Decomposition and Characterisation of Monodromy in 2-parameter Persistence Modules

Octave Mortain
supervised by Sara Scaramuccia

February-June 2024

Contents

About the internship and this report: My internship at the University of Roma Tor Vergata under the supervision of Sara Scaramuccia was split in three main phases. I first learned through reading and attending a class given by Paolo Salvatore the notions of Topological Data Analysis which where unknown to me before. I then spent the rest of the internship exploring two different ideas around this relatively new domain of mathematics. The first one did not, despite over a month of work, give any significant result. The second one did however yield two interesting and new results. The goal of this report is to present these two theorems and their proofs.

1 Introduction

Persistent homolgy is one of the most important tools in Topological Data Analysis (TDA). It was introduced by Edelsbrunner in [ELZ02] as a way to describe the topological properties of an object – most often a point cloud in the case of TDA – across different scales. Edelsbrunner also gave an efficient algorithm to compute persistent homology, which happens to also be the fastest way to compute homology. This makes persistent homology one of the best tool at our disposal to compute topological information.

It has been used in many different fields to describe topological properties of complex objects. Some example of the possible uses of persistent homology appear in robotics to study trajectories ([BGK15], [PHA16]), in neuroscience to study brain arteries ([Ben+16]), in bio-engineering to study the shape and properties of proteins ([XW14]) and in various ways in machine learning ([HMR21]).

However, persistent homology is a tool that produces a rarely used algebraic structure, persistence modules. This means that in order to properly study persistent homology, one needs to understand the structure of persistence modules. When persistence is done over a single parameter, this is relatively easy thanks to a powerful structure theorem (Theorem 3.2). But some applications such as [XW15] make use of multiple parameters at the same time.

The study of multi-parameter persistence is vastly more complicated because of the absence of any satisfying structure theorem. Indeed, while it is possible to decompose multi-parameter persistence modules (see [BC20]), the resulting invariant does not satisfy the conditions presented by Carlsson in [Car09]. In other words, we have very little information about what we are decomposing into.

To remedy this, some papers like [DX21] studied a subclass of persistence modules which can be decomposed into simpler elements: interval-decomposable modules.

Others have tried to reduce to the single-parameter structure theorem from the multi-parameter case by considering the restrictions to lines in the parameter space. In order to extract significant information however, the structures along each of these lines cannot be studied independently. Thankfully, when moving continuously among these lines, the structure of the restricted persistence module also changes continuously, and as a result, it is possible to lift it. When doing this, a phenomenon of monodromy appears, first described by Cerri, Ethier and Frosini in [CEF13]. Successive papers [CEF16] and [CEF19] develop a new notion of coherent matching which relies on the same lifting.

During my internship, I studied the phenomenon of monodromy in multi-parameter persistence using an algebraic point of view, while Cerri, Ethier and Frosini use a mainly topological point of view. This allows me to adapt the example of monodromy given in [CEF13] into a simpler example which I show in section 5.1.

Using this algebraic point of view I proved two main results. Firstly, in Section 5.2 I show that, if a module decomposes into direct summands, then monodromy happens separately within each summand. This is a particularly useful result as it shows that some modules like the interval-decomposable modules always have trivial monodromy, and it shows that the study of monodromy can be restricted to indecomposable modules.

Afterwards, I obtained a characterisation of monodromy which I present in section 5.3. The two results in this section give a complete description of the image in the monodromy group of a loop around a single line.

2 Introduction to persistence modules

2.1 Persistence

Definition 2.1: Persistence objects

Let P be an ordered set and $\mathcal{C}$ a category. A *P-persistence object* of $\mathcal{C}$ is a collection $\{M_x\}_{x\in P}$ of objects of $\mathcal{C}$ and a collection of arrows $\{f_{x,y} : M_x \to M_y\}_{x\leqslant y}$, such that $f_{y,z} \circ f_{x,y} = f_{x,z}$ for all $x, y, z \in P$.

Alternatively, a P-persistence object can be seen a a functor $\mathcal{P} \to \mathcal{C}$ where $\mathcal{P}$ is the category associated to the ordered set P.

In this paper we will mostly use $\mathbb{R}_+$-persistence, and $\mathbb{R}_+^p$-persistence objects. The order on $\mathbb{R}_+^p$ will always be the product order. We will refer to $\mathbb{R}_p$-persistence and $\mathbb{R}_+^p$-persistence as single-parameter and multi-parameter persistence respectively.

