Chapter 5
The Isometry Theorem

In this section we discuss the metric relationship between persistence modules and their persistence diagrams. As in Chap. 4, all persistence modules are indexed by **R** unless explicitly stated otherwise.

The principal result is the famous stability theorem of Cohen-Steiner, Edelsbrunner and Harer [19], in the generality established by [15]. The main difference is that we emphasise persistence measures, deriving the standard theorem from a more general statement about measures. The structure of the proof remains the same as in [19].

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 [42] for modules which satisfy $\dim(V_t) < \infty$ for all t, and independently in the work of Bubenik and Scott [5] for modules of finite type.

5.1 The Interleaving Distance

In this section we define the interleaving distance between persistence modules. This was introduced in [15].

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

$$\Phi' = \Phi 1_{\mathbb{U}}^{\varepsilon} = 1_{\mathbb{V}}^{\varepsilon}\Phi$$
$$\Psi' = \Psi 1_{\mathbb{V}}^{\varepsilon} = 1_{\mathbb{U}}^{\varepsilon}\Psi$$

provide the required interleaving.

F. Chazal et al., *The Structure and Stability of Persistence Modules*,
SpringerBriefs in Mathematics, DOI 10.1007/978-3-319-42545-0_5

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 $(\delta + \varepsilon)$-interleaved for all $\varepsilon > 0$. This does not imply that $\mathbb{U}, \mathbb{V}$ are δ-interleaved, as we see now:

Example 5.1 Two persistence modules are 0-interleaved if and only if they are isomorphic.

Example 5.2 A persistence module $\mathbb{V}$ is **ephemeral** if $v_t^s = 0$ for all $s < t$. An ephemeral module may be constructed by selecting an arbitrary family of spaces (V_t) and setting every v_t^s to be zero. Let $\mathbb{U}$ and $\mathbb{V}$ be a pair of non-isomorphic ephemeral modules. Then $\mathbb{U}, \mathbb{V}$ are 0^+-interleaved but not 0-interleaved. Indeed, $1_{\mathbb{U}}^{2\varepsilon} = 0$ and $1_{\mathbb{V}}^{2\varepsilon} = 0$ for all $\varepsilon > 0$, so the zero maps

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

constitute an ε-interleaving.

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

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

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

Proposition 5.3 *The interleaving distance satisfies the triangle inequality:*

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

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:

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

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

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

Now take the infimum over δ_1, δ_2. □

The proposition tells us that $\mathrm{d_i}$ is an extended pseudometric. It is not a true metric because $\mathrm{d_i}(\mathbb{U}, \mathbb{V}) = 0$ 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 straightforward proof is left to the reader (or see Proposition 5.6).

Example 5.4 The four interval modules

$$\mathbf{k}[p,q], \;\; \mathbf{k}[p,q), \;\; \mathbf{k}(p,q], \;\; \mathbf{k}(p,q)$$

are 0^+-interleaved but not isomorphic.

The following property of interleaving distance will be useful later.

Proposition 5.5 *Let* $\mathbb{U}_1, \mathbb{U}_2, \mathbb{V}_1, \mathbb{V}_2$ *be persistence modules. Then*

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

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

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

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 $\mathrm{d_i}(\mathbb{U}_\ell, \mathbb{V}_\ell)$ is an upper bound for $\mathrm{d_i}(\mathbb{U}, \mathbb{V})$. In particular, this is true for the least upper bound, or sup. □

5.2 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

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

is a multiset in the extended open half-plane

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

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:

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

Points at infinity are compared in the expected way:

$$\mathrm{d}^{\infty}((-\infty, q), (-\infty, s)) = |q - s|,$$
$$\mathrm{d}^{\infty}((p, +\infty), (r, +\infty)) = |p - r|,$$

and

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

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

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

be the corresponding interval modules. Then

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

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:

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

Again this is related to the behaviour of interval modules.

Proposition 5.7 *Let* (p^*, q^*) *be an interval, let*

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

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

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

(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 $\delta < \frac{1}{2}(q - p)$. □

We now use these two concepts to define the bottleneck distance between two multisets A, 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 $\alpha_1, \ldots, \alpha_k$. Henceforth we will do this implicitly, without comment.

A **partial matching** between A and B is a collection of pairs

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

such that:

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

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

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

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

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

In Sect. 5.3, we will show that 'inf' can be replaced by 'min' if A, B are locally finite.

Remark 5.8 In order for $d_b(\mathsf{A}, \mathsf{B}) < \infty$, it is necessary that the cardinalities of A, B agree over each of the three strata at infinity:

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

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 5.9 *The bottleneck distance satisfies the triangle inequality:*

$$\mathrm{d_b}(\mathsf{A}, \mathsf{C}) \le \mathrm{d_b}(\mathsf{A}, \mathsf{B}) + \mathrm{d_b}(\mathsf{B}, \mathsf{C})$$

for any three multisets $\mathsf{A}, \mathsf{B}, \mathsf{C}$.

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 $\mathsf{M}_1, \mathsf{M}_2$ to be

$$\mathsf{M} = ((\alpha, \gamma) \mid \text{there exists } \beta \in \mathsf{B} \text{ such that } (\alpha, \beta) \in \mathsf{M}_1 \text{ and } (\beta, \gamma) \in \mathsf{M}_2)\,.$$

It is a partial matching because $\mathsf{M}_1, \mathsf{M}_2$ are partial matchings. We verify that M is the required δ-matching:

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

  $$\mathrm{d}^\infty(\alpha, \gamma) \le \mathrm{d}^\infty(\alpha, \beta) + \mathrm{d}^\infty(\beta, \gamma) \le \delta_1 + \delta_2 = \delta$$

  where $\beta \in \mathsf{B}$ 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

  $$\mathrm{d}^\infty(\alpha, \Delta) \le \delta_1 \le \delta.$$

  Or α is matched in M_1, let's say $(\alpha, \beta) \in \mathsf{M}_1$. Then β must be unmatched in M_2, so

  $$\mathrm{d}^\infty(\alpha, \Delta) \le \mathrm{d}^\infty(\alpha, \beta) + \mathrm{d}^\infty(\beta, \Delta) \le \delta_1 + \delta_2 = \delta.$$

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

  $$\mathrm{d}^\infty(\gamma, \Delta) \le \delta.$$

This completes the proof. □

Remark 5.10 Because $\mathsf{A}, \mathsf{B}, \mathsf{C}$ 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 5.1 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.

