
 \downarrow 1_{\mathbb{U}}^4 \oplus 1_{\mathbb{V}} & & \downarrow \Psi & & \downarrow 1_{\mathbb{U}} \oplus 1_{\mathbb{V}}^4 \\
 \mathbb{U}[t+2] \oplus \mathbb{V} & \xrightarrow{\Omega_1} & \mathbb{U}[t+2] & \xrightarrow{\Omega_2} & \mathbb{U}[t+2] \oplus \mathbb{V}[t+4]
 \end{array}$$

shows that Ψ is the induced horizontal cross map. □

This completes the proof of Theorem 3.5. □

We note that Step 3 isn't necessary to deduce the interpolation lemma (3.4), because it is sufficient to demonstrate that \mathbb{U}, \mathbb{V} belong to some 1-parameter family of persistence modules. It is important that the $|x-y|$ -interleaving between each pair $\mathbb{U}_x, \mathbb{U}_y$ exists, but it doesn't matter what it is.

One can interpret the stronger Theorem 3.5 as a functor extension theorem. If we regard the posets $\Delta_{x_0} \cup \Delta_{x_1}$ and $\Delta_{[x_0, x_1]}$ as categories, then persistence modules over these posets are the same as functors to the category of vector spaces. (See the remark near the beginning of section 1.1.) The theorem asserts the existence of an extension $\overline{\mathbb{W}}$

$$\begin{array}{ccc}
 & \Delta_{[x_0, x_1]} & \\
 & \uparrow & \searrow \overline{\mathbb{W}} \\
 \Delta_{x_0} \cup \Delta_{x_1} & \xrightarrow{\mathbb{W}} & \mathbf{Vect}
 \end{array}$$

for any functor \mathbb{W} .

4. THE ISOMETRY THEOREM

In this section we discuss the metric relationship between persistence modules and their persistence diagrams. As in section 3, all persistence modules are indexed by \mathbf{R} unless explicitly stated otherwise.

The principal result is the famous stability theorem of Cohen–Steiner, Edelsbrunner and Harer [10], in the generality established by [7]. The main difference here is that we emphasise persistence measures in the proof, and indeed we derive the standard theorem from a more general statement about persistence measures. The basic structure of the proof is essentially the same as in [10].

The secondary result is the converse inequality, which together with the stability theorem implies that the space of q-tame persistence modules is isometric with the space of locally finite persistence diagrams. This isometry theorem appeared originally in the work of Lesnick [16] for modules which satisfy $\dim(V_t) < \infty$ for all t , and independently in the work of Bubenik and Scott [3] for modules of finite type.

4.1. The interleaving distance. In this section we define the interleaving distance between persistence modules. This was introduced in [7].

The first observation is that if \mathbb{U} and \mathbb{V} are δ -interleaved, then they are $(\delta + \epsilon)$ -interleaved for every $\epsilon > 0$. Indeed, the maps

$$\begin{aligned}\Phi' &= \Phi 1_{\mathbb{U}}^{\epsilon} = 1_{\mathbb{V}}^{\epsilon} \Phi \\ \Psi' &= \Psi 1_{\mathbb{V}}^{\epsilon} = 1_{\mathbb{U}}^{\epsilon} \Psi\end{aligned}$$

provide the required interleaving.

The challenge, then, if two persistence modules are interleaved, is to make the interleaving parameter as small as possible. The minimum is not necessarily attained, so we introduce some additional terminology.

We say that two persistence modules \mathbb{U}, \mathbb{V} are δ -**interleaved** if they are $(\delta + \epsilon)$ -interleaved for all $\epsilon > 0$. As we shall see, this does not imply that \mathbb{U}, \mathbb{V} are δ -interleaved.

Example 4.1. Two persistence modules are 0-interleaved if and only if they are isomorphic.

Example 4.2. A persistence module \mathbb{V} is **ephemeral** if $v_s^t = 0$ for all $s < t$. The spaces V_t can be completely arbitrary. Let \mathbb{U} and \mathbb{V} be two non-isomorphic ephemeral modules. Then \mathbb{U}, \mathbb{V} are 0^+ -interleaved but not 0-interleaved. Indeed, $1_{\mathbb{U}}^{2\epsilon} = 0$ and $1_{\mathbb{V}}^{2\epsilon} = 0$ for all $\epsilon > 0$, so the zero maps

$$\begin{aligned}\Phi &= 0 \in \text{Hom}^{\epsilon}(\mathbb{U}, \mathbb{V}) \\ \Psi &= 0 \in \text{Hom}^{\epsilon}(\mathbb{V}, \mathbb{U})\end{aligned}$$

give an ϵ -interleaving.

The **interleaving distance** between two persistence modules is defined as follows:

$$\begin{aligned} d_i(\mathbb{U}, \mathbb{V}) &= \inf\{\delta \mid \mathbb{U}, \mathbb{V} \text{ are } \delta\text{-interleaved}\} \\ &= \min\{\delta \mid \mathbb{U}, \mathbb{V} \text{ are } \delta^+\text{-interleaved}\} \end{aligned}$$

If there is no δ -interleaving between \mathbb{U}, \mathbb{V} for any value of δ , then $d_i(\mathbb{U}, \mathbb{V}) = \infty$.

Proposition 4.3. *The interleaving distance satisfies the triangle inequality:*

$$d_i(\mathbb{U}, \mathbb{W}) \leq d_i(\mathbb{U}, \mathbb{V}) + d_i(\mathbb{V}, \mathbb{W})$$

for any three persistence modules $\mathbb{U}, \mathbb{V}, \mathbb{W}$.

Proof. Given a δ_1 -interleaving between \mathbb{U}, \mathbb{V} and a δ_2 -interleaving between \mathbb{V}, \mathbb{W} one can construct a $\delta = (\delta_1 + \delta_2)$ -interleaving between \mathbb{U}, \mathbb{W} by composing the interleaving maps:

$$\begin{array}{ccc} \mathbb{U} & \xrightarrow{\Phi_1} \mathbb{V} & \xrightarrow{\Phi_2} \mathbb{W} \\ \mathbb{U} & \xleftarrow{\Psi_1} \mathbb{V} & \xleftarrow{\Psi_2} \mathbb{W} \end{array}$$

One easily verifies that $\Phi = \Phi_2\Phi_1$ and $\Psi = \Psi_1\Psi_2$ are interleaving maps. Explicitly:

$$\begin{aligned} \Psi\Phi &= \Psi_1\Psi_2\Phi_2\Phi_1 = \Psi_1 1_{\mathbb{V}}^{2\delta_2}\Phi_1 = \Psi_1\Phi_1 1_{\mathbb{U}}^{2\delta_2} = 1_{\mathbb{U}}^{2\delta_1} 1_{\mathbb{U}}^{2\delta_2} = 1_{\mathbb{U}}^{2\delta} \\ \Phi\Psi &= \Phi_2\Phi_1\Psi_1\Psi_2 = \Phi_2 1_{\mathbb{V}}^{2\delta_1}\Psi_2 = \Phi_2\Psi_2 1_{\mathbb{W}}^{2\delta_1} = 1_{\mathbb{W}}^{2\delta_2} 1_{\mathbb{W}}^{2\delta_1} = 1_{\mathbb{W}}^{2\delta} \end{aligned}$$

Now take the infimum over δ_1, δ_2 . □

The proposition tells us that d_i is a pseudometric. It is not a true metric because $d_i(\mathbb{U}, \mathbb{V}) = 0$ does not imply $\mathbb{U} \cong \mathbb{V}$, as we saw above. In fact, two q -tame persistence modules have interleaving distance 0 if and only if their undecorated persistence diagrams are the same. This is a consequence of the isometry theorem.

Here is the simplest instance. The direct proof is left to the reader (or see Proposition 4.6).

Example 4.4. The four interval modules

$$\mathbb{I}[p, q], \mathbb{I}[p, q], \mathbb{I}(p, q), \mathbb{I}(p, q)$$

are 0^+ -interleaved but not isomorphic.

The following property of interleaving distance will be useful later.

Proposition 4.5. *Let $\mathbb{U}_1, \mathbb{U}_2, \mathbb{V}_1, \mathbb{V}_2$ be persistence modules. Then*

$$d_i(\mathbb{U}_1 \oplus \mathbb{U}_2, \mathbb{V}_1 \oplus \mathbb{V}_2) \leq \max(d_i(\mathbb{U}_1, \mathbb{V}_1), d_i(\mathbb{U}_2, \mathbb{V}_2))$$

More generally, let $(\mathbb{U}_\ell \mid \ell \in L)$ and $(\mathbb{V}_\ell \mid \ell \in L)$ be families of persistence modules indexed by the same set L , and let

$$\mathbb{U} = \bigoplus_{\ell \in L} \mathbb{U}_\ell, \quad \mathbb{V} = \bigoplus_{\ell \in L} \mathbb{V}_\ell.$$

Then

$$d_i(\mathbb{U}, \mathbb{V}) \leq \sup(d_i(\mathbb{U}_\ell, \mathbb{V}_\ell) \mid \ell \in L).$$

Proof. Given δ -interleavings Φ_ℓ, Ψ_ℓ for each pair $\mathbb{U}_\ell, \mathbb{V}_\ell$, the direct sum maps $\Phi = \bigoplus \Phi_\ell, \Psi = \bigoplus \Psi_\ell$ constitute a δ -interleaving of \mathbb{U}, \mathbb{V} . Thus any upper bound on the $d_i(\mathbb{U}_\ell, \mathbb{V}_\ell)$ is an upper bound for $d_i(\mathbb{U}, \mathbb{V})$. In particular, this is true for the least upper bound, or sup. \square

4.2. The bottleneck distance. Now we define the metric on the other side of the isometry theorem, namely the bottleneck distance between undecorated persistence diagrams. For a q-tame persistence module \mathbb{V} , every rectangle not touching the diagonal has finite $\mu_{\mathbb{V}}$ -measure. This implies that the undecorated diagram

$$\text{dgm}(\mathbb{V}) = \text{dgm}(\mu_{\mathbb{V}})$$

is a multiset in the extended open half-plane

$$\overline{\mathcal{H}}^\circ = \{(p, q) \mid -\infty \leq p < q \leq +\infty\}.$$

In order to define the bottleneck distance, we need to specify the distance between any pair of points in $\overline{\mathcal{H}}^\circ$, as well as the distance between any point and the diagonal (the boundary of the half-plane). These distance functions are not arbitrary; they are defined as they are because of the interleaving properties of interval modules.

(point to point): The first idea is that two undecorated diagrams are close if there is a bijection between them which doesn't move any point too far. We use the ℓ^∞ -metric in the plane:

$$d^\infty((p, q), (r, s)) = \max(|p - r|, |q - s|)$$

Points at infinity are compared in the expected way:

$$\begin{aligned} d^\infty((-\infty, q), (-\infty, s)) &= |q - s|, \\ d^\infty((p, +\infty), (r, +\infty)) &= |p - r|, \end{aligned}$$

and

$$d^\infty((-\infty, +\infty), (-\infty, +\infty)) = 0.$$

Distances between points in different strata (e.g. between (p, q) and $(-\infty, s)$) are infinite.

The next result gives a relationship between the ℓ^∞ -metric and the interleaving of interval modules.

Proposition 4.6. *Let (p^*, q^*) and (r^*, s^*) be intervals (possibly infinite), and let*

$$\mathbb{U} = \mathbb{I}(p^*, q^*) \quad \text{and} \quad \mathbb{V} = \mathbb{I}(r^*, s^*)$$

be the corresponding interval modules. Then

$$d_i(\mathbb{U}, \mathbb{V}) \leq d^\infty((p, q), (r, s)).$$

The proof is postponed to the end of the section. We remark that equality holds provided that the intervals overlap sufficiently (the closure of each interval must meet the midpoint of the other), so the proposition is tight in that sense.

(point to diagonal): The next idea is that points which are close to the diagonal are allowed to be swallowed up by the diagonal. Again we use the ℓ^∞ -metric:

$$d^\infty((p, q), \Delta) = \frac{1}{2}(q - p)$$

Again this is related to the behaviour of interval modules.

Proposition 4.7. *Let (p^*, q^*) be an interval, let*

$$\mathbb{U} = \mathbb{I}(p^*, q^*),$$

be the corresponding interval module, and let 0 denote the zero persistence module. Then

$$d_i(\mathbb{U}, 0) = \frac{1}{2}(q - p).$$

(This is infinite if the interval is infinite.)

Proof. Let $\delta \geq 0$. When is there a δ -interleaving? The interleaving maps must be zero (no other maps exist to or from the module 0), so the only condition that needs checking is $\Psi\Phi = 1_{\mathbb{U}}^{2\delta}$, which is really $0 = 1_{\mathbb{U}}^{2\delta}$. This holds when $\delta > \frac{1}{2}(q - p)$ and fails when $\delta < \frac{1}{2}(q - p)$. \square

We now use these two concepts to define the bottleneck distance between two multisets \mathbf{A}, \mathbf{B} in the extended half-plane.

It is easier to work with sets rather than multisets. One way to do this is to attach labels to distinguish multiple instances of each repeated point. For instance, α with multiplicity k becomes $\alpha_1, \dots, \alpha_k$. Henceforth, we will do this without much comment.