Definition 2.2: Persistent morphisms

A morphism φ between two P-persistence objects M and N is a collection of morphisms $(\varphi_x : M_x \to N_x)_{x\in P}$ such that the following diagram commute:

$$\begin{array}{ccc} M_x & \longrightarrow & M_y \\ \downarrow{\scriptstyle \varphi_x} & & \downarrow{\scriptstyle \varphi_y} \\ N_x & \longrightarrow & N_y \end{array}$$

In the case where $P = \mathbb{R}_+^p$, a morphism of degree $d \in \mathbb{R}_+$ is a collection of morphisms $(\varphi_x : M_x \to N_{x+(d,\ldots,d)})_{x\in P}$ such that the following diagram commute:

$$\begin{array}{ccccc} M_x & & \longrightarrow & & M_y \\ & \searrow^{\varphi_x} & & & & \searrow^{\varphi_y} \\ & & N_{x+(d,\ldots,d)} & & \longrightarrow & & N_{y+(d,\ldots,d)} \end{array}$$

A morphism φ said to be injective (resp. surjective) if and only if φ_x is injective (resp. surjective) for every x.

For $x \in \mathbb{R}_+^p$ and $d \in \mathbb{R}_+$, we will simplify the notations by writing $x + d$ instead of $x + (d, \ldots, d)$.

Persistence objects also inherit notions such as direct sum and quotients if those exists in the original category. These are defined naturally by taking the direct sum (or quotient) for each $x \in P$, and checking that the morphisms can still be defined.

2.2 Persistence modules

In this paper, we will only use persistence objects of categories of modules.

Persistence modules are obviously persistence objects of a category of modules, however they can also be seen as graded modules over an appropriate graded ring. This equivalence between persistence modules and graded modules is well documented in papers such as [CK18] and was used extensively during the internship. However, the results shown in this paper do not rely in any way on the theory of graded modules. As such, and in order to avoid having to introduce numerous notions from graded algebra, we will only use the notation and concept of generators and relations, which are redefined here in the terms of persistence modules.

Note that this is not standard notation when it comes to persistence modules, and that it will not be used in the proofs, only in the examples and illustrations.

Definition 2.3

Given a partially ordered set P, a ring A and a generator g of degree $\deg(g) \in P$, the P-persistence A-module generated by g is the persistence module $\langle g \rangle$ defined by

$$\langle g \rangle_x = \begin{cases} A \text{ if } x \geqslant \deg(g) \\ 0 \text{ otherwise,} \end{cases}$$

with the arrows being Id_A whenever possible. The generator g will then coincide with the element $1_A \in \langle g \rangle_{\deg(g)}$

We define the persistence module generated by multiple generators $g_1, \ldots, g_n$ in the natural way: $\langle g_1, \ldots, g_n \rangle = \bigoplus_{i=1}^{n} \langle g_i \rangle$.

Finally, given an element $R \in \langle g_1, \ldots, g_n \rangle_d$ which we will call relation in degree d, we can define the persistence submodule $\langle R \rangle$ of $\langle g_1, \ldots, g_n \rangle$ by $\langle R \rangle_x = f_{d,x}(A \cdot R)$ if $d \leqslant x$ and 0 otherwise. Given multiple relations like this, we can define persistence modules by generators and relations like so:

$$\langle g_1, \ldots, g_n | R_1, \ldots, R_m \rangle = \langle g_1, \ldots, g_n \rangle / \sum_{i=1}^{m} \langle R_i \rangle.$$

If the degree $\deg(R)$ of a relation is clear, we will often write g instead of $f_{\deg(g),\deg(R)}(g)$. In order to underline the fact that R is a relation that we quotient by, we will write $R = 0$ instead of R , as in Figure 1.

For example, in order to describe the module $\langle g | R \rangle$ with $\deg(g) = 1$ and $R = f_{1,2}(g)$, we will simply do the drawing Figure 1:

More generally, given $b, d \in \mathbb{R}_+$ the degrees of birth and death respectively, we define the single-parameter *interval module* $I(b,d) = \langle g | R \rangle$ with $\deg(g) = b$ and $R = f_{b,d}(g)$.

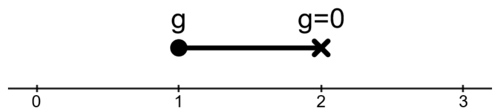

Figure 1: The interval module $I(1,2)$.