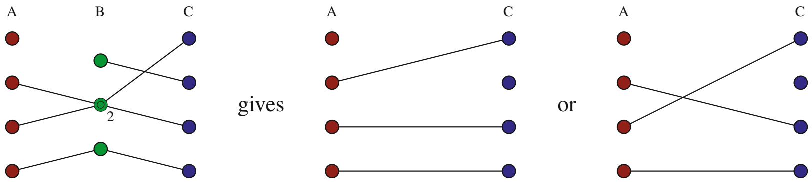

Fig. 5.1 The partial matchings between A, B and B, C (*left*) compose to give a partial matching between A, C in two different ways (*middle, right*). The non-uniqueness arises from the point of multiplicity 2 in B

Theorem 5.11 *Let $\mathbb{U}, \mathbb{V}$ be decomposable persistence modules (i.e. direct sums of interval modules). Then*

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

(We remind the reader that tameness is not required to define Dgm and dgm in this case: see Sect. 2.6.)

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

Let 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 M a partial matching between the interval summands of $\mathbb{U}$ and $\mathbb{V}$.

Rewrite $\mathbb{U}$ 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;
- $\mathbb{U}_\ell$ is an unmatched interval, $\mathbb{V}_\ell = 0$;
- $\mathbb{V}_\ell$ is an unmatched interval, $\mathbb{U}_\ell = 0$.

In each case, by Propositions 5.6 and 5.7, we have $\mathrm{d_i}(\mathbb{U}_\ell, \mathbb{V}_\ell) \leq \delta$. It follows from Proposition 5.5 that $\mathrm{d_i}(\mathbb{U}, \mathbb{V}) \leq \delta$. □

We complete this section with the postponed proof.

Proof (*Proposition* 5.6) We treat the case where p, q, r, s are all finite. We must show that if

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

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

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

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

$$\begin{tikzcd} U_t \arrow[rr] \arrow[dr] & & U_{t+\eta} \arrow[dr] & \\ & V_{t+\delta} \arrow[rr] & & V_{t+\eta+\delta} \end{tikzcd}$$

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:

$$\begin{tikzcd} \bullet \arrow[r] \arrow[dr] & \bullet \arrow[dr] & \\ & \circ \arrow[r] & \bullet \end{tikzcd} \quad \text{or} \quad \begin{tikzcd} \bullet \arrow[r] \arrow[dr] & \circ \arrow[dr] & \\ & \bullet \arrow[r] & \bullet \end{tikzcd}$$

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

$$p \le t \quad \text{and} \quad t + \delta \le r$$

which would imply $\delta \le r - t \le r - p$. This contradicts the hypothesis $\delta > r - p$. For the second situation to occur, one must have

$$q \le t + \eta \quad \text{and} \quad t + \eta + \delta \le s$$

which would imply $\delta \le s - t - \eta \le s - q$. This contradicts the hypothesis $\delta > s - q$. 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

$$\begin{tikzcd} U_t \arrow[rr] \arrow[dr] & & U_{t+2\delta} \\ & V_{t+\delta} \arrow[ur] & \end{tikzcd}$$

commutes for all t. This time the unique forbidden configuration is

$$\begin{tikzcd} \bullet \arrow[rr] \arrow[dr] & & \bullet \\ & \circ \arrow[ur] & \end{tikzcd}$$

and if this occurs then the top row implies

$$p \leq t \quad \text{and} \quad t + 2\delta \leq q.$$

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

$$r < p + \delta \leq t + \delta \leq q - \delta < s$$

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

5.3 The Bottleneck Distance (Continued)

If A, B are locally finite, it turns out that the 'inf' is attained in the definition

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

and can be replaced by 'min'. This will allow us to make a tighter statement of the stability theorem (5.14′) for q-tame modules. See Theorem 5.23.

Theorem 5.12 *Let* A, 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* A, B. *Then there exists a* δ*-matching between* A, B.

The assertion is obvious if A, B are finite. The general case is proved using a compactness argument, as follows.

Proof As usual we treat A, B as sets rather than multisets.

For every integer $n \geq 1$, let M_n be a $(\delta + \frac{1}{n})$-matching between A, B. The plan is to construct a δ-matching M from the sequence (M_n). In practice, we work with the indicator functions

$$\chi : \mathsf{A} \times \mathsf{B} \to \{0, 1\}$$
$$\chi_n : \mathsf{A} \times \mathsf{B} \to \{0, 1\}$$

of the partial matchings M, 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 $\mathsf{A} \times \mathsf{B}$. 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 $\mathbf{N}_\ell$ 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 $\mathbf{N}_\ell$. (If both, then either will do.) Repeat.

Lemma 5.13 *If* F *is any finite subset of* $\mathsf{A} \times \mathsf{B}$*, then there exists* $\ell \geq 1$ *such that*

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta)$$

for all $(\alpha, \beta) \in \mathsf{F}$ *and for all* $n \in \mathbf{N}_\ell$*.*

Proof Indeed, select ℓ such that $(\alpha_1, \beta_1), \ldots, (\alpha_\ell, \beta_\ell)$ include all of F. $\square$

The second step is to verify that χ is the indicator function of a δ-matching. There are several items to check.

- For $\alpha \in \mathsf{A}$ there is at most one $\beta \in \mathsf{B}$ such that $\chi(\alpha, \beta) = 1$.
  Proof Suppose $\chi(\alpha, \beta) = \chi(\alpha, \beta') = 1$ for distinct elements $\beta, \beta' \in \mathsf{B}$. 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. $\square$
- For $\alpha \in \mathsf{A}$ with $\mathrm{d}^\infty(\alpha, \Delta) > \delta$, there is at least one $\beta \in \mathsf{B}$ such that $\chi(\alpha, \beta) = 1$.

  Proof Select N such that $\mathrm{d}^\infty(\alpha, \Delta) > \delta + \frac{1}{N}$. Then the set

$$\mathsf{F}_\alpha = \left\{\beta \in \mathsf{B} \mid \mathrm{d}^\infty(\alpha, \beta) \leq \delta + \tfrac{1}{N}\right\}$$

is finite, since B is locally finite and these points lie in a square bounded away from the diagonal. By Lemma 5.13, there exists ℓ such that

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta)$$

for all $\beta \in \mathsf{F}_\alpha$ and for all $n \in \mathbf{N}_\ell$. On the other hand, if $n \geq N$, then M_n matches α with some $\beta \in \mathsf{F}_\alpha$. Combining these observations,

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta) = 1$$

for sufficiently large $n \in \mathbf{N}_\ell$ and for some $\beta \in \mathsf{F}_\alpha$. □

By symmetry we have:

- For $\beta \in \mathsf{B}$ there is at most one $\alpha \in \mathsf{A}$ such that $\chi(\alpha, \beta) = 1$.
- For $\beta \in \mathsf{B}$ with $\mathrm{d}^\infty(\beta, \Delta) > \delta$, there is at least one $\alpha \in \mathsf{A}$ such that $\chi(\alpha, \beta) = 1$.