A **partial matching** between \mathbf{A} and \mathbf{B} is a collection of pairs

$$\mathbf{M} \subset \mathbf{A} \times \mathbf{B}$$

such that:

- for every $\alpha \in \mathbf{A}$ there is at most one $\beta \in \mathbf{B}$ such that $(\alpha, \beta) \in \mathbf{M}$;
- for every $\beta \in \mathbf{B}$ there is at most one $\alpha \in \mathbf{A}$ such that $(\alpha, \beta) \in \mathbf{M}$.

We say that a partial matching \mathbf{M} is a **δ -matching** if all of the following are true:

- if $(\alpha, \beta) \in \mathbf{M}$ then $d^\infty(\alpha, \beta) \leq \delta$;
- if $\alpha \in \mathbf{A}$ is unmatched then $d^\infty(\alpha, \Delta) \leq \delta$;
- if $\beta \in \mathbf{B}$ is unmatched then $d^\infty(\beta, \Delta) \leq \delta$.

The **bottleneck distance** between two multisets \mathbf{A}, \mathbf{B} in the extended half-plane is

$$d_b(\mathbf{A}, \mathbf{B}) = \inf(\delta \mid \text{there exists a } \delta\text{-matching between } \mathbf{A} \text{ and } \mathbf{B}).$$

In section 4.3, we will show that ‘inf’ can be replaced by ‘min’ if \mathbf{A}, \mathbf{B} are locally finite.

Remark. In order for $d_b(\mathbf{A}, \mathbf{B}) < \infty$, it is necessary that the cardinalities of \mathbf{A}, \mathbf{B} agree over each of the three strata at infinity:

$$\begin{aligned} \text{card}(\mathbf{A}|_{\{-\infty\} \times \mathbf{R}}) &= \text{card}(\mathbf{B}|_{\{-\infty\} \times \mathbf{R}}) \\ \text{card}(\mathbf{A}|_{\mathbf{R} \times \{+\infty\}}) &= \text{card}(\mathbf{B}|_{\mathbf{R} \times \{+\infty\}}) \\ \text{card}(\mathbf{A}|_{\{-\infty\} \times \{+\infty\}}) &= \text{card}(\mathbf{B}|_{\{-\infty\} \times \{+\infty\}}) \end{aligned}$$

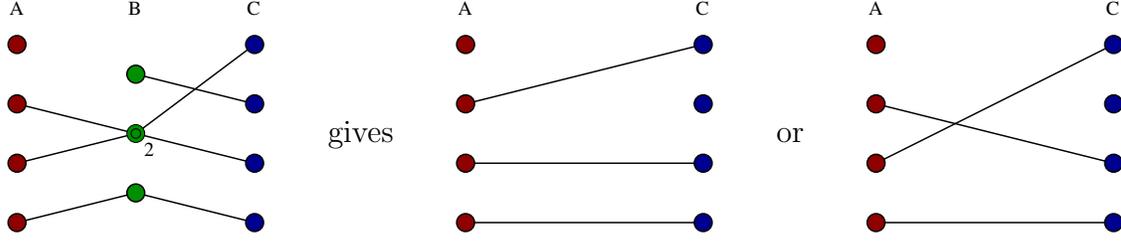


FIGURE 9. The partial matchings between A, B and B, C (left) compose to give a partial matching between A, C in two different ways (middle, right). The non-uniqueness arises from the point of multiplicity 2 in B .

Indeed, these points have infinite distance from the diagonal and from points in the other strata, and therefore they must be bijectively matched within each stratum.

Proposition 4.8. *The bottleneck distance satisfies the triangle inequality:*

$$d_b(A, C) \leq d_b(A, B) + d_b(B, C)$$

for any three multisets A, B, C .

Proof. Suppose M_1 is a δ_1 -matching between A, B , and M_2 is a δ_2 -matching between B, C . Let $\delta = \delta_1 + \delta_2$. We must show that there is a δ -matching between A, C .

Define the composition of M_1, M_2 to be

$$M = \{(\alpha, \gamma) \mid \text{there exists } \beta \in B \text{ such that } (\alpha, \beta) \in M_1 \text{ and } (\beta, \gamma) \in M_2\}.$$

We verify that M is the required δ -matching:

- If $(\alpha, \gamma) \in M$ then

$$d^\infty(\alpha, \gamma) \leq d^\infty(\alpha, \beta) + d^\infty(\beta, \gamma) \leq \delta_1 + \delta_2 = \delta$$

where $\beta \in B$ is the point linking α to γ .

- If α is unmatched in M then there are two possibilities. Either α is unmatched in M_1 , in which case

$$d^\infty(\alpha, \Delta) \leq \delta_1 \leq \delta.$$

Or α is matched in M_1 , let's say $(\alpha, \beta) \in M_1$. Then β must be unmatched in M_2 , so

$$d^\infty(\alpha, \Delta) \leq d^\infty(\alpha, \beta) + d^\infty(\beta, \Delta) \leq \delta_1 + \delta_2 = \delta.$$

- If γ is unmatched in M , then a similar argument shows that

$$d^\infty(\gamma, \Delta) \leq \delta.$$

This completes the proof. □

Remark. Because A, B, C are in truth multisets rather than sets, the composition operation between matchings is not uniquely defined, but depends on how the matchings are realised when labels are added. Figure 9 illustrates what can happen when B has points of multiplicity greater than 1. Since we are concerned only with the existence of the composite matching, this ambiguity does not trouble us.

Here is the first substantial-looking result comparing the interleaving and bottleneck distances.

Theorem 4.9. *Let \mathbb{U}, \mathbb{V} be decomposable persistence modules (i.e. direct sums of interval modules). Then*

$$d_i(\mathbb{U}, \mathbb{V}) \leq d_b(\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})).$$

(We remind the reader that tameness is not required to define \mathbf{Dgm} and \mathbf{dgm} in this case: see section 1.6.)

Proof. We show that whenever there exists a δ -matching between $\mathbf{dgm}(\mathbb{U})$ and $\mathbf{dgm}(\mathbb{V})$, we have $d_i(\mathbb{U}, \mathbb{V}) \leq \delta$. The result follows by taking the infimum over all such δ .

Let \mathbf{M} be a δ -matching between the two diagrams. Since the points in each diagram correspond to the interval summands of the module, we can construct from \mathbf{M} a partial matching between the interval summands of \mathbb{U} and \mathbb{V} .

Re-write \mathbb{U} and \mathbb{V} in the form

$$\mathbb{U} = \bigoplus_{\ell \in L} \mathbb{U}_\ell, \quad \mathbb{V} = \bigoplus_{\ell \in L} \mathbb{V}_\ell$$

so that each pair $(\mathbb{U}_\ell, \mathbb{V}_\ell)$ is one of the following:

- a pair of matched intervals;
- \mathbb{U}_ℓ is an unmatched interval, $\mathbb{V}_\ell = 0$;
- \mathbb{V}_ℓ is an unmatched interval, $\mathbb{U}_\ell = 0$.

In each case, by Propositions 4.6 and 4.7, we have $d_i(\mathbb{U}_\ell, \mathbb{V}_\ell) \leq \delta$. It follows from Proposition 4.5 that $d_i(\mathbb{U}, \mathbb{V}) \leq \delta$. \square

We complete this section with the postponed proof.

Proof of Proposition 4.6. We treat the case where p, q, r, s are all finite. We must show that if

$$\delta > \max(|p - r|, |q - s|)$$

then \mathbb{U}, \mathbb{V} are δ -interleaved. We define systems of linear maps

$$\begin{aligned} \Phi &= (\phi_t : U_t \rightarrow V_{t+\delta}) \\ \Psi &= (\psi_t : V_t \rightarrow U_{t+\delta}) \end{aligned}$$

and then show that the interleaving relations

$$\Phi 1_{\mathbb{U}}^\eta = 1_{\mathbb{V}}^\eta \Phi, \quad \Psi 1_{\mathbb{V}}^\eta = 1_{\mathbb{U}}^\eta \Psi, \quad \Psi \Phi = 1_{\mathbb{U}}^{2\delta}, \quad \Phi \Psi = 1_{\mathbb{V}}^{2\delta}$$

hold.

The definition of the maps ϕ_t, ψ_t is straightforward. Each vector space in \mathbb{U}, \mathbb{V} is equal to zero or to the field \mathbf{k} . If the domain and codomain equal \mathbf{k} , then the map is defined to be the identity $1 = 1_{\mathbf{k}}$. Otherwise, the map is necessarily 0.

The first step is to show that the systems of maps $\Phi = (\phi_t), \Psi = (\psi_t)$ are module homomorphisms. For Φ this entails verifying that the diagram

$$\begin{array}{ccc} U_t & \longrightarrow & U_{t+\eta} \\ & \searrow & \searrow \\ & & V_{t+\delta} \longrightarrow V_{t+\eta+\delta} \end{array}$$

commutes for all t and for all $\eta > 0$. Because of the special form of the vector spaces and maps, it is enough to show that the situation is not one of the following:

$$\begin{array}{ccc} \bullet & \longrightarrow & \bullet \\ & \searrow & \searrow \\ & & \circ \longrightarrow \bullet \end{array} \quad \text{or} \quad \begin{array}{ccc} \bullet & \longrightarrow & \circ \\ & \searrow & \searrow \\ & & \bullet \longrightarrow \bullet \end{array}$$

Here a filled circle \bullet indicates that the vector space is \mathbf{k} , and an open circle \circ denotes that it is zero. For the first situation to occur, one must have

$$p^* < t \quad \text{and} \quad t + \delta < r^*$$

but $\delta > r - p$ so this is impossible. For the second situation to occur, one must have

$$q^* < t + \eta \quad \text{and} \quad t + \eta + \delta < s^*$$

but $\delta > s - q$ so this too is impossible. It follows that Φ is a module homomorphism. By symmetry, so is Ψ .

The second step is to show that $\Psi\Phi = 1_{\mathbb{U}}^{2\delta}$ and $\Phi\Psi = 1_{\mathbb{V}}^{2\delta}$. For the first of these, we must verify that the diagram

$$\begin{array}{ccc} U_t & \longrightarrow & U_{t+2\delta} \\ & \searrow & \nearrow \\ & & V_{t+\delta} \end{array}$$

commutes for all t . This time the only forbidden configuration is:

$$\begin{array}{ccc} \bullet & \longrightarrow & \bullet \\ & \searrow & \nearrow \\ & & \circ \end{array}$$

If this occurs then the top row implies:

$$p^* < t \quad \text{and} \quad t + 2\delta < q^*$$

Since $\delta > r - p$ and $\delta > q - s$ we infer that

$$r^* < (p + \delta)^* < t + \delta < (q - \delta)^* < s^*$$

which implies that the circle on the bottom row is filled after all. It follows that $\Psi\Phi = 1_{\mathbb{U}}^{2\delta}$. By symmetry, $\Phi\Psi = 1_{\mathbb{V}}^{2\delta}$.

This completes the proof when p, q, r, s are finite. The infinite cases are similar. □

Remark. In the proof we have used and manipulated inequalities between decorated and undecorated real numbers. The definitions are the natural ones:

$$p^- < p < p^+ < q^- < q < q^+$$

whenever $p < q$ are undecorated real numbers. The proof can be written without this; for instance, the inequality $p^* < t$ can be replaced by the marginally weaker assertion $p \leq t$, and so on.

4.3. The bottleneck distance (continued). If \mathbf{A}, \mathbf{B} are locally finite, it turns out that the ‘inf’ is attained in the definition

$$d_b(\mathbf{A}, \mathbf{B}) = \inf(\delta \mid \text{there exists a } \delta\text{-matching between } \mathbf{A} \text{ and } \mathbf{B}),$$

and can be replaced by ‘min’. This will allow us to make a tighter statement of the stability theorem (4.11'') for q-tame modules. See Theorem 4.20.

Theorem 4.10. *Let \mathbf{A}, \mathbf{B} be locally finite multisets in the extended open half-plane $\overline{\mathcal{H}}^\circ$. Suppose for every $\eta > \delta$ there exists an η -matching between \mathbf{A}, \mathbf{B} . Then there exists a δ -matching between \mathbf{A}, \mathbf{B} .*

The assertion is obvious if \mathbf{A}, \mathbf{B} are finite. The general case is proved using a compactness argument, as follows.

Proof. As usual we treat \mathbf{A}, \mathbf{B} as sets rather than multisets.

For every integer $n \geq 1$, let \mathbf{M}_n be a $(\delta + \frac{1}{n})$ -matching between \mathbf{A}, \mathbf{B} . The plan is to construct a δ -matching \mathbf{M} from the sequence (\mathbf{M}_n) . In practice, we work with the indicator functions

$$\begin{aligned} \chi &: \mathbf{A} \times \mathbf{B} \rightarrow \{0, 1\} \\ \chi_n &: \mathbf{A} \times \mathbf{B} \rightarrow \{0, 1\} \end{aligned}$$

of the partial matchings \mathbf{M}, \mathbf{M}_n .