3 Persistence diagrams

3.1 Structure theorem

$\mathbb{N}$-persistence modules have a very powerful structure theorem presented in [Car09]. This work however focuses on persistence modules with real parameters, hence I consider a subset of $\mathbb{R}_+$-persistence modules that behave similarly to $\mathbb{N}$-persistence ones:

Definition 3.1: Discrete $\mathbb{R}_+$-persistence module

An $\mathbb{R}_+$-persistence module M is *discrete* if there exist reals $0 = c_0 < c_1 < \cdots < c_n < c_{n+1} = +\infty$ such that M is constant on the intervals $[c_i; c_{i+1}[$, or equivalently, if $c_i \leqslant x < y < c_{i+1}$ then $M_x = M_y$ and $f_{x,y} = \mathrm{Id}$. The set $\{c_0, \ldots, c_{n+1}\}$ will be called a grid of M.

Theorem 3.2: Structure Theorem

Let M be a discrete $\mathbb{R}_+$-persistence module over a field F. Suppose that $\forall x, \dim(M_x) < \infty$. ($M$ is then said to be *tame*). Then there exists an integer N and couples of reals (b_n, d_n) with $0 \leqslant b_n \leqslant d_n \leqslant +\infty$ such that:

$$M \cong \bigoplus_{n=1}^{N} I(b_n, d_n).$$

Additionally, the pairs $(b_n, d_n)_{n=1}^{N}$ are unique up to permutation. Furthermore, we have that the birth and death points b_n and d_n belong to a grid $\{c_i\}$ of the module.

This means that all the information about the persistence module can be represented in a multiset of points in $\mathbb{R} \times (\mathbb{R} \cup \{+\infty\})$, where each point (b, d) corresponds to a generator that appears at b and dies at d.

Definition 3.3: Diagram

Given a single-parameter persistence module M such that $M \cong \bigoplus_{n=1}^{N} U(b_n, d_n)$ we define the *persistence diagram* of M as the multiset of points in $\mathbb{R}_+ \times (\mathbb{R}_+ \cup \{+\infty\})$:

$$\mathrm{Dgm}(M) = \{(b_n, d_n), n \in [\![1; N]\!]\} \cup \Delta,$$

where Δ is the diagonal of $\mathbb{R}_+^2$ and has infinite multiplicity.

The structure theorem shows that discrete tame $\mathbb{R}_+$-persistence modules are completely described by their diagram.

Example: Consider the module M described in Figure 2a, which has two generators g and h with respective degrees 1 and 1.5 that become equal in degree 3 and die in degree 4.

This module is not at first glance the direct sum of interval modules. The structure theorem states that it is possible to write it as such a sum. Indeed, by using $h' = h - g$ instead of h (Figure 2b) we can see that $M \cong I(1, 4) \oplus I(1.5, 3)$.

As a result, we obtain the diagram of M in Figure 2c. Each non-diagonal point in the diagram corresponds to an interval in the previous decomposition.

Note that the closer a point in the diagram is to the diagonal, the shorter the corresponding segment is, and as such points closer to the diagonal are less significant. A point in the diagonal would be invisible in the diagram (because Δ is already included with infinite multiplicity in the diagram) because it would represent an interval of the form $I(x, x) = 0$.

Figure 2

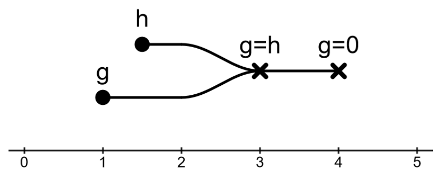

(a) M

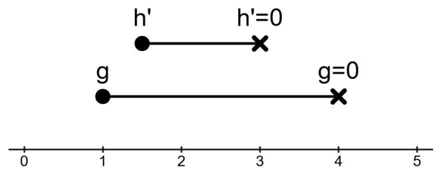

(b) M using g and $h' = h - g$

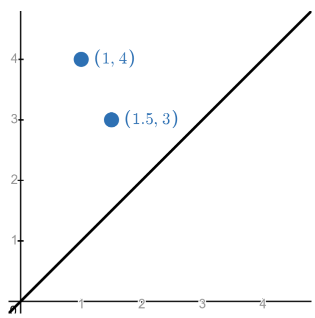

(c) $\mathrm{Dgm}(M)$

3.2 Distances

This section shows the two main ways of defining how close two persistence modules are. We will focus on $\mathbb{R}_+$-persistence and $\mathbb{R}_+^2$-persistence modules, which will sometimes be referred to as 1-parameter and 2-parameter persistence modules, respectively.

The first definition of a distance (Definition 3.5) applies to both types of modules but is hard to compute in practice because of it's very theoretical nature.

The second one (Definition 3.7) makes use of persistence diagrams as a way of representing 1-parameter persistence modules. This makes it much easier to use and compute, but restricts it to the single-parameter case.