Finally:

- If $\chi(\alpha, \beta) = 1$ then $\mathrm{d}^\infty(\alpha, \beta) \leq \delta$.

Proof By Lemma 5.13, there are infinitely many n for which $\chi_n(\alpha, \beta) = 1$. Then

$$\mathrm{d}^\infty(\alpha, \beta) \leq \delta + \tfrac{1}{n}$$

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. □

Although we have chosen to spell out a direct argument, in fact Theorem 5.12 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 to the interested reader.

5.4 The Isometry Theorem

Having defined the interleaving distance and the bottleneck distance, we can now state the main theorem.

Theorem 5.14 *Let* $\mathbb{U}, \mathbb{V}$ *be q-tame persistence modules. Then*

$$\mathrm{d_i}(\mathbb{U}, \mathbb{V}) = \mathrm{d_b}(\mathsf{dgm}(\mathbb{U}), \mathsf{dgm}(\mathbb{V}))$$

(Recall that dgm *denotes the undecorated persistence diagram).*

The result naturally falls into two parts: the 'stability theorem' [19], [14]

$$\mathrm{d_i}(\mathbb{U}, \mathbb{V}) \geq \mathrm{d_b}(\mathsf{dgm}(\mathbb{U}), \mathsf{dgm}(\mathbb{V})), \tag{5.14$'$}$$

and the 'converse stability theorem' [42]

$$\mathrm{d_i}(\mathbb{U}, \mathbb{V}) \leq \mathrm{d_b}(\mathsf{dgm}(\mathbb{U}), \mathsf{dgm}(\mathbb{V})). \tag{5.14$''$}$$

The proof of the converse stability theorem (5.14″) occupies Sect. 5.5. We have already seen the result for decomposable modules, in Theorem 5.11, 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 (5.14′) is given in Sects. 5.6 and 5.7.

5.5 The Converse Stability Theorem

In this section we deduce the converse stability inequality (5.14″) for q-tame modules from Theorem 5.11. A similar argument was given by Lesnick [42], to whom the theorem is due. The key idea is that persistence modules can be approximated by better-behaved persistence modules, using a procedure we call 'smoothing'.

Definition 5.15 Let $\mathbb{V}$ be a persistence module, and let $\varepsilon > 0$. The ε **-smoothing** of $\mathbb{V}$ is the persistence module $\mathbb{V}^\varepsilon$ defined to be the image of the map

$$1_{\mathbb{V}}^{2\varepsilon} : \mathbb{V}[-\varepsilon] \to \mathbb{V}[\varepsilon]$$

(using the 'shift' notation from Remark 4.1). Thus $(V^\varepsilon)_t$ is the image of the map

$$v_{t+\varepsilon}^{t-\varepsilon} : V_{t-\varepsilon} \to V_{t+\varepsilon},$$

and $(v^\varepsilon)_t^s$ is the restriction of $v_{t+\varepsilon}^{s+\varepsilon}$.

Then we have a factorisation of $1_{\mathbb{V}}^{2\varepsilon}$

$$\mathbb{V}[-\varepsilon] \longrightarrow \mathbb{V}^\varepsilon \longrightarrow \mathbb{V}[\varepsilon] \tag{5.1}$$

where the first map is surjective and the second map is injective. At a given index t this is the sequence:

$$V_{t-\varepsilon} \xrightarrow{v_{t-\varepsilon}^{t+\varepsilon}} V_t^\varepsilon \xrightarrow{1} V_{t+\varepsilon}$$

Proposition 5.16 *Let $\mathbb{V}$ be a persistence module. Then* $\mathrm{d_i}(\mathbb{V}, \mathbb{V}^\varepsilon) \leq \varepsilon$.

Proof One checks that the maps in (5.1) constitute an ε-interleaving. □

Smoothing changes the persistence diagram in a predictable way. Here is the atomic example (which the reader can easily verify):

Example 5.17 Let $\mathbb{V} = \mathbf{k}(p^*, q^*)$. Then:

$$\mathbb{V}^\varepsilon = \begin{cases} \mathbf{k}((p+\varepsilon)^*, (q-\varepsilon)^*) & \text{if } (p+\varepsilon)^* < (q-\varepsilon)^* \\ 0 & \text{otherwise} \end{cases}$$

In other words, ε-smoothing shrinks the interval by ε at both ends. □

Proposition 5.18 *The persistence diagram of* $\mathbb{V}^\varepsilon$ *is obtained from the persistence diagram of* $\mathbb{V}$ *by applying the translation* $T_\varepsilon : (p, q) \mapsto (p+\varepsilon, q-\varepsilon)$ *to the part of the extended half-plane that lies above the line* $\Delta_\varepsilon = \{(t-\varepsilon, t+\varepsilon) \mid t \in \mathbf{R}\}$.

In other words, the entire diagram is pushed towards the diagonal by the translation vector $(\varepsilon, -\varepsilon)$. Information in $\mathsf{Dgm}(\mathbb{V})$ that lies below the line Δ_ε is lost and cannot be retrieved from $\mathsf{Dgm}(\mathbb{V}^\varepsilon)$.

In the case where 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 Dgm is defined, is shifted by T_ε.

Proof We consider three different cases. Case (ii) is subsumed by case (iii), but we include it because it makes the proof 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:

$$\left[\bigoplus_{\ell \in L} \mathbb{V}_\ell\right]^\varepsilon = \bigoplus_{\ell \in L} \mathbb{V}_\ell^\varepsilon$$

By Example 5.17, the proposition is true for interval modules. It is therefore true for direct sums of interval modules.

(ii) $\mathbb{V}$ **is q-tame.** It is enough to show that the rank function of $\mathbb{V}^\varepsilon$ 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:

$$\operatorname{rank}[(V^\varepsilon)_b \to (V^\varepsilon)_c] = \operatorname{rank}[V_{b-\varepsilon} \to V_{c+\varepsilon}]$$

In fact, these maps are related by the sequence

$$V_{b-\varepsilon} \longrightarrow (V^\varepsilon)_b \longrightarrow (V^\varepsilon)_c \longrightarrow V_{c+\varepsilon}$$

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 $\mathbb{V}^{\varepsilon}$ is equal to the persistence measure of $\mathbb{V}$ shifted by T_{ε}. Writing

$$A = a - \varepsilon, \quad B = b - \varepsilon, \quad C = c + \varepsilon, \quad D = d + \varepsilon$$

this means showing that

$$\langle \circ_A — \bullet_B — \bullet_C — \circ_D \mid \mathbb{V} \rangle = \langle \circ_a — \bullet_b — \bullet_c — \circ_d \mid \mathbb{V}^{\varepsilon} \rangle$$

for all $a < b \leq c < d$.