The first step is to construct χ as a limit of the sequence (χ_n) . Take a fixed enumeration

$$((\alpha_\ell, \beta_\ell) \mid \ell \geq 1)$$

of the countable set $\mathbf{A} \times \mathbf{B}$. We will inductively construct a descending sequence

$$\mathbf{N} = \mathbf{N}_0 \supseteq \mathbf{N}_1 \supseteq \cdots \supseteq \mathbf{N}_\ell \supseteq \cdots$$

of infinite subsets of the natural numbers, with the property that $\chi_n(\alpha_\ell, \beta_\ell)$ takes the same value for all $n \in \mathbf{N}_\ell$. Having done so, we define $\chi(\alpha_\ell, \beta_\ell)$ to be this common value.

The construction of \mathbf{N}_ℓ is straightforward: once $\mathbf{N}_{\ell-1}$ is defined, at least one of the two sets

$$\{n \in \mathbf{N}_{\ell-1} \mid \chi_n(\alpha_\ell, \beta_\ell) = 0\} \quad \text{and} \quad \{n \in \mathbf{N}_{\ell-1} \mid \chi_n(\alpha_\ell, \beta_\ell) = 1\}$$

has infinite cardinality, and that will be our \mathbf{N}_ℓ . (If both, then either will do.) Repeat.

Lemma. *If \mathbf{F} is any finite subset of $\mathbf{A} \times \mathbf{B}$, then there exists $\ell \geq 1$ such that*

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta)$$

for all $(\alpha, \beta) \in \mathbf{F}$ and for all $n \in \mathbf{N}_\ell$.

Proof. Indeed, select ℓ such that $(\alpha_1, \beta_1), \dots, (\alpha_\ell, \beta_\ell)$ include all of F . □

The second step is to verify that χ is the indicator function of a δ -matching. There are several items to check.

- For $\alpha \in A$ there is at most one $\beta \in B$ such that $\chi(\alpha, \beta) = 1$.

Proof. Suppose $\chi(\alpha, \beta) = \chi(\alpha, \beta') = 1$ for distinct elements $\beta, \beta' \in B$. By the lemma, there exists n such that $\chi_n(\alpha, \beta) = \chi_n(\alpha, \beta') = 1$, which contradicts the fact that M_n is a partial matching. □

- For $\alpha \in A$ with $d^\infty(\alpha, \Delta) > \delta$, there is at least one $\beta \in B$ such that $\chi(\alpha, \beta) = 1$.

Proof. Select N such that $d^\infty(\alpha, \Delta) > \delta + \frac{1}{N}$. Then the set

$$F_\alpha = \left\{ \beta \in B \mid d^\infty(\alpha, \beta) \leq \delta + \frac{1}{N} \right\}$$

is finite, since B is locally finite and these points lie in a square bounded away from the diagonal. By the lemma, there exists ℓ such that

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta)$$

for all $\beta \in F_\alpha$ and for all $n \in \mathbf{N}_\ell$. On the other hand, if $n \geq N$, then M_n matches α with some $\beta \in F_\alpha$. Combining these observations,

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta) = 1$$

for sufficiently large $n \in \mathbf{N}_\ell$ and for some $\beta \in F_\alpha$. □

By symmetry we have:

- For $\beta \in B$ there is at most one $\alpha \in A$ such that $\chi(\alpha, \beta) = 1$.
- For $\beta \in B$ with $d^\infty(\beta, \Delta) > \delta$, there is at least one $\alpha \in A$ such that $\chi(\alpha, \beta) = 1$.

Finally:

- If $\chi(\alpha, \beta) = 1$ then $d^\infty(\alpha, \beta) \leq \delta$.

Proof. By the lemma, there are infinitely many n for which $\chi_n(\alpha, \beta) = 1$. Then

$$d^\infty(\alpha, \beta) \leq \delta + \frac{1}{n}$$

for these n . Since n may be arbitrarily large, the result follows. □

These five bullet points confirm that M , defined by its indicator function χ , is a δ -matching between A, B . □

Remark. Although we have chosen to give a direct argument, we point out that Theorem 4.10 is an instance of the compactness theorem in first-order logic. The set of constraints that must be satisfied by an η -matching can be formulated as a theory T_η on a collection of binary-valued variables $x_{\alpha\beta}$. An η -matching is precisely a model for that theory. The theory

T_δ is seen to be logically equivalent to the union of the theories $(T_\eta \mid \eta > \delta)$. If each T_η has a model, then any finite subtheory of this union has a model, therefore by compactness T_δ has a model. The details are left as an exercise for the interested reader.

4.4. The isometry theorem. Having defined the interleaving distance and the bottleneck distance, we can now state the main theorem.

Theorem 4.11. *Let \mathbb{U}, \mathbb{V} be q -tame persistence modules. Then*

$$d_i(\mathbb{U}, \mathbb{V}) = d_b(\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V}))$$

(Recall that \mathbf{dgm} denotes the undecorated persistence diagram.)

The result naturally falls into two parts: the ‘stability theorem’ [10, 6]

$$(4.11') \quad d_i(\mathbb{U}, \mathbb{V}) \geq d_b(\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})),$$

and the ‘converse stability theorem’ [16]

$$(4.11'') \quad d_i(\mathbb{U}, \mathbb{V}) \leq d_b(\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})).$$

The proof of the converse stability theorem (4.11'') occupies section 4.5. We have already seen the result for decomposable modules, in Theorem 4.9, so it is a matter of extending the result to q -tame modules that are not known to be decomposable.

The proof of the stability theorem (4.11') is given in sections 4.6 and 4.7.

4.5. The converse stability theorem. In this section we deduce the converse stability inequality (4.11'') for q -tame modules from Theorem 4.9. A similar argument is given by Lesnick [16]. It may turn out that q -tame modules are decomposable, in which case this extra work becomes unnecessary.

The key idea is that persistence modules can be approximated by better-behaved persistence modules, using a procedure we call ‘smoothing’.

Definition 4.12. Let \mathbb{V} be a persistence module, and let $\epsilon > 0$. The ϵ -**smoothing** of \mathbb{V} is the persistence module \mathbb{V}^ϵ defined to be the image of the following map:

$$1_{\mathbb{V}}^{2\epsilon} : \mathbb{V}[t - \epsilon] \rightarrow \mathbb{V}[t + \epsilon]$$

In other words, V_t^ϵ is the image of the map

$$v_{t-\epsilon}^{t+\epsilon} : V_{t-\epsilon} \rightarrow V_{t+\epsilon},$$

and, for $s < t$, the shift map $(v^\epsilon)_s^t$ is the restriction of $v_{s+\epsilon}^{t+\epsilon}$.

Thus we have a factorisation of $1_{\mathbb{V}}^{2\epsilon}$

$$(4.13) \quad \mathbb{V}[t - \epsilon] \longrightarrow \mathbb{V}^\epsilon \longrightarrow \mathbb{V}[t + \epsilon]$$

where the first map is surjective and the second map is injective (at all t). At a given index t this is the sequence:

$$V_{t-\epsilon} \xrightarrow{v_{t-\epsilon}^{t+\epsilon}} V_t^\epsilon \xrightarrow{1} V_{t+\epsilon}$$

Proposition 4.14. *Let \mathbb{V} be a persistence module. Then $d_i(\mathbb{V}, \mathbb{V}^\epsilon) \leq \epsilon$.*

Proof. Indeed, the maps in (4.13) give an ϵ -interleaving. □

Smoothing changes the persistence diagram in a predictable way. Here is the atomic example (which the reader can easily verify):

Example 4.15. Let $\mathbb{V} = \mathbb{I}(p^*, q^*)$. Then:

$$\mathbb{V}^\epsilon = \begin{cases} \mathbb{I}((p + \epsilon)^*, (q - \epsilon)^*) & \text{if } (p + \epsilon)^* < (q - \epsilon)^* \\ 0 & \text{otherwise} \end{cases}$$

In other words, ϵ -smoothing shrinks the interval by ϵ at both ends.

Proposition 4.16. *The persistence diagram of \mathbb{V}^ϵ is obtained from the persistence diagram of \mathbb{V} by applying the translation $T_\epsilon : (p, q) \mapsto (p + \epsilon, q - \epsilon)$ to the part of the extended half-plane that lies above the line $\Delta_\epsilon = \{(t - \epsilon, t + \epsilon) \mid t \in \mathbf{R}\}$.*

In the case where \mathbf{Dgm} is not everywhere defined, the proposition is understood to include the assertion that the finite r -interior of the persistence measure, and hence the region where \mathbf{Dgm} is defined, is shifted by T_ϵ .

Information on $\mathbf{Dgm}(\mathbb{V})$ that lies below the line Δ_ϵ is lost in $\mathbf{Dgm}(\mathbb{V}^\epsilon)$.

Proof. We consider three different cases. Case (ii) is subsumed by case (iii), but the proof is easier to digest.

(i) **\mathbb{V} is decomposable.** The image of a direct sum of maps is the direct sum of the images of the maps; therefore ϵ -smoothing commutes with direct sums:

$$\left[\bigoplus_{\ell \in L} \mathbb{V}_\ell \right]^\epsilon = \bigoplus_{\ell \in L} \mathbb{V}_\ell^\epsilon$$

By Example 4.15, the proposition is true for interval modules. It is therefore true for direct sums of interval modules.

(ii) **\mathbb{V} is \mathbf{q} -tame.** It is enough to show that the rank function of \mathbb{V}^ϵ is equal to the rank function of \mathbb{V} shifted by T_ϵ , since this determines the persistence measure and hence the persistence diagram. Specifically, for all $b < c$ we require:

$$\text{rank}[V_b^\epsilon \rightarrow V_c^\epsilon] = \text{rank}[V_{b-\epsilon} \rightarrow V_{c+\epsilon}]$$

In fact, these maps are related by the sequence

$$V_{b-\epsilon} \longrightarrow V_b^\epsilon \longrightarrow V_c^\epsilon \longrightarrow V_{c+\epsilon}$$

where the first map is surjective and the third map is injective. Since the rank of a linear map is unchanged by pre-composing with a surjective map, or post-composing with an injective map, it follows that the rank of the middle map is equal to the rank of the composite. This is what we wished to prove.

(iii) **general case.** We show that the persistence measure of \mathbb{V}^ϵ is equal to the persistence measure of \mathbb{V} shifted by T_ϵ . Writing

$$A = a - \epsilon, \quad B = b - \epsilon, \quad C = c + \epsilon, \quad D = d + \epsilon$$

this means showing that

$$\langle \circ_A - \bullet_B - \bullet_C - \circ_D \mid \mathbb{V} \rangle = \langle \circ_a - \bullet_b - \bullet_c - \circ_d \mid \mathbb{V}^\epsilon \rangle$$

for all $a < b \leq c < d$.

The proof is based on the following commutative diagram

$$\begin{array}{ccccccc} & & V_a^\epsilon & \longrightarrow & V_b^\epsilon & \longrightarrow & V_c^\epsilon & \longrightarrow & V_d^\epsilon \\ & \nearrow & & & \nearrow & & \searrow & & \searrow \\ V_A & \longrightarrow & V_B & \longrightarrow & V_C & \longrightarrow & V_D \end{array}$$

in which the maps \nearrow are surjective and the maps \searrow are injective. We will construct various quiver diagrams from this. Surjectivity means that

$$\langle \circ_A - \bullet_a \rangle = 0 \quad \text{and} \quad \langle \circ_B - \bullet_b \rangle = 0,$$

and injectivity means that

$$\langle \bullet_c - \circ_C \rangle = 0 \quad \text{and} \quad \langle \bullet_d - \circ_D \rangle = 0.$$

Moreover, by the restriction principle, interval types containing any of these ‘forbidden’ configurations occur with multiplicity zero.

Then

$$\begin{aligned} \langle \circ_A \text{ --- } \bullet_b \text{ --- } \bullet_c \text{ --- } \circ_D \rangle &= \langle \circ_A - \bullet_B - \bullet_b - \bullet_c - \bullet_C - \circ_D \rangle \\ &\quad + \text{three other terms} \\ &= \langle \circ_A - \bullet_B - \bullet_b - \bullet_c - \bullet_C - \circ_D \rangle \\ &= \langle \circ_A - \bullet_B \text{ --- } \bullet_C - \circ_D \rangle \\ &= \langle \circ_A - \bullet_B - \bullet_C - \circ_D \mid \mathbb{V} \rangle \end{aligned}$$

and at the same time

$$\begin{aligned} \langle \circ_A \text{ --- } \bullet_b \text{ --- } \bullet_c \text{ --- } \circ_D \rangle &= \langle \circ_A - \circ_a - \bullet_b - \bullet_c - \circ_d - \circ_D \rangle \\ &\quad + \text{three other terms} \\ &= \langle \circ_A - \circ_a - \bullet_b - \bullet_c - \circ_d - \circ_D \rangle \\ &= \langle \text{--- } \circ_a \text{ --- } \bullet_b \text{ --- } \bullet_c \text{ --- } \circ_d \text{ ---} \rangle \\ &= \langle \circ_a - \bullet_b - \bullet_c - \circ_d \mid \mathbb{V}^\epsilon \rangle \end{aligned}$$

so we get the required equality. The six ‘other terms’ are all zero because they contain forbidden configurations. \square

Corollary 4.17. *Let \mathbb{V} be a q -tame persistence module. Then $d_b(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{V}^\epsilon)) \leq \epsilon$.*

Proof. Indeed, an ϵ -matching is defined as follows:

$$(p, q) \in \mathbf{dgm}(\mathbb{V}^\epsilon) \quad \leftrightarrow \quad (p - \epsilon, q + \epsilon) \in \mathbf{dgm}(\mathbb{V})$$

This is bijective except for the unmatched points of $\mathbf{dgm}(\mathbb{V})$, which lie on or below the line Δ_ϵ , and therefore have distance at most ϵ from the diagonal. \square

Our claim that smoothing makes a persistence module ‘better-behaved’ amounts to the following fact.