3.2.1 Interleaving distance

Definition 3.4: Interleaving

Let M and N be two $\mathbb{R}_+^p$-persistence modules and $\delta \geqslant 0$. M and N are *δ-interleaved* if there exist two morphisms $\phi : M \to N$ and $\psi : N \to M$ of degree δ such that the following diagrams commute:

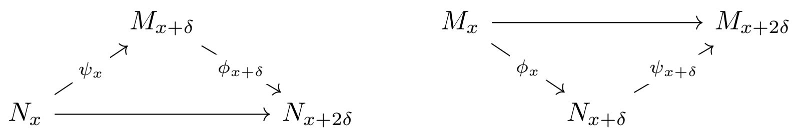

Note that being 0-interleaved means being isomorphic. Additionally, we can see that if $\delta < \delta'$, being δ-interleaved is stronger than being δ'-interleaved (by using $\phi'_x : M_x \to N_{x+\delta'} = \phi_{x+\delta'-\delta} \circ f_{x,x+\delta'-\delta}$). As a result, we can define a reasonable way to measure distances between persistence modules.

Definition 3.5: Interleaving distance

Let M and N be two persistence modules. The *interleaving distance* between them is defined as:

$$d_I(M, N) = \inf\{\delta, M \text{ and } N \text{ are } \delta\text{-interleaved}\}.$$

Note that the interleaving distance can be infinite: we will for instance show in Proposition 3.6 that $d_I(0, I(1, +\infty)) = +\infty$.

3.2.2 Bottleneck distance

When computing the interleaving distance between two of the interval modules that appear in the structure theorem, we get the following.

Proposition 3.6

Let (a, b) and (c, d) be points in $\mathbb{R}_+ \times (\mathbb{R}_+ \cup \{+\infty\})$ with $a < b$ and $c < d$, and let $M = I(a, b)$ and $N = I(c, d)$. Then

$$d_I(M, N) = \min\left\{d_\infty((a, b), (c, d)), \max\{d_\infty((a, b), \Delta), d_\infty((c, d), \Delta)\}\right\},$$

where Δ is the diagonal in $\mathbb{R}_+^2$ and d_∞ is the infinity distance in $\mathbb{R}_+ \times (\mathbb{R}_+ \cup \{+\infty\})$.

Computing the interleaving distance can be difficult, even for simple modules. However, we see that in the case of the single-parameter indecomposable modules, one can interpret the distance between $I(a, b)$ and $I(c, d)$ geometrically in the plane as the ∞-distance between the points, if they are closer to one another than to the diagonal, or their ∞-distance to the diagonal otherwise.

Given that we can identify modules to sets of points in the plane using their persistence diagrams, it is natural to try to express the interleaving distance geometrically using the diagrams. We can do that using the bottleneck distance associated to d_∞:

Definition 3.7: Bottleneck distance

Let A and B be multisets of points in $\mathbb{R}^2$. The *bottleneck distance* is defined as:

$$d_B(A, B) = \inf_{\varphi : A \to B} \sup_{a \in A} d_\infty(a, \varphi(a)),$$

where φ ranges over all bijections from A to B.

The bottleneck distance between two diagrams is always defined because both diagrams are composed of a finite number of points above Δ, and Δ itself with infinite multiplicity, hence it is always possible to find a bijection.

Additionally, it is always possible to match points within Δ to themselves, and as a result all that matters is how we match the points above the diagonal. The non-diagonal points of a diagram can each be matched either with a non-diagonal point of the other diagram, or with the diagonal. An example is shown in Figure 3: the distance between the orange and blue diagrams is the length of the longest red segment (for the infinite norm).

In a more formal way, we have that the distance $d_B(\mathrm{Dgm}(M), \mathrm{Dgm}(N))$ is the infimum over all partial bijections $\varphi : A \subset \mathrm{Dgm}(M) \to B \subset \mathrm{Dgm}(N)$ of the maximum of the distances between matched points: $d_\infty(a, \varphi(a))$, $a \in A$ and the distances from non matched points to the diagonal: $d_\infty(x, \Delta)$, $x \in \mathrm{Dgm}(M) \setminus A$ or $x \in \mathrm{Dgm}(N) \setminus B$.

The following theorem presented by Lesnick in [Les15] shows that computing the bottleneck distance between the diagrams yields the interleaving distance.