The proof is based on the following commutative diagram

$$\begin{tikzcd}
 & (V^{\varepsilon})_a \arrow[rr] & & (V^{\varepsilon})_b \arrow[r] & (V^{\varepsilon})_c \arrow[rr] \arrow[dr] & & (V^{\varepsilon})_d \arrow[dr] & \\
V_A \arrow[ur] \arrow[rr] & & V_B \arrow[ur] \arrow[rrr] & & & V_C \arrow[rr] & & V_D
\end{tikzcd}$$

in which the maps $\nearrow$ are surjective and the maps $\searrow$ are injective. The diagram can be thought of as a persistence module over an 8-element poset, the partial order being implied by the arrows. We will carry out quiver calculations by restricting to various totally-ordered subsets of this poset.

To begin with, surjectivity of the maps $\nearrow$ means that

$$\langle \circ_A — \bullet_a \rangle = 0 \quad \text{and} \quad \langle \circ_B — \bullet_b \rangle = 0,$$

and injectivity of the maps $\searrow$ 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

$$\begin{aligned}
\langle \circ_A ——— \bullet_b — \bullet_c ——— \circ_D \rangle &= \langle \circ_A — \bullet_B — \bullet_b — \bullet_c — \bullet_C — \circ_D \rangle \\
&\quad + \text{three other terms} \\
&= \langle \circ_A — \bullet_B — \bullet_b — \bullet_c — \bullet_C — \circ_D \rangle \\
&= \langle \circ_A — \bullet_B ——————— \bullet_C — \circ_D \rangle \\
&= \langle \circ_A — \bullet_B — \bullet_C — \circ_D \mid \mathbb{V} \rangle
\end{aligned}$$

and at the same time

$$\begin{aligned}
\langle \circ_A \text{———} \bullet_b \text{—} \bullet_c \text{———} \circ_D \rangle &= \langle \circ_A \text{—} \circ_a \text{—} \bullet_b \text{—} \bullet_c \text{—} \circ_d \text{—} \circ_D \rangle \\
&\quad + \text{three other terms} \\
&= \langle \circ_A \text{—} \circ_a \text{—} \bullet_b \text{—} \bullet_c \text{—} \circ_d \text{—} \circ_D \rangle \\
&= \langle \text{——} \circ_a \text{—} \bullet_b \text{—} \bullet_c \text{—} \circ_d \text{——} \rangle \\
&= \langle \circ_a \text{—} \bullet_b \text{—} \bullet_c \text{—} \circ_d \mid \mathbb{V}^\varepsilon \rangle
\end{aligned}$$

so we get the required equality. The six 'other terms' are all zero because they contain forbidden configurations. $\square$

Corollary 5.19 *Let* $\mathbb{V}$ *be q-tame. Then* $\mathrm{d_b}(\mathsf{dgm}(\mathbb{V}), \mathsf{dgm}(\mathbb{V}^\varepsilon)) \leq \varepsilon$.

Proof Indeed, an ε-matching is defined as follows:

$$(p, q) \in \mathsf{dgm}(\mathbb{V}^\varepsilon) \quad \leftrightarrow \quad (p - \varepsilon, q + \varepsilon) \in \mathsf{dgm}(\mathbb{V})$$

This is bijective except for the unmatched points of $\mathsf{dgm}(\mathbb{V})$, which lie on or below the line Δ_ε, and therefore have distance at most ε from the diagonal. $\square$

Proposition 5.20 *If* $\mathbb{V}$ *is q-tame then* $\mathbb{V}^\varepsilon$ *is locally finite.*

Proof Since $\dim((V^\varepsilon)_t) = \mathrm{rank}[V_{t-\varepsilon} \to V_{t+\varepsilon}] < \infty$, it follows from Theorem 2.8 (2) that $\mathbb{V}^\varepsilon$ is decomposable into interval modules. We claim that the collection of intervals is locally finite. Specifically, for any $t \in \mathbf{R}$, we estimate

$$\begin{aligned}
&\#\{\text{intervals which meet } [t - \tfrac{1}{2}\varepsilon, t + \tfrac{1}{2}\varepsilon]\} \\
&\quad = \#\{\text{points of } \mathsf{Dgm}(\mathbb{V}^\varepsilon) \text{ in the upper-left quadrant at } (t + \tfrac{1}{2}\varepsilon, t - \tfrac{1}{2}\varepsilon)\} \\
&\quad \leq \#\{\text{points of } \mathsf{Dgm}(\mathbb{V}) \text{ in the upper-left quadrant at } (t - \tfrac{1}{2}\varepsilon, t + \tfrac{1}{2}\varepsilon)\} \\
&\quad = \mathrm{rank}[V_{t-\frac{1}{2}\varepsilon} \to V_{t+\frac{1}{2}\varepsilon}]
\end{aligned}$$

which is finite. The '$\leq$' in the third line is a consequence of Proposition 5.18. $\square$

We are now ready to prove the converse stability theorem for q-tame persistence modules, using the triangle inequalities for $\mathrm{d_i}$, $\mathrm{d_b}$ and our results on ε-smoothing.

Proof (5.14″) Let $\mathbb{U}, \mathbb{V}$ be q-tame persistence modules. For any $\varepsilon > 0$, the ε-smoothings $\mathbb{U}^\varepsilon, \mathbb{V}^\varepsilon$ are decomposable, so the converse stability theorem applies to them. Then:

$$\begin{aligned}
\mathrm{d_i}(\mathbb{U}, \mathbb{V}) &\leq \mathrm{d_i}(\mathbb{U}^\varepsilon, \mathbb{V}^\varepsilon) + 2\varepsilon && \text{by Proposition 5.16} \\
&\leq \mathrm{d_b}(\mathsf{dgm}(\mathbb{U}^\varepsilon), \mathsf{dgm}(\mathbb{V}^\varepsilon)) + 2\varepsilon && \text{by Theorem 5.11} \\
&\leq \mathrm{d_b}(\mathsf{dgm}(\mathbb{U}), \mathsf{dgm}(\mathbb{V})) + 4\varepsilon && \text{by Corollary 5.19}
\end{aligned}$$

Since this is true for all $\varepsilon > 0$, we deduce that

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

The converse stability theorem for q-tame modules is proved. □

We finish this section with a characterisation of q-tame modules.

Theorem 5.21 *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 $\mathbb{V}^\varepsilon$, which are locally finite by Proposition 5.20. 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 $\varepsilon < (c - b)/2$. Then

$$\mathrm{r}_b^c = \operatorname{rank}[V_b \to V_c] = \operatorname{rank}[V_b \to W_{b+\varepsilon} \to W_{c-\varepsilon} \to V_c] \leq \dim(W_{b+\varepsilon})$$

which is finite. It follows that $\mathbb{V}$ is q-tame. □

Example 5.22 It is easy to see that there are q-tame modules which are not locally finite. For instance:

$$\bigoplus_{n=1}^{\infty} \mathbf{k}[0, \tfrac{1}{n}] \quad \text{and} \quad \prod_{n=1}^{\infty} \mathbf{k}[0, \tfrac{1}{n}]$$

The latter is the example of Crawley–Boevey [24] with no interval decomposition discussed in Remark 2.9. Incidentally, one can verify directly that the two modules are 0-interleaved; and also that their persistence measures, and hence their persistence diagrams, are equal away from the unique singular point $(0^-, 0^+)$.