Proposition 4.18. *Let \mathbb{V} be a q -tame persistence module over \mathbf{R} . Then \mathbb{V}^ϵ is locally finite. In particular, \mathbb{V}^ϵ is decomposable into intervals.*

Proof. We make use of the characterisation in Proposition 2.21. Note first that

$$\dim(V_t^\epsilon) = \text{rank}[V_{t-\epsilon} \rightarrow V_{t+\epsilon}] < \infty$$

so condition (i) is satisfied.

For condition (ii), we must show that there is a locally finite set $\mathbf{S} \subset \mathbf{R}$ of ‘singular values’ with the property that \mathbb{V}^ϵ is constant over each interval of the open set $\mathbf{R} - \mathbf{S}$.

The set \mathbf{S} is easily described: consider all points $(p^*, q^*) \in \mathbf{Dgm}(\mathbb{V})$ such that $p^* + 2\epsilon < q^*$. Define \mathbf{S} to be the union of the sets $\{p + \epsilon, q - \epsilon\}$ over all such points of $\mathbf{Dgm}(\mathbb{V})$.

In other words, we take all points of $\mathbf{Dgm}(\mathbb{V})$ which lie on⁵ or above the line

$$\Delta_\epsilon = \{(t - \epsilon, t + \epsilon) \mid t \in \mathbf{R}\}$$

and then project each point both vertically and horizontally onto Δ_ϵ , abandoning the decoration. The resulting subset of Δ_ϵ defines $\mathbf{S} \subset \mathbf{R}$, using the identification $(t - \epsilon, t + \epsilon) \leftrightarrow t$. See Figure 10 (left).

We show that \mathbf{S} is locally finite. Indeed, for any t and for any $\eta < \epsilon$,

$$\begin{aligned} \text{card}(S \cap (t - \eta, t + \eta)) &\leq 2\mu_{\mathbb{V}}([-\infty, t + \eta - \epsilon] \times [t - \eta + \epsilon, +\infty]) \\ &= 2 \text{rank}[V_{t+\eta-\epsilon} \rightarrow V_{t-\eta+\epsilon}] \end{aligned}$$

which is finite.

We show that \mathbb{V} is constant over each component of $\mathbf{R} - \mathbf{S}$.

Claim. The map $V_b^\epsilon \rightarrow V_c^\epsilon$ is an isomorphism if and only if $\mathbf{Dgm}(\mathbb{V})$ does not meet the union of the rectangles

$$[-\infty, b - \epsilon] \times [b + \epsilon, c + \epsilon] \quad \text{and} \quad [b - \epsilon, c - \epsilon] \times [c + \epsilon, +\infty].$$

See Figure 10 (right).

We split the claim into two statements.

⁵ Points on the line must be decorated $((t - \epsilon)^-, (t + \epsilon)^+)$.

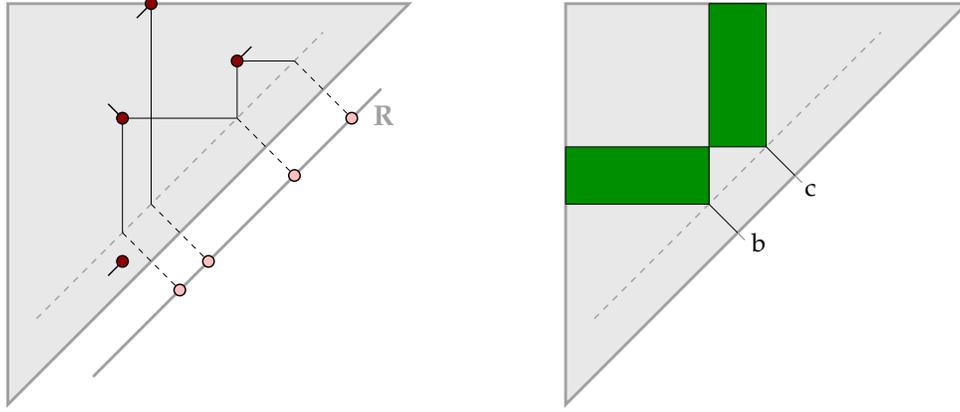


FIGURE 10. Constructing the set \mathbf{S} from the persistence diagram (left). The exclusion zone for $V_b^\epsilon \rightarrow V_c^\epsilon$ to be an isomorphism (right) .

Part 1. The map $V_b^\epsilon \rightarrow V_c^\epsilon$ is injective if and only if

$$\mu_{\mathbb{V}}([-\infty, b - \epsilon] \times [b + \epsilon, c + \epsilon]) = 0.$$

Part 2. The map $V_b^\epsilon \rightarrow V_c^\epsilon$ is surjective if and only if

$$\mu_{\mathbb{V}}([b - \epsilon, c - \epsilon] \times [c + \epsilon, +\infty]) = 0.$$

We approach this in the usual way. Consider the following diagram of vector spaces:

$$\begin{array}{ccccc}
 & & V_b^\epsilon & \longrightarrow & V_c^\epsilon \\
 & \nearrow & & \searrow & \searrow \\
 & & V_{b-\epsilon} & \longrightarrow & V_{c-\epsilon} \\
 & & & & \nearrow \\
 & & & & V_{b+\epsilon} & \longrightarrow & V_{c+\epsilon}
 \end{array}$$

where the maps \nearrow are surjective and the maps \searrow are injective. Using the abbreviations

$$b_- = b - \epsilon, \quad b_+ = b + \epsilon, \quad c_- = c - \epsilon, \quad c_+ = c + \epsilon$$

we can write down the ‘forbidden’ configurations

$$\circ_{b_-} \text{---} \bullet_b \quad \bullet_b \text{---} \circ_{b_+} \quad \circ_{c_-} \text{---} \bullet_c \quad \bullet_c \text{---} \circ_{c_+}$$

which occur with multiplicity zero.

For Part 1, note that $V_b^\epsilon \rightarrow V_c^\epsilon$ is injective if and only if

$$\langle \bullet_b \text{---} \circ_c \mid \mathbb{V}^\epsilon \rangle = 0.$$

Taking the forbidden configurations into account, we calculate:

$$\begin{aligned}
\langle \bullet_b - \circ_c \mid \mathbb{V}^\epsilon \rangle &= \langle \text{---} \bullet_b - \circ_c \text{---} \rangle \\
&= \langle \text{---} \bullet_b - \circ_c - \circ_{c_+} \rangle \\
&= \langle \text{---} \bullet_b \text{---} \circ_{c_+} \rangle \\
&= \langle \bullet_{b_-} - \bullet_b - \bullet_{b_+} - \circ_{c_+} \rangle \\
&= \langle \bullet_{b_-} \text{---} \bullet_{b_+} - \circ_{c_+} \rangle
\end{aligned}$$

which is exactly $\mu_{\mathbb{V}}([-\infty, b - \epsilon] \times [b + \epsilon, c + \epsilon])$.

For Part 2, note that $V_b^\epsilon \rightarrow V_c^\epsilon$ is surjective if and only if

$$\langle \circ_b - \bullet_c \mid \mathbb{V}^\epsilon \rangle = 0.$$

We calculate:

$$\begin{aligned}
\langle \circ_b - \bullet_c \mid \mathbb{V}^\epsilon \rangle &= \langle \text{---} \circ_b - \bullet_c \text{---} \rangle \\
&= \langle \circ_{b_-} - \circ_b - \bullet_c \text{---} \rangle \\
&= \langle \circ_{b_-} \text{---} \bullet_c \text{---} \rangle \\
&= \langle \circ_{b_-} - \bullet_{c_-} - \bullet_c - \bullet_{c_+} \rangle \\
&= \langle \circ_{b_-} - \bullet_{c_-} \text{---} \bullet_{c_+} \rangle
\end{aligned}$$

which is exactly $\mu_{\mathbb{V}}([b - \epsilon, c - \epsilon] \times [c + \epsilon, +\infty])$.

This completes the proof of the claim.

The claim implies that \mathbb{V} is constant on each connected component of $\mathbf{R} - \mathbf{S}$. Indeed, if $\mathbf{S} \cap [b, c] = \emptyset$ then the exclusion zone does not contain any points of $\mathbf{Dgm}(\mathbb{V})$, and therefore $(v^\epsilon)_b^c$ is an isomorphism.

Thus \mathbb{V}^ϵ satisfies both conditions of Proposition 2.21 and is therefore locally finite. \square

We are now ready to prove the converse stability theorem for q-tame persistence modules, using the triangle inequalities for d_i, d_b and our results on ϵ -smoothing.

Proof of (4.11''). Let \mathbb{U}, \mathbb{V} be q-tame persistence modules. For any $\epsilon > 0$, the ϵ -smoothings $\mathbb{U}^\epsilon, \mathbb{V}^\epsilon$ are decomposable, so the converse stability theorem applies to them. Then:

$$\begin{aligned}
d_i(\mathbb{U}, \mathbb{V}) &\leq d_i(\mathbb{U}^\epsilon, \mathbb{V}^\epsilon) + 2\epsilon && \text{by Proposition 4.14} \\
&\leq d_b(\mathbf{dgm}(\mathbb{U}^\epsilon), \mathbf{dgm}(\mathbb{V}^\epsilon)) + 2\epsilon && \text{by Theorem 4.9} \\
&\leq d_b(\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})) + 4\epsilon && \text{by Corollary 4.17}
\end{aligned}$$

Since this is true for all $\epsilon > 0$, we deduce that

$$d_i(\mathbb{U}, \mathbb{V}) \leq d_b(\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})).$$

The converse stability theorem for q-tame modules is proved. \square

We finish this section with a characterisation of q-tame modules.

Theorem 4.19. *A persistence module \mathbb{V} is q -tame if and only if it can be approximated, in the interleaving distance, by locally finite modules.*

Proof. If \mathbb{V} is q -tame then it is approximated by the modules \mathbb{V}^ϵ , which are locally finite by Proposition 4.18. Conversely, suppose \mathbb{V} is approximated by locally finite modules. Suppose $b < c$ is given. Let \mathbb{W} be a locally finite module which is ϵ -interleaved with \mathbb{V} , for some $\epsilon < (c - b)/2$. Then

$$\begin{aligned} r_b^c &= \text{rank}[V_b \rightarrow V_c] \\ &= \text{rank}[V_b \rightarrow W_{b+\epsilon} \rightarrow W_{c-\epsilon} \rightarrow V_c] \\ &\leq \dim(W_{b+\epsilon}) \\ &< \infty. \end{aligned}$$

It follows that \mathbb{V} is q -tame. □

There certainly exist q -tame modules which are not locally finite, such as:

$$\bigoplus_{n=1}^{\infty} \mathbb{I}[-\frac{1}{n}, \frac{1}{n}]$$

As stated before, we do not know whether all q -tame modules can be decomposed into intervals.

4.6. The stability theorem. The inequality (4.11'') can be expressed in the following form:

Theorem 4.20. *Let \mathbb{U}, \mathbb{V} be q -tame persistence modules which are δ^+ -interleaved. Then there exists a δ -matching between the multisets $\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})$.*

It is easier to prove the following. (Notice the missing $+$.)

Theorem 4.21. *Let \mathbb{U}, \mathbb{V} be q -tame persistence modules which are δ -interleaved. Then there exists a δ -matching between the multisets $\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})$.*

Theorem 4.10 allows us to deduce Theorem 4.20 from Theorem 4.21: if \mathbb{U}, \mathbb{V} are δ^+ -interleaved then there is an η -matching between their diagrams for every $\eta > \delta$, hence there is a δ -matching.

The proof of Theorem 4.21 depends on two main ingredients:

- The interpolation lemma (3.4), which embeds \mathbb{U}, \mathbb{V} within a 1-parameter family.
- The box lemma (4.22), which relates the persistence measures of \mathbb{U}, \mathbb{V} locally.

Once these ingredients are in place, the theorem can be proved using the continuity method of [10]. Our persistence diagrams may have infinite cardinality, so we will need an additional compactness argument to finish off the proof.

Let $R = [a, b] \times [c, d]$ be a rectangle in $\overline{\mathbf{R}}^2$. The δ -**thickening** of R is the rectangle

$$R^\delta = [a - \delta, b + \delta] \times [c - \delta, d + \delta].$$

For convenience we will write

$$A = a - \delta, \quad B = b + \delta, \quad C = c - \delta, \quad D = d + \delta$$

in this situation. For infinite rectangles, note that $-\infty - \delta = -\infty$ and $+\infty + \delta = +\infty$.

We can also thicken an individual point: if $\alpha = (p, q)$ then

$$\alpha^\delta = [p - \delta, p + \delta] \times [q - \delta, q + \delta]$$

for $\delta > 0$.