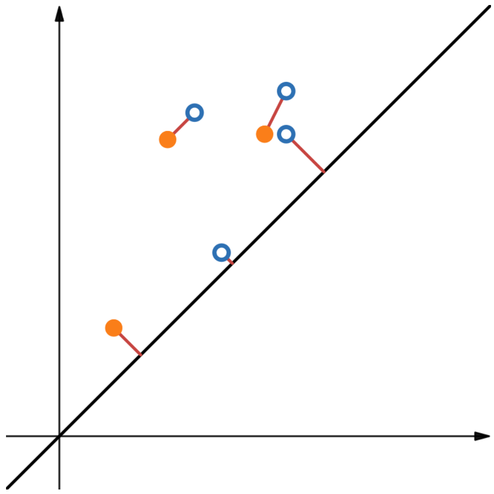

Figure 3: A matching between two diagrams

Theorem 3.8: Isometry Theorem

Let M and N be tame single-parameter persistence modules. We have:

$$d_I(M, N) = d_B(\mathrm{Dgm}(M), \mathrm{Dgm}(N)).$$

4 Lifting diagrams

From now on, we will be focusing only on 2-parameter persistence modules. Our objective in this section is to use the structure of 1-parameter persistence modules to study the 2-parameter ones.

4.1 Line parametrisation

Consider a 2-parameter persistence module M. In order to apply the structure theorem for single-parameter persistence modules, we can restrict ourselves to a line in the parameter plane.

We will parameterise the lines as so: for $0 < a < 1$ and $b \in \mathbb{R}$, the line $\ell = (a, b)$ is the line with slope $\frac{1-a}{a}$ that passes through the point $(b, -b)$.

Definition 4.1: Restriction to a line

Given a $\mathbb{R}^2_+$-persistence module M, $0 < a < 1$ and $b \in \mathbb{R}$, $M^{(a,b)}$ is the $\mathbb{R}_+$-persistence module defined by:

$$\tilde{M}^{(a,b)}_t = M_{(at+b,(1-a)t-b)}$$

and reparametrised by:

$$M^{(a,b)}_t = \tilde{M}^{(a,b)}_{\frac{t}{\max(a,1-a)}}$$

with the arrows naturally defined from those in M, and with the convention that if x or y are negative, $M_{(x,y)} = M_{(\max(x,0),\max(y,0))}$.

Note that the factor $\frac{1}{\max(a,1-a)}$ is here to make sure that the distance between the values of t correspond to the restriction of the infinite norm to the line.

In order to apply the structure theorem to the modules obtained by restricting along a line, we need some hypothesis on M. We will consider the following:

Definition 4.2: Discrete $\mathbb{R}_+^2$-persistence module

A $\mathbb{R}_+^2$-persistence module M is *discrete* if for $k = 1, 2$, there exists a finite number of reals $0 = c_0^k < c_1^k < \cdots < c_{n_k}^k < c_{n_k+1}^k = +\infty$ such that M is constant in the rectangles $[c_i^1; c_{i+1}^1[\times[c_j^2; c_{j+1}^2[$.

The pair of sets $(\{c_0^1, \ldots, c_{n_1+1}^1\}, \{c_0^2, \ldots, c_{n_2+1}^2\})$ will be called a grid of M.

Proposition 4.3

If M is discrete, then so is $M^{(a,b)}$. Additionally, one can give an expression for a grid of $M^{(a,b)}$ given a grid of M.

Proof. If $(C^1, C^2) = (\{c_0^1, \ldots, c_{n_1+1}^1\}, \{c_0^2, \ldots, c_{n_2+1}^2\})$ is a grid of M, then

$$\{0, +\infty\} \cup \left\{ \frac{c_k^1 - b}{a} \max(a, 1-a), k \in [\![0; n_1]\!] \right\} \cup \left\{ \frac{c_k^2 + b}{1-a} \max(a, 1-a), k \in [\![0; n_2]\!] \right\}$$

is a grid of $M^{(a,b)}$. ∎

Note that the expression of the resulting grid can be written as the union of the images of C^1 and C^2 by two different affine transformations.

$$\{0\} \cup \frac{C^1 - b}{a} \max(a, 1-a) \cup \frac{C^2 + b}{1-a} \max(a, 1-a).$$

The space of lines $]0; 1[\times\mathbb{R}$ behaves in a dual way to the parameter space:

Proposition 4.4

The set of lines that go through the point (x, y) is the open segment between $(0, x)$ and $(1, -y)$. The set of lines that go through the rectangle $[x_1; x_2[\times[y_1, y_2[$ is the open quadrilateral with vertices $(0, x_1)$, $(0, x_2)$,$(1, -y_1)$,$(1, -y_2)$.