5.6 The Stability Theorem

The inequality (5.14′) can be expressed in the following form:

Theorem 5.23 *Let $\mathbb{U}, \mathbb{V}$ be q-tame persistence modules which are δ^+-interleaved. Then there exists a δ-matching between the multisets* $\mathsf{dgm}(\mathbb{U})$, $\mathsf{dgm}(\mathbb{V})$.

It is easier to prove the following. (Notice the missing $^+$.)

Theorem 5.24 *Let $\mathbb{U}, \mathbb{V}$ be q-tame persistence modules which are δ-interleaved. Then there exists a δ-matching between the multisets* $\mathsf{dgm}(\mathbb{U})$, $\mathsf{dgm}(\mathbb{V})$.

Theorem 5.12 allows us to deduce Theorem 5.23 from Theorem 5.24: 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 5.24 depends on two main ingredients:

- The interpolation Lemma 4.6, which embeds $\mathbb{U}$, $\mathbb{V}$ within a 1-parameter family.
- The box inequalities 5.26, which relate 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 [19]. Our persistence diagrams may have infinite cardinality, so we will need an additional compactness argument to finish off the proof.

Definition 5.25 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, we use $-\infty - \delta = -\infty$ 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 5.26 (Box inequalities [19]) *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 region extension convention (Remark 3.22) we can state the lemma without the requirement that R^δ lies above the diagonal, since 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

$$\mathbb{U}_{a,b,c,d} \;:\; U_a \to U_b \to U_c \to U_d$$

and

$$\mathbb{V}_{A,B,C,D} \;:\; V_A \to V_B \to V_C \to V_D$$

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 $\mu_{\mathbb{U}}(R) \leq \mu_{\mathbb{V}}(R^{\delta})$. Then $\mu_{\mathbb{V}}(R) \leq \mu_{\mathbb{U}}(R^{\delta})$ follows by symmetry. □

Recall the measures at infinity defined in Sect. 3.6. By considering the appropriate limits, we immediately have:

Proposition 5.27 (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 $R \in \mathrm{Rect}(\overline{\mathbf{R}}^2)$. Then

$$\begin{aligned}
&\mu([a,b],-\infty) \leq \nu([A,B],-\infty), \qquad && \mu(-\infty,[c,d]) \leq \nu(-\infty,[C,D]),\\
&\mu([a,b],+\infty) \leq \nu([A,B],+\infty), && \mu(+\infty,[c,d]) \leq \nu(+\infty,[C,D]),
\end{aligned}$$

for all $a < b$ and $c < d$; and

$$\begin{aligned}
&\mu(-\infty,-\infty) \leq \nu(-\infty,-\infty), \qquad && \mu(+\infty,-\infty) \leq \nu(+\infty,-\infty),\\
&\mu(-\infty,+\infty) \leq \nu(-\infty,+\infty), && \mu(+\infty,+\infty) \leq \nu(+\infty,+\infty).
\end{aligned}$$

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 $\mu_{\mathbb{U}}(-\infty,+\infty) = \mu_{\mathbb{V}}(-\infty,+\infty)$.

5.7 The Measure Stability Theorem

We now embed Theorem 5.24 as a special case of a stability theorem for the diagrams of abstract r-measures. The more general statement is no more difficult[1] 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

$$\mathrm{ex}^\infty(\alpha, \mathcal{D}) = \mathrm{d}^\infty(\alpha, \overline{\mathbf{R}}^2 - \mathcal{D}) = \min\left(\mathrm{d}^\infty(\alpha, x) \mid x \in \overline{\mathbf{R}}^2 - \mathcal{D}\right). \qquad (5.2)$$

For instance, for the extended half-plane we have $\mathrm{ex}^\infty(\alpha, \overline{\mathcal{H}}) = \mathrm{d}^\infty(\alpha, \Delta)$.

Let A, B be multisets in $\mathcal{D}$. A δ**-matching** between A, B is a partial matching $\mathsf{M} \subset \mathsf{A} \times \mathsf{B}$ such that

$$\begin{aligned} \mathrm{d}^\infty(\alpha, \beta) &\le \delta \quad \text{if } \alpha, \beta \text{ are matched,} \\ \mathrm{ex}^\infty(\alpha, \mathcal{D}) &\le \delta \quad \text{if } \alpha \in \mathsf{A} \text{ is unmatched,} \\ \mathrm{ex}^\infty(\beta, \mathcal{D}) &\le \delta \quad \text{if } \beta \in \mathsf{B} \text{ is unmatched.} \end{aligned}$$

If $\mathcal{D}$ is not clear from the context, we refer to M as a 'δ-matching between $(\mathsf{A}, \mathcal{D})$ and $(\mathsf{B}, \mathcal{D})$'.

With the same proof as Proposition 5.9, we have:

Proposition 5.28 (triangle inequality) *If* $\mathsf{A}, \mathsf{B}, \mathsf{C}$ *are multisets in* $\mathcal{D}$ *and there exist a* δ_1*-matching between* $(\mathsf{A}, \mathcal{D}), (\mathsf{B}, \mathcal{D})$ *and a* δ_2*-matching between* $(\mathsf{B}, \mathcal{D}), (\mathsf{C}, \mathcal{D})$, *then there exists a* $(\delta_1 + \delta_2)$*-matching between* $(\mathsf{A}, \mathcal{D}), (\mathsf{C}, \mathcal{D})$. □

Now for the main theorem.

Theorem 5.29 (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*

$$\mu_x(R) \le \mu_y(R^{|y-x|})$$

holds for all rectangles R *whose* $|y - x|$*-thickening* $R^{|y-x|}$ *belongs to* $\mathrm{Rect}(\mathcal{D})$. *Then there exists a* δ*-matching between the undecorated diagrams* $(\mathsf{dgm}(\mu_0), \mathcal{D})$ *and* $(\mathsf{dgm}(\mu_\delta), \mathcal{D})$.

In view of the interpolation Lemma 4.6, this implies Theorem 5.24 (take $\mu_x = \mu(\mathbb{U}_x)$ and $\mathcal{D} = \overline{\mathcal{H}}^\circ$) and therefore the stability theorem (5.14′) for q-tame modules.