Lemma 4.22 (Box lemma [10]). *Let \mathbb{U}, \mathbb{V} be a δ -interleaved pair of persistence modules. Let R be a rectangle whose δ -thickening R^δ lies above the diagonal. Then $\mu_{\mathbb{U}}(R) \leq \mu_{\mathbb{V}}(R^\delta)$ and $\mu_{\mathbb{V}}(R) \leq \mu_{\mathbb{U}}(R^\delta)$.*

If we use the extension convention (section 2.5) we can state the lemma without the requirement that R^δ lies above the diagonal, because the convention gives $\mu(R^\delta) = \infty$ if it doesn't.

Proof. Write $R = [a, b] \times [c, d]$ and $R^\delta = [A, B] \times [C, D]$ as above. Thanks to the interleaving, the finite modules

$$\mathbb{U}_{a,b,c,d} : U_a \rightarrow U_b \rightarrow U_c \rightarrow U_d$$

and

$$\mathbb{V}_{A,B,C,D} : V_A \rightarrow V_B \rightarrow V_C \rightarrow V_D$$

are restrictions of the following 8-term module

$$\mathbb{W} : V_A \xrightarrow{\Psi} U_a \longrightarrow U_b \xrightarrow{\Phi} V_B \longrightarrow V_C \xrightarrow{\Psi} U_c \longrightarrow U_d \xrightarrow{\Phi} V_D$$

where Φ, Ψ are the interleaving maps.

Using the restriction principle, we calculate:

$$\begin{aligned} \mu_{\mathbb{V}}([A, B] \times [C, D]) &= \langle \circ_A \text{-----} \bullet_B \text{---} \bullet_C \text{-----} \circ_D \mid \mathbb{V} \rangle \\ &= \langle \circ_A \text{-----} \bullet_B \text{---} \bullet_C \text{-----} \circ_D \mid \mathbb{W} \rangle \\ &= \langle \circ_A \text{---} \circ_a \text{---} \bullet_b \text{---} \bullet_B \text{---} \bullet_C \text{---} \bullet_c \text{---} \circ_d \text{---} \circ_D \mid \mathbb{W} \rangle \\ &\quad + \text{eight other terms} \\ &\geq \langle \circ_A \text{---} \circ_a \text{---} \bullet_b \text{---} \bullet_B \text{---} \bullet_C \text{---} \bullet_c \text{---} \circ_d \text{---} \circ_D \mid \mathbb{W} \rangle \\ &= \langle \text{---} \circ_a \text{---} \bullet_b \text{-----} \bullet_c \text{---} \circ_d \text{-----} \mid \mathbb{W} \rangle \\ &= \langle \text{---} \circ_a \text{---} \bullet_b \text{-----} \bullet_c \text{---} \circ_d \text{-----} \mid \mathbb{U} \rangle \\ &= \mu_{\mathbb{U}}([a, b] \times [c, d]) \end{aligned}$$

This proves $\mu_{\mathbb{U}}(R) \leq \mu_{\mathbb{V}}(R^\delta)$. The inequality $\mu_{\mathbb{V}}(R) \leq \mu_{\mathbb{U}}(R^\delta)$ follows by symmetry. \square

Recall the measures at infinity defined in section 2.6. By considering the appropriate limits, we immediately have:

Proposition 4.23 (Box inequalities at infinity). *Let μ, ν be r -measures on $\overline{\mathbf{R}^2}$ which satisfy a one-sided box inequality with parameter δ*

$$\mu(R) \leq \nu(R^\delta)$$

for all rectangles $R \in \text{Rect}(\overline{\mathbf{R}^2})$. Then

$$\begin{aligned} \mu([a, b], -\infty) &\leq \nu([A, B], -\infty), & \mu(-\infty, [c, d]) &\leq \nu(-\infty, [C, D]), \\ \mu([a, b], +\infty) &\leq \nu([A, B], +\infty), & \mu(+\infty, [c, d]) &\leq \nu(+\infty, [C, D]), \end{aligned}$$

for all $a < b$ and $c < d$; and

$$\begin{aligned} \mu(-\infty, -\infty) &\leq \nu(-\infty, -\infty), & \mu(+\infty, -\infty) &\leq \nu(+\infty, -\infty), \\ \mu(-\infty, +\infty) &\leq \nu(-\infty, +\infty), & \mu(+\infty, +\infty) &\leq \nu(+\infty, +\infty). \end{aligned}$$

Here $A = a - \delta$, $B = b + \delta$, $C = c - \delta$, $D = d + \delta$. □

Consequently, if \mathbb{U}, \mathbb{V} are δ -interleaved persistence modules then $\mu_{\mathbb{U}}, \mu_{\mathbb{V}}$ satisfy (two-sided) box inequalities on $(-\infty, \mathbf{R})$ and $(\mathbf{R}, +\infty)$ as well as the equality $\mu_{\mathbb{U}}(-\infty, +\infty) = \mu_{\mathbb{V}}(-\infty, +\infty)$.

4.7. The measure stability theorem. We now embed Theorem 4.21 as a special case of a stability theorem for the diagrams of abstract r -measures. The more general statement is no more difficult⁶ to prove, and seems to be the natural home for the result.

Let \mathcal{D} be an open subset of $\overline{\mathbf{R}^2}$. For $\alpha \in \mathcal{D}$, define the **exit distance** of α to be

$$\text{ex}^\infty(\alpha, \mathcal{D}) = d^\infty(\alpha, \overline{\mathbf{R}^2} - \mathcal{D}) = \min(d^\infty(\alpha, x) \mid x \in \overline{\mathbf{R}^2} - \mathcal{D}).$$

For instance, for the extended half-plane we have $\text{ex}^\infty(\alpha, \overline{\mathcal{H}}) = d^\infty(\alpha, \Delta)$.

Let \mathbf{A}, \mathbf{B} be multisets in \mathcal{D} . A δ -**matching** between \mathbf{A}, \mathbf{B} is a partial matching $\mathbf{M} \subset \mathbf{A} \times \mathbf{B}$ such that

$$\begin{aligned} d^\infty(\alpha, \beta) &\leq \delta && \text{if } \alpha, \beta \text{ are matched,} \\ \text{ex}^\infty(\alpha, \mathcal{D}) &\leq \delta && \text{if } \alpha \in \mathbf{A} \text{ is unmatched,} \\ \text{ex}^\infty(\beta, \mathcal{D}) &\leq \delta && \text{if } \beta \in \mathbf{B} \text{ is unmatched.} \end{aligned}$$

If \mathcal{D} is not clear from the context, we refer to \mathbf{M} as a ‘ δ -matching between $(\mathbf{A}, \mathcal{D})$ and $(\mathbf{B}, \mathcal{D})$ ’.

With the same proof as Proposition 4.8, we have:

Proposition 4.24 (triangle inequality). *If $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are multisets in \mathcal{D} and there exists a δ_1 -matching between $(\mathbf{A}, \mathcal{D}), (\mathbf{B}, \mathcal{D})$ and a δ_2 -matching between $(\mathbf{B}, \mathcal{D}), (\mathbf{C}, \mathcal{D})$, then there exists a $(\delta_1 + \delta_2)$ -matching between $(\mathbf{A}, \mathcal{D}), (\mathbf{C}, \mathcal{D})$.* □

Now for the main theorem.

⁶In fact it’s a little easier to prove, because the compactness argument for diagrams with infinitely many points can be conducted more cleanly in this generality.

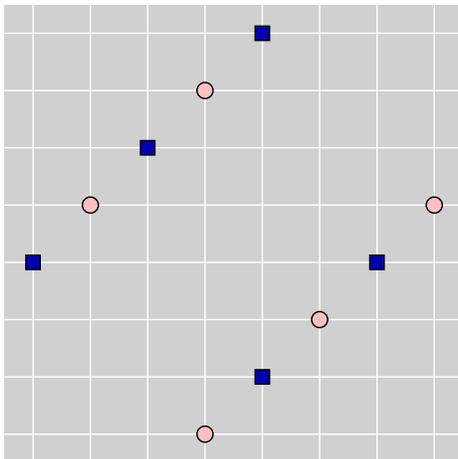


FIGURE 11. The two diagrams (5 dark blue squares; 5 light pink circles) have box distance 1 and bottleneck distance 3. Generalising this example, one can exhibit a pair of diagrams with $4k + 1$ points each, which have box distance 1 and bottleneck distance $2k + 1$.

Theorem 4.25 (stability for finite measures). *Suppose $(\mu_x \mid x \in [0, \delta])$ is a 1-parameter family of finite r -measures on an open set $\mathcal{D} \subseteq \overline{\mathbf{R}^2}$. Suppose for all $x, y \in [0, \delta]$ the box inequality*

$$\mu_x(R) \leq \mu_y(R^{|y-x|})$$

holds for all rectangles R whose $|y - x|$ -thickening $R^{|y-x|}$ belongs to $\text{Rect}(\mathcal{D})$. Then there exists a δ -matching between the undecorated diagrams $(\text{dgm}(\mu_0), \mathcal{D})$ and $(\text{dgm}(\mu_\delta), \mathcal{D})$.

In view of the Interpolation Lemma (3.4), this implies Theorem 4.21 (take $\mu_x = \mu(\mathbb{U}_x)$ and $\mathcal{D} = \overline{\mathcal{H}^0}$) and therefore the stability theorem (4.11'') for q -tame modules.

Example 4.26. The existence of a 1-parameter family interpolating between μ_0 and μ_δ may seem unnecessarily strong. It is natural to hope that two measures μ, ν which satisfy the (two-sided) box inequality with parameter δ will have diagrams $\text{dgm}(\mu), \text{dgm}(\nu)$ which are δ -matched. This is simply not true, and in fact there is no universal bound on the bottleneck distance between the two diagrams. See Figure 11.

Our goal for the rest of this section is to prove Theorem 4.25. Parts 1 and 2 closely follow the method of Cohen-Steiner, Edelsbrunner and Harer [10].

Afterwards, in section 4.8, we generalise the theorem to r -measures that are not finite.

Initial remark. Because the metric d^∞ separates $\overline{\mathbf{R}^2}$ into nine strata (the standard plane, the four lines at infinity, and the four points at infinity), we seek separate δ -matchings for each stratum that meets \mathcal{D} . We begin with the points in the standard plane.

Temporary hypothesis. Suppose initially that $\mathcal{D} \subseteq \mathbf{R}^2$.

Part 1. *The Hausdorff distance between $(\text{dgm}(\mu_x), \mathcal{D})$, and $(\text{dgm}(\mu_y), \mathcal{D})$ is at most $|y - x|$.*

Write $\mathbf{A} = \text{dgm}(\mu_x)$, $\mathbf{B} = \text{dgm}(\mu_y)$, and $\eta = |y - x|$. The assertion is understood to mean:

- If $\alpha \in \mathbf{A}$ and $\text{ex}^\infty(\alpha, \mathcal{D}) > \eta$, then there exists $\beta \in \mathbf{B}$ with $d^\infty(\alpha, \beta) \leq \eta$.
- If $\beta \in \mathbf{B}$ and $\text{ex}^\infty(\beta, \mathcal{D}) > \eta$, then there exists $\alpha \in \mathbf{A}$ with $d^\infty(\alpha, \beta) \leq \eta$.

Proof. By symmetry, it is enough to prove the first statement.

Given such an α , let $\epsilon > 0$ be small enough that $\eta + \epsilon < \text{ex}^\infty(\alpha, \mathcal{D})$. Then the box inequality gives

$$1 \leq \mu_x(\alpha^\epsilon) \leq \mu_y(\alpha^{\eta+\epsilon})$$

so there is at least one point of \mathbf{B} in the square $\alpha^{\eta+\epsilon}$. This is true for all sufficiently small $\epsilon > 0$, and moreover \mathbf{B} is locally finite. Therefore there is at least one point of \mathbf{B} in α^η . \square

Henceforth, we will write $\mathbf{A}_x = \text{dgm}(\mu_x)$ for all x .

Part 2. *The theorem is true if \mathbf{A}_x has finite cardinality for all x .*

Proof. (i) The triangle inequality for matchings includes the implication

$$\left. \begin{array}{l} \mathbf{A}_0, \mathbf{A}_x \text{ are } x\text{-matched} \\ \mathbf{A}_x, \mathbf{A}_y \text{ are } (y-x)\text{-matched} \end{array} \right\} \Rightarrow \mathbf{A}_0, \mathbf{A}_y \text{ are } y\text{-matched}$$

whenever $0 < x < y$.

(ii) We claim that for every $x \in [0, \delta]$ there exists $\rho(x) > 0$ such that $\mathbf{A}_x, \mathbf{A}_y$ are $|y-x|$ -matched whenever $y \in [0, \delta]$ with $|y-x| < \rho(x)$.

Suppose $\alpha_1, \dots, \alpha_k$ is an enumeration of the distinct points of \mathbf{A}_x , with respective multiplicities n_1, \dots, n_k . Let $\rho(x)$ be chosen to satisfy the following finite set of constraints:

$$0 < \rho(x) \leq \begin{cases} \frac{1}{2}\text{ex}^\infty(\alpha_i, \mathcal{D}) & \text{all } i \\ \frac{1}{2}d^\infty(\alpha_i, \alpha_j) & \text{all } i, j \text{ distinct} \end{cases}$$

We must show that if $|y-x| < \rho(x)$ then $\mathbf{A}_x, \mathbf{A}_y$ are $|y-x|$ -matched.