4.2 Continuity

We now want to study the evolution of $\mathrm{Dgm}(M^{(a,b)})$ as a and b vary continuously. In the following, we prove that its evolution is continuous with respect to the topology induced by the bottleneck distance. Thanks to the isometry theorem, it is equivalent the evolution of $M^{(a,b)}$ being continuous with respect to the interleaving distance.

Lemma 4.5

Let $0 < a < 1$ and $b, b' \in \mathbb{R}$. $M^{(a,b)}$ and $M^{(a,b')}$ are $\left(|b - b'| \max(\frac{a}{1-a}, \frac{1-a}{a})\right)$-interleaved.

Proof. Suppose for example $b < b'$ and $a \geqslant 0.5$. Computation shows that we can define

$$\phi_t : M_t^{(a,b)} \to M_{t+\frac{a}{1-a}|b-b'|}^{(a,b')},$$

$$\psi_t : M_t^{(a,b)} \to M_{t+|b-b'|}^{(a,b')}$$

that commute appropriately. Given $\frac{a}{1-a} \geqslant 1$, we have that $M^{(a,b)}$ and $M^{(a,b')}$ are $\frac{a}{1-a}|b - b'|$-interleaved. The proof is very similar if $b > b'$ or if $a \leqslant 0.5$. ∎

A similar result can be obtained with a similar proof for the parameter a:

Lemma 4.6

Let $0.5 \leqslant a, a' < 1$ and $b \in \mathbb{R}$. Suppose that M_x is a constant module for $\|x\| \geqslant C$. Then $M^{(a,b)}$ and $M^{(a',b)}$ are $\max(C\frac{a-a'}{(1-a)-a'}, C\frac{a'-a}{(1-a')-a})$-interleaved.

The result also holds if $0 < a, a' \leqslant 0.5$. From this we can immediately deduce what we wanted:

Proposition 4.7

The map

$$\ell \in]0;1[\times\mathbb{R} \longmapsto \mathrm{Dgm}(M^{(\ell)})$$

is continuous with respect to the bottleneck distance.

4.3 Lifting

Given a continuous path of lines $\gamma : [0;1] \to]0;1[\times\mathbb{R}$, we can track the diagram $\mathrm{Dgm}(M^{(\gamma(t)})$ which will evolve continuously. From this, we would like to be able to track individually each point moving in the diagram.

This is clearly impossible once two points cross or once a point collides with the diagonal (i.e. appears or disappears). As such we need to restrict ourselves to certain lines by defining the following.

Definition 4.8: Admissibility

Let $n \in \mathbb{N}$. We say that a persistence diagram D is *n-admissible* if $D \setminus \Delta$ contains exactly n elements and has no elements with multiplicity above 1. We will write $\mathcal{D}_n$ the set of n-admissible diagrams.

Given a 2-parameter persitence module M, we say that the line $\ell \in]0;1[\times\mathbb{R}$ is *n-admissible* for M if the diagram $\mathrm{Dgm}(M^{(\ell)})$ is. We will write $\mathcal{A}_n(M)$ (or simply $\mathcal{A}_n$ if the context is clear) the set of n-admissible lines for M.

Once we take a continuous path with values in $\mathcal{A}_n$, we obtain a continuous path with values in $\mathcal{D}_n$ which we want to lift into the n different paths of the n points is the diagram.

Proposition 4.9

Let $\Delta^+ = \{(b,d), 0 \leqslant b < d \leqslant +\infty\}$ be the open half-plane above the diagonal. Let $n \in \mathbb{N}$. The application

$$p : \left\{ \begin{array}{l} \{(z_1,\dots,z_n) \in (\Delta^+)^n, z_1,\dots,z_n \text{ are all distinct}\} \to \mathcal{D}_n \\ (z_1,\dots,z_n) \longmapsto \{z_1,\dots,z_n\} \cup \Delta \end{array} \right.$$

that takes a tuple of points (ordered) and turns them into an unordered diagram, is a covering.

Proof. The preimage of any admissible diagram $D \in \mathcal{D}_n$ is of cardinally $n!$. Additionaly, p is a local homeomorphism. Hence p is a covering. ■

As a result, given a continuous path $\gamma : [0;1] \to \mathcal{A}_n$ and an ordering $(z_1,\dots,z_n)$ to the elements of $\mathrm{Dgm}(M^{(\gamma(0))}) \setminus \Delta$, one can lift the path $\mathrm{Dgm}(M^{(\gamma)})$ into a path in $(\Delta^+)^n$, which we will write $\tilde{\gamma} = (\tilde{\gamma}_1,\dots,\tilde{\gamma}_n)$ with $\tilde{\gamma}_i(0) = z_i$. This $\tilde{\gamma}$ is continuous and verifies $p \circ \tilde{\gamma} = \mathrm{Dgm}(M^{(\gamma)})$.