Example 5.30 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

[1] In fact it's a little easier to prove, because the compactness argument for diagrams with infinitely many points can be made more cleanly in this generality.

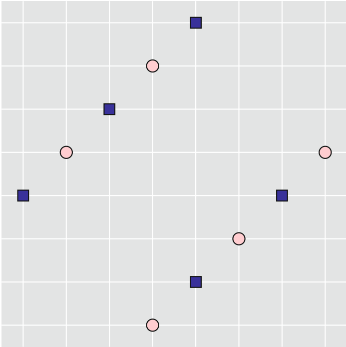

Fig. 5.2 The box inequalities do not control the bottleneck distance. The two diagrams (5 *dark blue squares*; 5 *light pink circles*) have box distance 1 and bottleneck distance 3. Generalising this example, one can exhibit a pair of diagrams with $4k + 1$ points each, whose box distance is 1 and whose bottleneck distance is $2k + 1$. By 'box distance 1' we mean that every rectangle R covers at most as many pink points as its 1-thickening R^1 covers blue points, and vice versa

$\mathsf{dgm}(\mu)$, $\mathsf{dgm}(\nu)$ which are δ-matched. This is not true, and in fact there is no universal bound on the bottleneck distance between the two diagrams. See Fig. 5.2.

Our goal for the rest of this section is to prove Theorem 5.29. Parts 1 and 2 closely follow the method of Cohen-Steiner, Edelsbrunner and Harer [19]. Afterwards, in Sect. 5.8, 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* $(\mathsf{dgm}(\mu_x), \mathcal{D})$, *and* $(\mathsf{dgm}(\mu_y), \mathcal{D})$ *is at most* $|y - x|$.

Write $\mathsf{A} = \mathsf{dgm}(\mu_x)$, $\mathsf{B} = \mathsf{dgm}(\mu_y)$, and $\eta = |y - x|$. The assertion is understood to mean:

- If $\alpha \in \mathsf{A}$ and $\mathrm{ex}^\infty(\alpha, \mathcal{D}) > \eta$, then there exists $\beta \in \mathsf{B}$ with $\mathrm{d}^\infty(\alpha, \beta) \leq \eta$.
- If $\beta \in \mathsf{B}$ and $\mathrm{ex}^\infty(\beta, \mathcal{D}) > \eta$, then there exists $\alpha \in \mathsf{A}$ with $\mathrm{d}^\infty(\alpha, \beta) \leq \eta$.

Proof By symmetry, it is enough to prove the first statement. Given α, let $\varepsilon > 0$ be small enough that $\eta + \varepsilon < \mathrm{ex}^\infty(\alpha, \mathcal{D})$. Then the box inequality gives

$$1 \leq \mu_x(\alpha^\varepsilon) \leq \mu_y(\alpha^{\eta+\varepsilon})$$

so there is at least one point of B in the square $\alpha^{\eta+\varepsilon}$. This is true for all sufficiently small $\varepsilon > 0$, and moreover B is locally finite. Therefore there is at least one point of B in α^η. $\square$

Henceforth, we will write $\mathsf{A}_x = \mathsf{dgm}(\mu_x)$ for all x.

Part 2. *The theorem is true if* A_x *has finite cardinality for all* x.

Proof (i) The triangle inequality for matchings includes the implication

$$\left.\begin{array}{l}\mathsf{A}_0, \mathsf{A}_x \text{ are } x\text{-matched} \\ \mathsf{A}_x, \mathsf{A}_y \text{ are } (y-x)\text{-matched}\end{array}\right\} \Rightarrow \mathsf{A}_0, \mathsf{A}_y \text{ are } y\text{-matched}$$

whenever $0 < x < y$.

(ii) We claim that for every $x \in [0, \delta]$ there exists $\rho(x) > 0$ such that $\mathsf{A}_x, \mathsf{A}_y$ 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 A_x, with respective multiplicities $n_1, \dots, n_k$. Let $\rho(x)$ be chosen to satisfy the following finite set of constraints:

$$0 < \rho(x) \leq \begin{cases} \frac{1}{2}\mathrm{ex}^\infty(\alpha_i, \mathcal{D}) & \text{all } i \\ \frac{1}{2}\mathrm{d}^\infty(\alpha_i, \alpha_j) & \text{all } i,j \text{ distinct}\end{cases}$$

We must show that if $|y - x| < \rho(x)$ then $\mathsf{A}_x, \mathsf{A}_y$ are $|y - x|$-matched. Write $\eta = |y - x|$ and let

$$(\mathbf{R}^2 - \mathcal{D})^\eta = \left\{\alpha \in \mathcal{D} \mid \mathrm{ex}^\infty(\alpha, \mathcal{D}) \leq \eta\right\}.$$

It follows from Part 1 that 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 A_y in each square α_i^η. Let $\varepsilon > 0$ be small enough that $2\eta + \varepsilon < 2\rho(x)$. Then the box inequality gives

$$n_i = \mu_x(\alpha_i^\varepsilon) \leq \mu_y(\alpha_i^{\eta+\varepsilon}) \leq \mu_x(\alpha_i^{2\eta+\varepsilon}) = n_i.$$

Thus $\mu_y(\alpha_i^{\eta+\varepsilon}) = n_i$ for all small $\varepsilon > 0$. We conclude that the square α_i^η contains precisely n_i points of 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 A_y in the square α_i^η, for each i, to define an η-matching between $(\mathsf{A}_x, \mathcal{D})$, $(\mathsf{A}_y, \mathcal{D})$. All points of A_x are matched, and the only unmatched points of A_y lie in $\mathbf{R}^2 - \mathcal{D}$ and do not need to be matched.

Items (i) and (ii) formally imply that $\mathsf{A}_0, \mathsf{A}_\delta$ are δ-matched, using a standard Heine–Borel argument. Indeed, let

$$m = \sup(x \in [0, \delta] \mid \mathsf{A}_0 \text{ and } \mathsf{A}_x \text{ are } x\text{-matched}).$$

First, m is positive; specifically $m \geq \rho(0)$. Applying (i) to $0 < m' < m$, where $\mathsf{A}_0, \mathsf{A}_{m'}$ are m'-matched and $m - m' < \rho(m)$, we deduce that $\mathsf{A}_0, \mathsf{A}_m$ are m-matched. Suppose

$m < \delta$. Applying (i) to $0 < m < m''$, where $m'' - m < \rho(m)$, we deduce that $\mathsf{A}_0, \mathsf{A}_{m''}$ are m''-matched. This contradicts the definition of m. Therefore $m = \delta$, and $\mathsf{A}_0, \mathsf{A}_\delta$ are δ-matched. □

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 A_x is locally finite, it follows that $\mathsf{A}_x \cap \mathcal{D}_n$ 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 M_n between $(\mathsf{A}_0 \cap \mathcal{D}_n, \mathcal{D}_n)$ and $(\mathsf{A}_\delta \cap \mathcal{D}_n, \mathcal{D}_n)$.