Write $\eta = |y-x|$ and let

$$(\mathbf{R}^2 - \mathcal{D})^\eta = \{\alpha \in \mathcal{D} \mid \text{ex}^\infty(\alpha, \mathcal{D}) \leq \eta\}.$$

It follows from Part 1 that \mathbf{A}_y is contained entirely in the closed set

$$(\mathbf{R}^2 - \mathcal{D})^\eta \cup \alpha_1^\eta \cup \dots \cup \alpha_k^\eta$$

and it follows from the definition of $\rho(x)$ that the terms in the union are disjoint. It is easy to count the points of \mathbf{A}_y in each square α_i^η . Let $\epsilon > 0$ be small enough that $2\eta + \epsilon < 2\rho(x)$. Then the box inequality gives

$$n_i = \mu_x(\alpha_i^\epsilon) \leq \mu_y(\alpha_i^{\eta+\epsilon}) \leq \mu_x(\alpha_i^{2\eta+\epsilon}) = n_i.$$

Thus $\mu_y(\alpha_i^{\eta+\epsilon}) = n_i$ for all small $\epsilon > 0$. We conclude that the square α_i^η contains precisely n_i points of \mathbf{A}_y .

This completes the proof of (ii), because we can match the n_i copies of α_i with the n_i points of \mathbf{A}_y in the square α_i^η , for each i , to define an η -matching between $(\mathbf{A}_x, \mathcal{D})$, $(\mathbf{A}_y, \mathcal{D})$. All points of \mathbf{A}_x are matched, and the only unmatched points of \mathbf{A}_y lie in $\mathbf{R}^2 - \mathcal{D}$ and do not need to be matched.

Items (i) and (ii) formally imply that $\mathbf{A}_0, \mathbf{A}_\delta$ are δ -matched, using the standard Heine–Borel argument. Indeed, let

$$m = \sup(x \in [0, \delta] \mid \mathbf{A}_0 \text{ and } \mathbf{A}_x \text{ are } x\text{-matched}).$$

First, m is positive; specifically $m \geq \rho(0)$. Applying (i) to $0 < m' < m$, where $\mathbf{A}_0, \mathbf{A}_{m'}$ are m' -matched and $m - m' < \rho(m)$, we deduce that $\mathbf{A}_0, \mathbf{A}_m$ are m -matched. Suppose $m < \delta$. Applying (i) to $0 < m < m''$, where $m'' - m < \rho(m)$, we deduce that $\mathbf{A}_0, \mathbf{A}_{m''}$ are m'' -matched. This contradicts the definition of m . Therefore $m = \delta$, and $\mathbf{A}_0, \mathbf{A}_\delta$ are δ -matched. \square

Part 3. *The theorem is true without assuming finite cardinality.*

Proof. Let (\mathcal{D}_n) be an increasing sequence of open subsets of \mathcal{D} whose union equals \mathcal{D} and such that each \mathcal{D}_n has compact closure. Because \mathbf{A}_x is locally finite, it follows that $\mathbf{A}_x \cap \mathcal{D}_n$ is finite for all x, n . We can therefore restrict the family of measures to each \mathcal{D}_n in turn, and apply Part 2 to get a δ -matching \mathbf{M}_n between $(\mathbf{A}_0 \cap \mathcal{D}_n, \mathcal{D}_n)$ and $(\mathbf{A}_\delta \cap \mathcal{D}_n, \mathcal{D}_n)$.

We now take a limit \mathbf{M} of the partial matchings \mathbf{M}_n , using the construction in the proof of Theorem 4.10. (This works because $\mathbf{A}_0, \mathbf{A}_\delta$ are locally finite and therefore countable.) Let χ, χ_n denote the indicator functions of \mathbf{M}, \mathbf{M}_n . As before, we have the Lemma that for any finite subset $F \subset \mathbf{A}_0 \times \mathbf{A}_\delta$, there are infinitely many $n \in \mathbf{N}$ for which

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta)$$

for all $(\alpha, \beta) \in F$.

We must show that \mathbf{M} is a δ -matching between $(\mathbf{A}_0, \mathcal{D})$ and $(\mathbf{A}_\delta, \mathcal{D})$. It is immediate that each matched pair is separated by at most δ , since this is true for every \mathbf{M}_n . The argument that each α is matched with at most one β , and vice versa, is the same as before.

Finally, suppose $\alpha \in \mathbf{A}_0$ with $\text{ex}^\infty(\alpha, \mathcal{D}) > \delta$. The square α^δ is contained in \mathcal{D} and is compact, and therefore is contained in \mathcal{D}_n for sufficiently large n . This means that $\text{ex}^\infty(\alpha, \mathcal{D}_n) > \delta$ and hence α is matched in \mathbf{M}_n for sufficiently large n . Now α has only finitely many δ -neighbours β_1, \dots, β_k in the locally finite set \mathbf{A}_δ , so by the Lemma there are infinitely many n such that $\chi(\alpha, \beta_i) = \chi_n(\alpha, \beta_i)$ for all i . By taking a sufficiently large such n , we conclude that

$$\chi(\alpha, \beta_i) = \chi_n(\alpha, \beta_i) = 1$$

for some i . Thus α is matched.

By symmetry, any $\beta \in \mathbf{A}_\delta$ with $\text{ex}^\infty(\beta, \mathcal{D}) > \delta$ is matched in \mathbf{M} to some $\alpha \in \mathbf{A}_0$.

It follows that \mathbf{M} is the required δ -matching. \square

The theorem at infinity. Now suppose $\mathcal{D} \subseteq \overline{\mathbf{R}^2}$ meets any of the strata at infinity. For each of the four lines at infinity, the 3-part proof given above works almost verbatim, if we replace \mathcal{D} with its intersection with the chosen line, and each r-measure μ_x with the corresponding measure at infinity. The other change is to replace the word ‘square’ with the word ‘interval’. The necessary box inequality at infinity is found in Proposition 4.23.

For the four corners $(\pm\infty, \pm\infty)$, it is easier still: the box inequality at each corner implies that μ_0, μ_δ have the same multiplicity there. The interpolating measures are not needed. \square

This completes the proof of the stability theorem for finite measures on an open domain \mathcal{D} , and hence the stability theorem for q-tame persistence modules, and hence the isometry theorem for q-tame persistence modules.

4.8. The measure stability theorem (continued). The stability theorem generalises to measures that are not necessarily finite. By the extension convention, we may suppose that the measures are defined on $\overline{\mathbf{R}^2}$ (rather than just a subset of $\overline{\mathbf{R}^2}$). Given a 1-parameter family $(\mu_x \mid x \in [0, \delta])$, the finite interiors

$$\mathcal{F}_x = \mathcal{F}^\circ(\mu_x)$$

now depend on x ; whereas previously we had $\mathcal{F}_x = \mathcal{D}$ for all x .

For $\mathcal{F} \subset \overline{\mathbf{R}^2}$ an open set and $\delta \geq 0$, the ‘reverse offset’ is the open set

$$\mathcal{F}^{-\delta} = \{\alpha \in \mathcal{F} \mid \text{ex}^\infty(\alpha, \mathcal{F}) > \delta\} = \{\alpha \in \mathcal{F} \mid \alpha^\delta \subset \mathcal{F}\}.$$

Intuitively, this shrinks \mathcal{F} by δ at the boundary. Clearly $\mathcal{F} \supseteq \mathcal{G}$ implies $\mathcal{F}^{-\delta} \supseteq \mathcal{G}^{-\delta}$, and $(\mathcal{F}^{-\delta_1})^{-\delta_2} = \mathcal{F}^{-(\delta_1 + \delta_2)}$. Note also that $(\mathcal{F} \cap \mathcal{G})^{-\delta} = \mathcal{F}^{-\delta} \cap \mathcal{G}^{-\delta}$. This is easiest to see from the second characterisation.

Remark. The operation $[\cdot]^{-\delta}$ has no effect on the corners at infinity, and acts independently on the standard plane and on the four lines at infinity.

We define δ -matchings for multisets in unequal domains. Let \mathcal{F}, \mathcal{G} be open subsets of $\overline{\mathbf{R}^2}$, let A, B be multisets in \mathcal{F}, \mathcal{G} respectively, and let $\delta > 0$. A **δ -matching between $(A, \mathcal{F}), (B, \mathcal{G})$** is a partial matching M between A, B such that the following four conditions hold:

- $\mathcal{F} \supseteq \mathcal{G}^{-\delta}$ and $\mathcal{G} \supseteq \mathcal{F}^{-\delta}$,
- if $(\alpha, \beta) \in M$ then $d^\infty(\alpha, \beta) \leq \delta$,
- every $\alpha \in A \cap \mathcal{G}^{-\delta}$ is matched with some $\beta \in B$,
- every $\beta \in B \cap \mathcal{F}^{-\delta}$ is matched with some $\alpha \in A$.

The first of these is a compatibility condition between the domains: they cannot be too unequal. This is automatic if $\mathcal{F} = \mathcal{G}$, which is why we haven’t seen it before. Notice the cross-over in the last two conditions: a point in A is allowed to be unmatched only if it is close to the boundary of B ’s domain \mathcal{G} , and vice versa.

Proposition 4.27 (triangle inequality). *If A, B, C are multisets in $\mathcal{F}, \mathcal{G}, \mathcal{H}$ respectively, and there exist a δ_1 -matching between $(A, \mathcal{F}), (B, \mathcal{G})$ and a δ_2 -matching between $(B, \mathcal{G}), (C, \mathcal{H})$, then there exists a $(\delta_1 + \delta_2)$ -matching between $(A, \mathcal{F}), (C, \mathcal{H})$.*

Proof. As usual, compose the two partial matchings to get a partial matching \mathbf{M} between \mathbf{A}, \mathbf{C} . Writing $\delta = \delta_1 + \delta_2$, we must check that this is a δ -matching between $(\mathbf{A}, \mathcal{F}), (\mathbf{C}, \mathcal{H})$. For first condition we see that

$$\mathcal{F} \supseteq \mathcal{G}^{-\delta_1} \supseteq (\mathcal{H}^{-\delta_2})^{-\delta_1} = \mathcal{H}^{-\delta} \quad \text{and} \quad \mathcal{H} \supseteq \mathcal{G}^{-\delta_2} \supseteq (\mathcal{F}^{-\delta_1})^{-\delta_2} = \mathcal{F}^{-\delta}.$$

The second condition follows from the triangle inequality for d^∞ . For the third condition, if $\alpha \in \mathbf{A}$ lies in $\mathcal{H}^{-\delta}$ then by the inclusion above it lies in $\mathcal{G}^{-\delta_1}$. Therefore α is matched with $\beta \in \mathbf{B}$. Moreover β must then lie in $\mathcal{H}^{-(\delta-\delta_1)} = \mathcal{H}^{-\delta_2}$ and so is matched with $\gamma \in \mathbf{C}$. The fourth condition follows by symmetry. \square

Remark. There is no triangle inequality if the condition on the domains is dropped.

Here is the main theorem of this section and the final theorem of the paper. Again we use the abbreviation $\mathcal{F}_x = \mathcal{F}^\circ(\mu_x)$ for the finite interiors.

Theorem 4.28 (stability for measures). *Suppose $(\mu_x \mid x \in [0, \delta])$ is a 1-parameter family of r -measures on $\overline{\mathbf{R}^2}$. Suppose for all $x, y \in [0, \delta]$ the box inequality*

$$\mu_x(R) \leq \mu_y(R^{|y-x|})$$

holds for all rectangles $R \in \text{Rect}(\overline{\mathbf{R}^2})$. Then there exists a δ -matching between the undecorated diagrams $(\text{dgm}(\mu_0), \mathcal{F}_0)$ and $(\text{dgm}(\mu_\delta), \mathcal{F}_\delta)$.

The first condition for a δ -matching, on the domains $\mathcal{F}_0, \mathcal{F}_\delta$, can be checked easily:

Proposition 4.29. *Under the hypotheses of Theorem 4.28, we have inclusions*

$$\mathcal{F}_x \supseteq \mathcal{F}_y^{-|y-x|}$$

for all $x, y \in [0, \delta]$.

Proof. Suppose $\alpha \in \mathcal{F}_y^{-|y-x|}$, then equivalently $\alpha^{|y-x|} \subset \mathcal{F}_y$. Since the square $\alpha^{|y-x|}$ is compact and \mathcal{F}_y is open, there exists $\epsilon > 0$ such that $\alpha^{|y-x|+\epsilon} \subset \mathcal{F}_y$. The box inequality gives

$$\mu_x(\alpha^\epsilon) \leq \mu_y(\alpha^{|y-x|+\epsilon})$$

and the right-hand side is finite by Proposition 2.16. Thus $\alpha \in \mathcal{F}^\circ(\mu_x) = \mathcal{F}_x$. \square

Proof of Theorem 4.28. The argument closely follows the proof of the stability theorem for finite measure, so we will confine ourselves to indicating the necessary modifications. We use the abbreviation $\mathbf{A}_x = \text{dgm}(\mu_x)$.