5 Monodromy

Given the previous results, one could hope to obtain a way to globally "number" each point in each diagram that respects the continuity, using the previously described lifting. The example is Section 5.1 shows that this is not possible, as a phenomenon of monodromy appears: when lifting the different points in a diagram along a loop, some may get swapped.

In order to study this phenomenon, I present the two main results of my internship, the Monodromy Decomposition Theorem (Theorem 5.2) and characterisation, in Section 5.2 and Section 5.3 respectively.

5.1 Definition and example

Consider the module M in Figure 4a, defined by two generators g and h of degrees $\deg(g) = (1,0)$ and $\deg(h) = (0,1)$ and by the relations $g = 0$ in degree $(2,0)$ and $g = h$ in degree $(1,2)$. The figure shows the dimension of the module in each degree (0, 1, and 2 in white, light blue and dark blue respectively).

The lines ℓ such that $\mathrm{Dgm}(M^{(\ell)})$ contains exactly two non-diagonal points are the ones that cross the the bigrades in the rectangle $[1;2[\times[0;2[$. In the space of lines, they form the quadrilateral in Figure 4b. One of these lines (the dotted orange line in Figure 4a) is however not admissible because the two points in it's diagram are equal: they are both born in degree 1 and dead in degree 2. As a result this line is not 2-admissible for M. This shows that $\mathcal{A}_2(M)$ is an open quadrilateral minus a point.

Figure 4: A 2-parameter module and it's set of 2-admissible lines

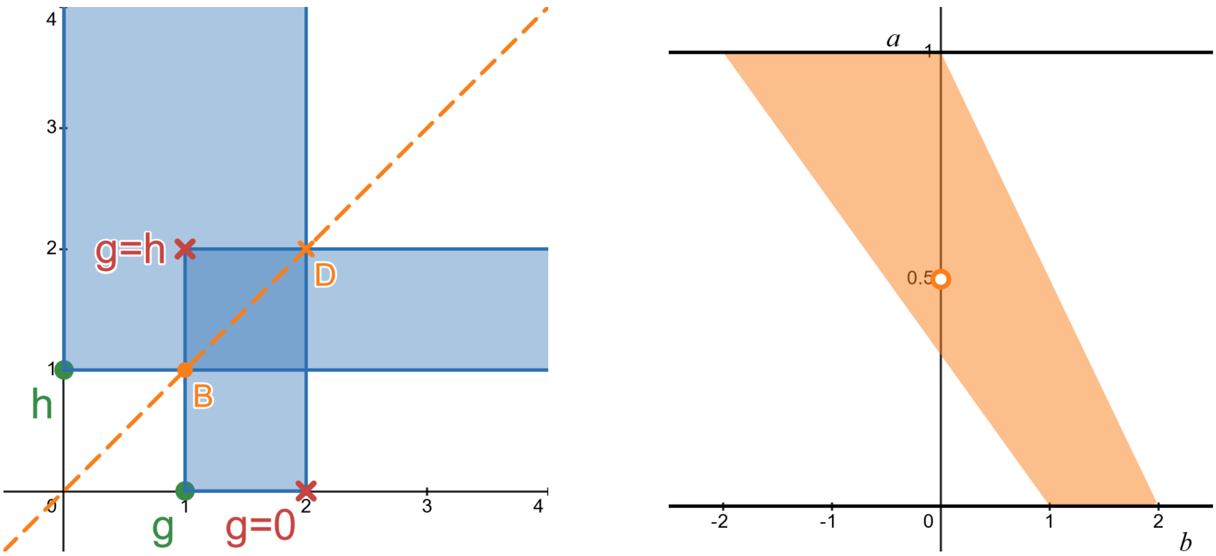

(a) M (b) $\mathcal{A}_2(M)$

This means that $\mathcal{A}_2(M)$ has a non-trivial fundamental group. As a result, if we consider a path $\gamma : [0;1] \to \mathcal{A}_2$ with $\gamma(0) = \gamma(1)$, we cannot be sure that $\tilde{\gamma}(0) = \tilde{\gamma}(1)$. We do know that $p(\tilde{\gamma}(0)) = p(\tilde{\gamma}(1))$ which means that $\tilde{\gamma}(0)$ and $\tilde{\gamma}(1)$ have the same elements, possibly permutated. In other words, the fundamental group $\Pi_1(\mathcal{A}_2) = \mathbb{Z}$ acts on the set of points in the diagram at a given line.