We now take a limit M of the partial matchings M_n, using the construction in the proof of Theorem 5.12. This works because $\mathsf{A}_0, \mathsf{A}_\delta$ are locally finite and therefore countable. Let χ, χ_n denote the indicator functions of M, M_n. Recall Lemma 5.13: for any finite subset $\mathsf{F} \subset \mathsf{A}_0 \times \mathsf{A}_\delta$, there are infinitely many $n \in \mathbf{N}$ for which

$$\chi(\alpha, \beta) = \chi_n(\alpha, \beta)$$

for all $(\alpha, \beta) \in \mathsf{F}$.

We must show that M is a δ-matching between $(\mathsf{A}_0, \mathcal{D})$ and $(\mathsf{A}_\delta, \mathcal{D})$. It is immediate that each matched pair is separated by at most δ, since this is true for every M_n. The argument that each α is matched with at most one β, and vice versa, is the same as in Lemma 5.13.

Finally, suppose $\alpha \in \mathsf{A}_0$ with $\mathrm{ex}^\infty(\alpha, \mathcal{D}) > \delta$. 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 $\mathrm{ex}^\infty(\alpha, \mathcal{D}_n) > \delta$ and hence α is matched in M_n for sufficiently large n. Now α has only finitely many δ-neighbours $\beta_1, \dots, \beta_k$ in the locally finite set A_δ, so by Lemma 5.13 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 $\beta \in \mathsf{A}_\delta$ with $\mathrm{ex}^\infty(\beta, \Delta) > \delta$ is matched in M to some $\alpha \in \mathsf{A}_0$. It follows that M is the required δ-matching. □

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 5.27.

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. □

This completes the proof of the stability theorem for finite measures on an open region $\mathcal{D}$, and hence the stability theorem for q-tame persistence modules, and hence the isometry theorem for q-tame persistence modules.

5.8 The Measure Stability Theorem (Continued)

The stability theorem generalises to measures that are not necessarily finite. By the region extension convention (Remark 3.22), 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

$$\mathcal{F}^{-\delta} = \left\{\alpha \in \mathcal{F} \mid \mathrm{ex}^{\infty}(\alpha, \mathcal{F}) > \delta\right\} = \left\{\alpha \in \mathcal{F} \mid \alpha^{\delta} \subset \mathcal{F}\right\}.$$

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 5.31 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 regions. 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** $(\mathsf{A}, \mathcal{F}), (\mathsf{B}, \mathcal{G})$ 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 $(\alpha, \beta) \in \mathsf{M}$ then $\mathrm{d}^{\infty}(\alpha, \beta) \leq \delta$,
- every $\alpha \in \mathsf{A} \cap \mathcal{G}^{-\delta}$ is matched with some $\beta \in \mathsf{B}$,
- every $\beta \in \mathsf{B} \cap \mathcal{F}^{-\delta}$ is matched with some $\alpha \in \mathsf{A}$.

The first of these is a compatibility condition between the regions: 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 region $\mathcal{G}$, and vice versa.

Proposition 5.32 (triangle inequality) *If* $\mathsf{A}, \mathsf{B}, \mathsf{C}$ *are multisets in* $\mathcal{F}, \mathcal{G}, \mathcal{H}$ *respectively, and there exist a* δ_1*-matching between* $(\mathsf{A}, \mathcal{F}), (\mathsf{B}, \mathcal{G})$ *and a* δ_2*-matching between* $(\mathsf{B}, \mathcal{G}), (\mathsf{C}, \mathcal{H})$*, then there exists a* $(\delta_1 + \delta_2)$*-matching between* $(\mathsf{A}, \mathcal{F})$, $(\mathsf{C}, \mathcal{H})$.

Proof As usual, compose the two partial matchings to get a partial matching M between A, C. Writing $\delta = \delta_1 + \delta_2$, we must check that this is a δ-matching between $(\mathsf{A}, \mathcal{F}), (\mathsf{C}, \mathcal{H})$. For the 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 $\alpha \in \mathsf{A}$ lies in $\mathcal{H}^{-\delta}$ then by the inclusion above it lies in $\mathcal{G}^{-\delta_1}$. Therefore α is matched with $\beta \in \mathsf{B}$. Moreover β must then lie in $\mathcal{H}^{-(\delta-\delta_1)} = \mathcal{H}^{-\delta_2}$ and so is matched with $\gamma \in \mathsf{C}$. The fourth condition follows by symmetry. □

Remark 5.33 There is no triangle inequality if the compatibility condition between the regions is dropped in the definition of δ-matching.

Here is the main theorem of this section and the last new theorem of the monograph. Again we use the abbreviation $\mathcal{F}_x = \mathcal{F}^\circ(\mu_x)$ for the finite interiors.

Theorem 5.34 (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 $R \in \mathrm{Rect}(\overline{\mathbf{R}}^2)$. *Then there exists a* δ*-matching between the undecorated diagrams* $(\mathsf{dgm}(\mu_0), \mathcal{F}_0)$ *and* $(\mathsf{dgm}(\mu_\delta), \mathcal{F}_\delta)$.

Remark 5.35 This version of measure stability allows us to dispense with the final assertion in Theorem 4.7 when deducing stability (5.14′) for q-tame modules: we no longer need to assume that the interpolating modules are themselves q-tame. The reader may wish to consider why this works. It results from careful management of the boundary.

An easy first step is to verify the condition on the regions $\mathcal{F}_0, \mathcal{F}_\delta$:

Proposition 5.36 *Under the hypotheses of Theorem 5.34, 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 $\varepsilon > 0$ such that $\alpha^{|y-x|+\varepsilon} \subset \mathcal{F}_y$. The box inequality gives

$$\mu_x(\alpha^\varepsilon) \leq \mu_y(\alpha^{|y-x|+\varepsilon})$$

and the right-hand side is finite by Proposition 3.18. Thus $\alpha \in \mathcal{F}^\circ(\mu_x) = \mathcal{F}_x$. □

Proof (*Theorem* 5.34) 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 $\mathsf{A}_x = \mathsf{dgm}(\mu_x)$.

Initial remark. Recall that the proof is carried out separately for each of the nine strata. The four corners at infinity are handled easily (each corner belongs to both $\mathcal{F}_0$ and $\mathcal{F}_\delta$, 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* $(\mathsf{A}_x, \mathcal{F}_x)$, *and* $(\mathsf{A}_y, \mathcal{F}_y)$ *is at most* $\eta = |y - x|$.