Initial remark. Recall that the proof is conducted separately for each of the nine strata. The four corners at infinity are handled easily (each corner belongs to both \mathcal{F}_0 and \mathcal{F}_δ , or to neither; in the former case the μ_0, μ_δ multiplicities agree). The proof is described for the points in the standard plane. The same proof applies to each of the four lines at infinity, replacing each μ_x with the corresponding measure at infinity.

Part 1. *The Hausdorff distance between $(\mathbf{A}_x, \mathcal{F}_x)$, and $(\mathbf{A}_y, \mathcal{F}_y)$ is at most $\eta = |y - x|$.*

The assertion is understood to mean:

- If $\alpha \in \mathbf{A}_x$ and $\text{ex}^\infty(\alpha, \mathcal{F}_y) > \eta$, then there exists $\beta \in \mathbf{A}_y$ with $d^\infty(\alpha, \beta) \leq \eta$.
- If $\beta \in \mathbf{A}_y$ and $\text{ex}^\infty(\beta, \mathcal{F}_x) > \eta$, then there exists $\alpha \in \mathbf{A}_x$ with $d^\infty(\alpha, \beta) \leq \eta$.

Proof: for all $\epsilon > 0$ with $\eta + \epsilon < \text{ex}^\infty(\alpha, \mathcal{F}_y)$, we have $1 \leq \mu_x(\alpha^\epsilon) \leq \mu_y(\alpha^{\eta+\epsilon})$ so there is at least one point of \mathbf{A}_y in α^η .

Part 2. *The theorem is true if \mathbf{A}_x has finite cardinality for all x .*

Item (i) is given by the triangle inequality (Proposition 4.27).

Item (ii) uses the same strategy as before. Let (α_i) be a finite enumeration of the distinct points of \mathbf{A}_x , with respective multiplicities (n_i) . Then $\rho(x)$ is chosen to satisfy:

$$0 < \rho(x) \leq \begin{cases} \frac{1}{2}\text{ex}^\infty(\alpha_i, \mathcal{F}_x) & \text{all } i \\ \frac{1}{2}d^\infty(\alpha_i, \alpha_j) & \text{all } i, j \text{ distinct} \end{cases}$$

If $\eta = |y - x| < \rho(x)$, then Part 1 implies that \mathbf{A}_y is contained in the disjoint union

$$(\mathbf{R}^2 - \mathcal{F}_x)^\eta \cup \alpha_1^\eta \cup \dots \cup \alpha_k^\eta.$$

The box inequality is then used to count precisely n_i points of \mathbf{A}_y in the square α_i^η . This defines a partial matching where all points of \mathbf{A}_x are matched and all points of $\mathbf{A}_y \cap \mathcal{F}_x^{-\eta}$ are matched.

The formal deduction of Part 2 from (i) and (ii) is unchanged, since it is a formal deduction.

Part 3. *The theorem is true without assuming finite cardinality.*

The idea is to restrict each measure μ_x to a relatively compact open subset $\hat{\mathcal{F}}_x \subset \mathcal{F}_x = \mathcal{F}^\circ(\mu_x)$. The subsets satisfy the compatibility condition

$$\hat{\mathcal{F}}_x \supseteq \hat{\mathcal{F}}_y^{-|y-x|}$$

for all $x, y \in [0, \delta]$.

Specifically, for $\epsilon > 0$ and $r > \delta$, let

$$\hat{\mathcal{F}}_x = \mathcal{F}_x^{-\epsilon} \cap \mathcal{Q}^r$$

where $\mathcal{Q}^r = (-r, r) \times (-r, r)$ is the open d^∞ -disk of radius r . Define a function on rectangles as follows:

$$\hat{\mu}_x(R) = \begin{cases} \mu_x(R) & \text{if } R \subset \hat{\mathcal{F}}_x \\ \infty & \text{otherwise} \end{cases}$$

It is easy to check that $\hat{\mu}_x$ is an r-measure (additivity still holds), that $\mathcal{F}^\circ(\hat{\mu}_x) = \hat{\mathcal{F}}_x$, and that $\text{dgm}(\hat{\mu}_x) = \text{dgm}(\mu_x) \cap \hat{\mathcal{F}}_x$.

Lemma. *The family $(\hat{\mu}_x)$ satisfies the box inequality $\hat{\mu}_x(R) \leq \hat{\mu}_y(R^{|y-x|})$ for all $x, y \in [0, \delta]$.*

Proof. Since the box inequality is assumed to hold for (μ_x) , it will automatically hold for $(\hat{\mu}_x)$; except possibly for rectangles R where the left-hand side of the inequality has become infinite while the right-hand side hasn't. This happens when $R \not\subset \hat{\mathcal{F}}_x$ while $R^{|y-x|} \subset \hat{\mathcal{F}}_y$, and we can prevent it by ensuring that $\hat{\mathcal{F}}_x \supseteq \hat{\mathcal{F}}_y^{-|y-x|}$. And, indeed,

$$\hat{\mathcal{F}}_y^{-|y-x|} = (\mathcal{F}_y^{-\epsilon} \cap \mathcal{Q}^r)^{-|y-x|} = \mathcal{F}_y^{-(\epsilon+|y-x|)} \cap \mathcal{Q}^{r-|y-x|} \subseteq \mathcal{F}_x^{-\epsilon} \cap \mathcal{Q}^r = \hat{\mathcal{F}}_x. \quad \square$$

Since $\hat{\mathcal{F}}_x$ has compact closure in \mathcal{F}_x , and \mathbf{A}_x is locally finite, it follows that $\hat{\mathbf{A}}_x = \mathbf{dgm}(\hat{\mu}_x) = \mathbf{A}_x \cap \hat{\mathcal{F}}_x$ has finite cardinality. We can therefore apply Part 2 to the family $(\hat{\mu}_x)$ to get a δ -matching between $(\hat{\mathbf{A}}_0, \hat{\mathcal{F}}_0)$ and $(\hat{\mathbf{A}}_\delta, \hat{\mathcal{F}}_\delta)$. This can be interpreted as a partial δ -matching between $\mathbf{A}_0, \mathbf{A}_\delta$ where:

- $\alpha \in \mathbf{A}_0$ is matched whenever $\alpha \in (\mathcal{F}_\delta^{-\epsilon} \cap \omega^r)^{-\delta} = \mathcal{F}_\delta^{-(\delta+\epsilon)} \cap \omega^{r-\delta}$
- $\beta \in \mathbf{A}_\delta$ is matched whenever $\beta \in (\mathcal{F}_0^{-\epsilon} \cap \omega^r)^{-\delta} = \mathcal{F}_0^{-(\delta+\epsilon)} \cap \omega^{r-\delta}$

Repeat this argument for a sequence (ϵ_n, r_n) where $\epsilon_n \rightarrow 0$ and $r_n \rightarrow +\infty$. This gives a sequence of δ -matchings \mathbf{M}_n , and we can form a limit \mathbf{M} as before.

If $\alpha \in \mathbf{A}_0 \cap \mathcal{F}_\delta^{-\delta}$ then eventually $\alpha \in \mathcal{F}_\delta^{-(\delta+\epsilon_n)} \cap \omega^{r_n-\delta}$ and so α is matched by \mathbf{M}_n for all sufficiently large n . The same is true for $\beta \in \mathbf{A}_\delta \cap \mathcal{F}_0^{-\delta}$. With this information, we can complete the usual proof that \mathbf{M} is a δ -matching between $(\mathbf{A}_0, \mathcal{F}_0)$ and $(\mathbf{A}_\delta, \mathcal{F}_\delta)$. \square

This concludes the proof of Theorem 4.28. Here is a sample consequence.

Example 4.30 (Stability of the Webb module). Let \mathbb{V} be a persistence module which is δ -interleaved with the module \mathbb{W} of Example 2.20. By interpolation and the box lemma, the measure stability theorem applies here. We get a δ -matching between $(\mathbf{dgm}(\mu_{\mathbb{V}}), \mathcal{F}^\circ(\mu_{\mathbb{V}}))$ and $(\mathbf{dgm}(\mu_{\mathbb{W}}), \mathcal{F}^\circ(\mu_{\mathbb{W}}))$. This amounts to the following.

- In the finite part \mathcal{H} of the half-plane:

The singular support of $\mu_{\mathbb{V}}$ is contained in the diagonal strip $\Delta_{[0,\delta]}$. Each point of $\mathbf{dgm}(\mu_{\mathbb{V}})$ outside this strip is matched with some point $(-n, 0) \in \mathbf{dgm}(\mu_{\mathbb{W}})$. Conversely, the only unmatched points of $\mathbf{dgm}(\mu_{\mathbb{W}})$ must lie within distance δ of the diagonal or of the singular support of $\mu_{\mathbb{V}}$.

In particular, if $\delta < \frac{1}{4}$ then all points of $\mathbf{dgm}(\mu_{\mathbb{W}})$ are matched.

- On the line $(-\infty, \mathbf{R})$: All points and singularities of $\mu_{\mathbb{V}}$ are contained in the interval $(-\infty, [-\delta, +\delta])$. There is at least one singular point.
- On $(\mathbf{R}, +\infty)$ and at $(-\infty, +\infty)$: The measure $\mu_{\mathbb{V}}$ has no points or singularities.

5. EXAMPLES

5.1. Partial interleavings. In some practical data analysis situations, one considers persistence modules which are only partially interleaved. One such scenario is presented by

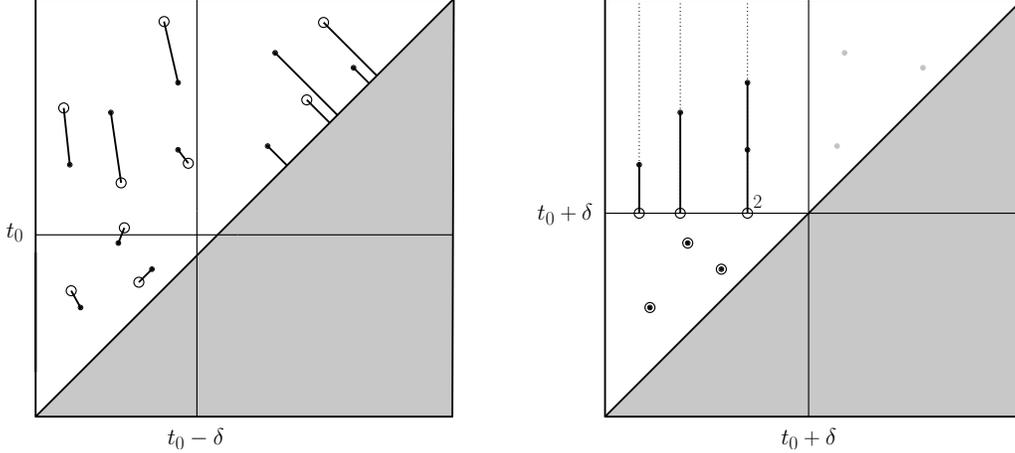


FIGURE 12. Left: the partial matching of Theorem 5.1 between $\mathbf{dgm}(\mathbb{U})$ (\bullet) and $\mathbf{dgm}(\mathbb{V})$ (\circ). Right: projection from $\mathbf{dgm}(\mathbb{U})$ (\bullet and \bullet) to $\mathbf{dgm}(\tilde{\mathbb{U}})$ (\circ). The grey points are the ones that disappear.

Chazal et al. in the context of clustering by mode-seeking [9]. A filtered simplicial complex on an input point cloud is compared with the sublevelset filtration of the density function it was sampled from. In low-density regions, the sample is too sparse to expect there to be an interleaving. Nevertheless, there is interleaving when the density is sufficiently high.

This leads to the following notion of partial interleaving, adapted from [9]. Two persistence modules \mathbb{U} and \mathbb{V} are said to be δ -**interleaved up to time** t_0 if there are maps $\phi_t : U_t \rightarrow V_{t+\delta}$ and $\psi_t : V_t \rightarrow U_{t+\delta}$ defined for all $t \leq t_0$, such that the diagrams of equation 3.1 commute for all values $s < t \leq t_0$; that is, for all values where the maps are defined.

A weaker version of the stability theorem, illustrated in Figure 12 (left), can be proven:

Theorem 5.1 (from [9]). *Let \mathbb{U} and \mathbb{V} be two q -tame persistence modules that are δ -interleaved up to time t_0 . Then, there is a partial matching $\mathbf{M} \subset \mathbf{dgm}(\mathbb{U}) \times \mathbf{dgm}(\mathbb{V})$ with the following properties:*

- *Points (p, q) in either diagram for which $\frac{1}{2}|p - q| \leq \delta$ are not required to be matched.*
- *Points (p, q) in either diagram for which $p \geq t_0 - \delta$ are not required to be matched.*

All other points must be matched. Then:

- *If α, β are matched, then the p -coordinates of α, β differ by at most δ .*
- *If α, β are matched and one of α, β lies below the line $q = t_0$, then $d^\infty(\alpha, \beta) \leq \delta$.*

For the proof, we introduce two new persistence modules $\tilde{\mathbb{U}}, \tilde{\mathbb{V}}$.

$$\begin{aligned} \tilde{U}_t &= U_t & \text{if } t \leq t_0 + \delta & & \text{and} & & \tilde{U}_t = 0 & \text{otherwise} \\ \tilde{V}_t &= V_t & \text{if } t \leq t_0 + \delta & & \text{and} & & \tilde{V}_t = 0 & \text{otherwise} \end{aligned}$$

with maps

$$\begin{aligned} \tilde{u}_s^t &= u_s^t & \text{if } t \leq t_0 + \delta & & \text{and} & & \tilde{u}_s^t = 0 & \text{otherwise} \\ \tilde{v}_s^t &= v_s^t & \text{if } t \leq t_0 + \delta & & \text{and} & & \tilde{v}_s^t = 0 & \text{otherwise} \end{aligned}$$

for all $s \leq t$. We may call $\hat{\mathbb{U}}, \hat{\mathbb{V}}$ the **truncations** of \mathbb{U}, \mathbb{V} to $(-\infty, T]$, where $T = t_0 + \delta$.

Proof. There are three steps.

Step 1. The decorated diagram of a persistence module \mathbb{U} determines the decorated diagram of its truncation $\tilde{\mathbb{U}}$, in a straightforward way. Specifically, transform each point $(p^*, q^*) \in \mathbf{Dgm}(\mathbb{U})$ as follows:

$$(\S) \quad (p^*, q^*) \mapsto \begin{cases} (p^*, q^*) & \text{if } q^* < T \\ (p^*, T^+) & \text{if } p^* < T < q^* \\ \text{disappears} & \text{if } T < p^* \end{cases}$$

Then $\mathbf{Dgm}(\tilde{\mathbb{U}})$ is the result of this transformation. The consequent relationship between the undecorated diagrams is illustrated in Figure 12 (right).

Step 2. If \mathbb{U}, \mathbb{V} are δ -interleaved up to time t_0 , then $\tilde{\mathbb{U}}, \tilde{\mathbb{V}}$ are δ -interleaved.

Combining the first two steps we get the third.

Step 3. The stability theorem gives a δ -matching between $\mathbf{dgm}(\tilde{\mathbb{U}}), \mathbf{dgm}(\tilde{\mathbb{V}})$. This lifts to a matching between $\mathbf{dgm}(\mathbb{U}), \mathbf{dgm}(\mathbb{V})$ which has the properties stated in the theorem.

The second and third steps are straightforward. Only the first step requires any technical input, intuitively plausible as it may be. The framework developed in [6] leads to a 2-page argument, presented in the appendix of [9].

Here is a shorter proof. Write $\mu = \mu_{\mathbb{U}}$ and $\tilde{\mu} = \mu_{\tilde{\mathbb{U}}}$. Let \mathbf{A} denote the multiset obtained from $\mathbf{Dgm}(\mathbb{U})$ by applying the transformation (\S). Consider an arbitrary rectangle $[a, b] \times [c, d] \in \text{Rect}(\overline{\mathcal{H}})$. We easily see:

$$\text{card}(\mathbf{A}|_{[a,b] \times [c,d]}) = \begin{cases} \mu([a, b] \times [c, d]) & \text{if } d \leq T \\ \mu([a, b] \times [c, +\infty]) & \text{if } c \leq T < d \\ 0 & \text{if } T < c \end{cases}$$

To show that we have correctly determined $\mathbf{Dgm}(\tilde{\mathbb{U}})$, it suffices to show that $\text{card}(\mathbf{A}|_{[a,b] \times [c,d]}) = \tilde{\mu}([a, b] \times [c, d])$ for all rectangles. And indeed:

- If $d \leq T$, then:

$$\begin{aligned} \tilde{\mu}([a, b] \times [c, d]) &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \tilde{\mathbb{U}} \rangle \\ &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \mathbb{U} \rangle = \mu([a, b] \times [c, d]) \end{aligned}$$

- If $c \leq T < d$, then:

$$\begin{aligned}\tilde{\mu}([a, b] \times [c, d]) &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \tilde{\mathbb{U}} \rangle \\ &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \text{---} \mid \mathbb{U} \rangle = \mu([a, b] \times [c, +\infty])\end{aligned}$$

since $\tilde{U}_d = 0$.

- If $T < c$, then:

$$\tilde{\mu}([a, b] \times [c, d]) = \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \mid \tilde{\mathbb{U}} \rangle = 0$$

since $\tilde{U}_c = 0$.

It follows that $\text{Dgm}(\tilde{\mathbb{U}}) = \mathbf{A}$ as claimed. □

5.2. Extended persistence. Cohen-Steiner, Edelsbrunner and Harer [11] introduced extended persistence to capture the homological information carried by a pair (X, f) . Some but not all of this information is recovered by the sublevelset persistence $H(\mathbb{X}_{\text{sub}})$. The idea is to grow the space from the bottom up, through sublevelsets; and then to relativise the space from the top down, with superlevelsets. Extended persistence is the persistent homology of this sequence of spaces and pairs.

It is usually assumed that (X, f) has finitely many homological critical points (a_i) . One applies a homology functor to the finite sequence⁷

$$X^{a_0} \rightarrow X^{a_1} \rightarrow \cdots \rightarrow X^{a_{n-1}} \rightarrow X \rightarrow (X, X_{a_n}) \rightarrow \cdots \rightarrow (X, X_{a_2}) \rightarrow (X, X_{a_1})$$

to get a quiver representation. The indecomposable summands of this representation are interpreted as features, and are drawn as points in the ‘extended persistence diagram’. There are three kinds of feature:

- ordinary features (which are born and die before the central X);
- relative features (which are born and die after the central X);
- extended features (which are born before the X and die after it).

The finiteness assumption is satisfied when (X, f) is a compact manifold with a Morse function, or a compact polyhedron with a piecewise-linear map. In the former situation, there are extra symmetries (Poincaré, Lefschetz) which are explored in [11].

In fact, it is perfectly straightforward to define the extended persistence diagram under a weaker hypothesis. Suppose X is a compact polyhedron and f is a continuous real-valued function on X . Then:

- $\text{rank}(H(X^s) \rightarrow H(X^t)) < \infty$ whenever $s < t$; and
- $\text{rank}(H(X, X_s) \rightarrow H(X, X_t)) < \infty$ whenever $s > t$.

⁷We write $X^t = (X, f)^t = f^{-1}(-\infty, t]$ and $X_t = (X, f)_t = f^{-1}[t, +\infty)$ for sublevelsets and superlevelsets.

The first of these facts is Theorem 2.22. The second is proved similarly, by factorising the map through some $H(X, Y)$, where Y is a subpolyhedron of X nested between X_s, X_t .

Define the ordered set

$$\mathbf{J} = \{\underline{t} \mid t \in \mathbf{R}\} \quad \text{ordered by} \quad \underline{s} \leq \underline{t} \Leftrightarrow s \geq t,$$

thought of as a ‘backwards’ copy of the real line, with bars under numbers to remind us. For extended persistence we may work with the set

$$\mathbf{R}_{\text{EP}} = \mathbf{R} \cup \{+\infty\} \cup \mathbf{J}$$

with the ordering $s < +\infty < \underline{t}$ for all s, \underline{t} .

The extended persistence module $\mathbb{X}_{\text{EP}} = \mathbb{X}_{\text{EP}}^f$ for (X, f) is defined as follows:

$$\begin{aligned} V_t &= H(X^t) & \text{for } t \in \mathbf{R} \\ V_{+\infty} &= H(X) \\ V_{\underline{t}} &= H(X, X_t) & \text{for } \underline{t} \in \mathbf{J} \end{aligned}$$

Since \mathbf{R}_{EP} is order-isomorphic to the real line, we may interpret \mathbb{X}_{EP} it as a persistence module over \mathbf{R} . The two facts cited above imply that it is q-tame, so the decorated diagram is defined away from the diagonal.

Alternatively, we can define the extended persistence diagram in three pieces:

$$\begin{aligned} \mu_{\text{ord}}([a, b] \times [c, d]) &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_c \text{---} \circ_d \rangle & \text{for } -\infty \leq a < b \leq c < d \leq +\infty \\ \mu_{\text{rel}}([\underline{a}, \underline{b}] \times [\underline{c}, \underline{d}]) &= \langle \circ_{\underline{a}} \text{---} \bullet_{\underline{b}} \text{---} \bullet_{\underline{c}} \text{---} \circ_{\underline{d}} \rangle & \text{for } +\infty \leq \underline{a} < \underline{b} \leq \underline{c} < \underline{d} \leq -\infty \\ \mu_{\text{ext}}([a, b] \times [\underline{c}, \underline{d}]) &= \langle \circ_a \text{---} \bullet_b \text{---} \bullet_{\underline{c}} \text{---} \circ_{\underline{d}} \rangle & \text{for } -\infty \leq a < b \leq +\infty \text{ and } +\infty \leq \underline{c} < \underline{d} \leq -\infty \end{aligned}$$

taking $V_{-\infty} = 0$ and $V_{-\infty} = 0$ whenever needed.

The measures $\mu_{\text{ord}}, \mu_{\text{rel}}$ are defined over the half-plane $\overline{\mathcal{H}}$, whereas μ_{ext} is defined over $\overline{\mathbf{R}^2}$.

Stability for $\text{dgm}_{\text{ord}}, \text{dgm}_{\text{rel}}$ and dgm_{ext} is proved individually for each diagram. Given two functions f, g which are δ -close in the supremum norm, there are inclusions

$$\begin{aligned} (X, f)^t &\subseteq (X, g)^{t+\delta} & (X, f)_t &\subseteq (X, g)_{t-\delta} \\ (X, g)^t &\subseteq (X, f)^{t+\delta} & (X, g)_t &\subseteq (X, f)_{t-\delta} \end{aligned}$$

using which we can prove the box lemma for each measure. Since linear combinations of continuous functions are continuous, we can interpolate between f and g to satisfy the hypotheses required by the measure stability theorem.

Remark. In the spirit of Theorem 2.23, one may treat the case where X is a locally compact polyhedron and f is proper. We leave it as an exercise for the sufficiently persistent reader to carry this out and locate the potential singularities of the measures.

ACKNOWLEDGEMENTS

Various people helped improve this paper. We deeply appreciate the intellectual influence of David Cohen-Steiner and Leo Guibas, as co-authors of [7], and Gunnar Carlsson, as co-author of [4]. We thank William Crawley-Boevey for some invaluable correspondence regarding the decomposition of persistence modules. Michael Lesnick generously made available early versions of [16]; among other consequences, we learned of the paper of Webb [19].

The authors have been supported by the following research grants.

- The Digiteo project C3TTA and the Digiteo Chair (held by the second author) which made this collaboration possible.
- European project CG-Learning EC: contract 255827.
- ANR GIGA: contract ANR-09-BLAN-0331-01.
- DARPA project Sensor Topology and Minimal Planning (SToMP): HR0011-07-1-0002.

We thank our home institutions, for their ongoing support: INRIA Saclay – Ile-de-France (FC, MG, SO) and Pomona College (VdS).

REFERENCES

- [1] Maurice Auslander. Representation theory of artin algebras ii. *Communications in Algebra*, 1(4):269–310, 1974.
- [2] Gorô Azumaya. Corrections and supplementaries to my paper concerning Krull–Remak–Schmidt’s theorem. *Nagoya Mathematical Journal*, 1:117–124, 1950.
- [3] Peter Bubenik and Jonathan A. Scott. Categorification of persistent homology. arXiv:1205.3669v1, 2012.
- [4] Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, aug 2010.
- [5] Gunnar Carlsson, Vin de Silva, and Dmitriy Morozov. Zigzag persistent homology and real-valued functions. In *Proceedings 25th ACM Symposium on Computational Geometry (SoCG)*, pages 247–256, 2009.
- [6] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leo Guibas, and Steve Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th Annual ACM Symposium on Computational Geometry (SoCG)*, pages 237–246, 2009.
- [7] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Oudot. Proximity of Persistence Modules and their Diagrams. Research Report RR-6568, INRIA, 2008.
- [8] Frédéric Chazal, Vin de Silva, and Steve Oudot. Persistence stability of geometric complexes. (in progress).
- [9] Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Persistence-based clustering in Riemannian manifolds. Under review. Short version published in the Proceedings of the 27th Annual ACM Symposium on Computational Geometry, Paris, France, June 2011.
- [10] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
- [11] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Foundations of Computational Mathematics*, 2008.
- [12] Harm Derksen and Jerzy Weyman. Quiver representations. *Notices of the American Mathematical Society*, 52(2):200–206, February 2005.

- [13] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.
- [14] Patrizio Frosini. Discrete computation of size functions. *Journal of Combinatorics, Information & System Sciences*, 17(3–4):232–250, 1992.
- [15] P. Gabriel. Unzerlegbare Darstellungen I. *Manuscripta Mathematica*, 6:71–103, 1972.
- [16] Michael Lesnick. The optimality of the interleaving distance on multidimensional persistence modules. arXiv:1106.5305.
- [17] Vanessa Robins. Towards computing homology from finite approximations. *Topology Proceedings*, 24:503–532, 1999.
- [18] H. Tachikawa and Claus Michael Ringel. Qf-3 rings. *Journal für die reine und angewandte Mathematik (Crelle's Journal)*, 1975(272):49–72, 2011/09/17 1975.
- [19] Cary Webb. Decomposition of graded modules. *Proceedings of the American Mathematical Society*, 94(4):565–571, 1985.
- [20] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33(2):249–274, 2005.