The assertion is understood to mean:

- If $\alpha \in \mathsf{A}_x$ and $\mathrm{ex}^\infty(\alpha, \mathcal{F}_y) > \eta$, then there exists $\beta \in \mathsf{A}_y$ with $\mathrm{d}^\infty(\alpha, \beta) \le \eta$.
- If $\beta \in \mathsf{A}_y$ and $\mathrm{ex}^\infty(\beta, \mathcal{F}_x) > \eta$, then there exists $\alpha \in \mathsf{A}_x$ with $\mathrm{d}^\infty(\alpha, \beta) \le \eta$.

Proof By symmetry it is enough to prove the first statement. Let $\alpha \in \mathsf{A}_x$. For all $\varepsilon > 0$ with $\eta + \varepsilon < \mathrm{ex}^\infty(\alpha, \mathcal{F}_y)$, we have $1 \le \mu_x(\alpha^\varepsilon) \le \mu_y(\alpha^{\eta+\varepsilon})$ so there is at least one point of A_y in α^η. $\square$

Part 2. *The theorem is true if* A_x *has finite cardinality for all* x.

Proof Item (i) is given by the triangle inequality (Proposition 5.32).

Item (ii) uses the same strategy as before. Let (α_i) be a finite enumeration of the distinct points of A_x, with respective multiplicities (n_i). Then $\rho(x)$ is chosen to satisfy

$$0 < \rho(x) \le \begin{cases} \frac{1}{2}\mathrm{ex}^\infty(\alpha_i, \mathcal{F}_x) & \text{all } i \\ \frac{1}{2}\mathrm{d}^\infty(\alpha_i, \alpha_j) & \text{all } i,\, j \text{ distinct.} \end{cases}$$

If $\eta = |y - x| < \rho(x)$, then Part 1 implies that A_y is contained in the disjoint union

$$(\mathbf{R}^2 - \mathcal{F}_x)^\eta \cup \alpha_1^\eta \cup \cdots \cup \alpha_k^\eta.$$

The box inequality is then used to count precisely n_i points of A_y in the square α_i^η. This defines a partial matching where all points of A_x are matched and all points of $\mathsf{A}_y \cap \mathcal{F}_x^{-\eta}$ 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.*

Proof 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 $\varepsilon > 0$ and $r > \delta$, let

$$\hat{\mathcal{F}}_x = \mathcal{F}_x^{-\varepsilon} \cap \mathcal{Q}^r$$

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 $\mathsf{dgm}(\hat{\mu}_x) = \mathsf{dgm}(\mu_x) \cap \hat{\mathcal{F}}_x$.

Lemma 5.37 *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,

$$\hat{\mathcal{F}}_y^{-|y-x|} = (\mathcal{F}_y^{-\varepsilon} \cap \mathcal{Q}^r)^{-|y-x|} = \mathcal{F}_y^{-(\varepsilon+|y-x|)} \cap \mathcal{Q}^{r-|y-x|} \subseteq \mathcal{F}_x^{-\varepsilon} \cap \mathcal{Q}^r = \hat{\mathcal{F}}_x$$

as required. □

We resume the proof of Part 3. Since $\hat{\mathcal{F}}_x$ has compact closure in $\mathcal{F}_x$, and A_x is locally finite, it follows that $\hat{\mathsf{A}}_x = \mathsf{dgm}(\hat{\mu}_x) = \mathsf{A}_x \cap \hat{\mathcal{F}}_x$ has finite cardinality. We can therefore apply Part 2 to the family $(\hat{\mu}_x)$ to get a δ-matching between $(\hat{\mathsf{A}}_0, \hat{\mathcal{F}}_0)$ and $(\hat{\mathsf{A}}_\delta, \hat{\mathcal{F}}_\delta)$. This can be interpreted as a partial δ-matching between $\mathsf{A}_0, \mathsf{A}_\delta$ where:

- $\alpha \in \mathsf{A}_0$ is matched whenever $\alpha \in (\mathcal{F}_\delta^{-\varepsilon} \cap \mathcal{Q}^r)^{-\delta} = \mathcal{F}_\delta^{-(\delta+\varepsilon)} \cap \mathcal{Q}^{r-\delta}$
- $\beta \in \mathsf{A}_\delta$ is matched whenever $\beta \in (\mathcal{F}_0^{-\varepsilon} \cap \mathcal{Q}^r)^{-\delta} = \mathcal{F}_0^{-(\delta+\varepsilon)} \cap \mathcal{Q}^{r-\delta}$

Repeat this argument for a sequence (ε_n, r_n) where $\varepsilon_n \to 0$ and $r_n \to +\infty$. This gives a sequence of δ-matchings M_n, and we can form a limit M as before.

If $\alpha \in \mathsf{A}_0 \cap \mathcal{F}_\delta^{-\delta}$ then eventually $\alpha \in \mathcal{F}_\delta^{-(\delta+\varepsilon_n)} \cap \mathcal{Q}^{r_n-\delta}$ and so α is matched by M_n for all sufficiently large n. The same is true for $\beta \in \mathsf{A}_\delta \cap \mathcal{F}_0^{-\delta}$. With this information, we can complete the usual proof that M is a δ-matching between $(\mathsf{A}_0, \mathcal{F}_0)$ and $(\mathsf{A}_\delta, \mathcal{F}_\delta)$. This completes the proof of Part 3, and hence of Theorem 5.34. □

Here is a sample consequence.

Example 5.38 (*Stability of the Webb module*) Let $\mathbb{V}$ be a persistence module which is δ-interleaved with the module $\mathbb{W}$ of Example 3.31. By interpolation (Lemma 4.6) and the box inequalities (Lemma 5.26), we can apply the measure stability theorem (Theorem 5.34): there exists a δ-matching between the undecorated diagrams $(\mathsf{dgm}(\mu_\mathbb{V}), \mathcal{F}^\circ(\mu_\mathbb{V}))$ and $(\mathsf{dgm}(\mu_\mathbb{W}), \mathcal{F}^\circ(\mu_\mathbb{W}))$. This amounts to the following.

- **In the finite part $\mathcal{H}$ of the half-plane**: Any singular points of $\mu_{\mathbb{V}}$ are confined to the diagonal strip $\Delta_{[0,\delta]}$. Each point of $\mathsf{dgm}(\mu_{\mathbb{V}})$ outside this strip is matched with some point $(-n, 0) \in \mathsf{dgm}(\mu_{\mathbb{W}})$. Conversely, the only unmatched points of $\mathsf{dgm}(\mu_{\mathbb{W}})$ must lie within distance δ of the diagonal or a singular point of $\mu_{\mathbb{V}}$. In particular, if $\delta < \frac{1}{4}$ then all points of $\mathsf{dgm}(\mu_{\mathbb{W}})$ are matched.
- **On the line** $(-\infty, \mathbf{R})$: All points and singular points 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 singular points.