A generator of $\Pi_1(\mathcal{A}_2)$ is obtained by going around the single non admissible line in the middle. We can see in figure Figure 5 four different lines (orange) in this loop. Along each line, we obtain the decomposition in interval modules (black).

Figure 5: Four lines in a loop generating $\Pi_1(\mathcal{A}_2)$

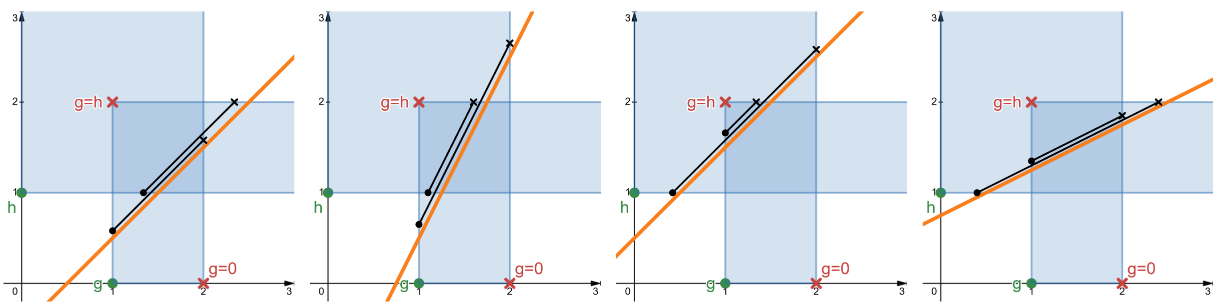

$\mathcal{A}$ $\mathcal{B}$ $\mathcal{C}$ $\mathcal{D}$

We can then lift this path as previously described in order to track the two points individually. In Figure 6, we can see how the diagram changes when going from one of the four lines above to the next.

Figure 6: The evolution of the diagram

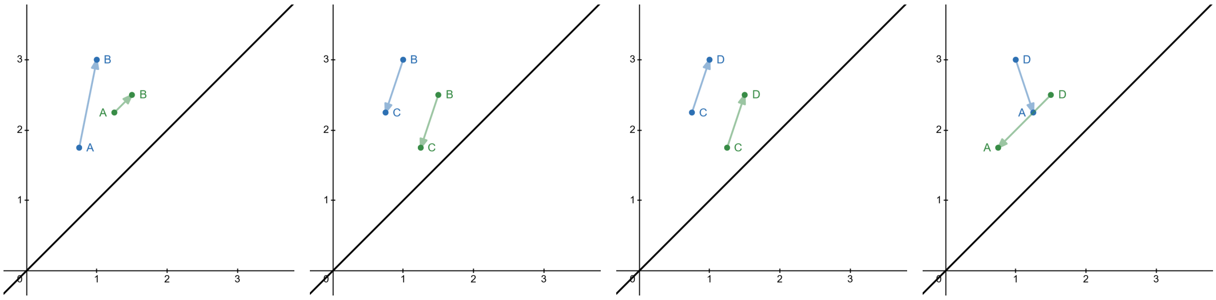

$\mathcal{A} \to \mathcal{B}$ $\qquad$ $\mathcal{B} \to \mathcal{C}$ $\qquad$ $\mathcal{C} \to \mathcal{D}$ $\qquad$ $\mathcal{D} \to \mathcal{A}$

We can see that $\tilde{\gamma}$ is not a loop: it swaps the two points.

Definition 5.1: Monodromy group

The *monodromy group* $G(M, \ell)$ of a discrete 2-parameter persistence module M in the n-admissible line ℓ is the image of the morphism from $\Pi_1(\mathcal{A}_n, \ell)$ to the group of permutations $\mathfrak{S}_n$ induced by the action of the fundamental group on $\mathrm{Dgm}(M^{(\ell)}) \setminus \Delta$.

This is well defined because $\mathcal{A}_n$ is locally connected. Additionaly, note that the monodromy group is the same for all lines in the same connected component of $\mathcal{A}_n$. In the previous example, we had $G(M, \ell) = \mathfrak{S}_2 \cong \mathbb{Z}_2$ for all $\ell \in \mathcal{A}_2$.

5.2 Monodromy Decomposition Theorem

Given a decomposition of a module as a direct sum, it is easy to see that the diagram along each line will be the union of the diagrams of the summands along the same line. The following result shows that the monodromy does not swap points from one to the other.

Theorem 5.2

Let M be a discrete 2-parameter persistence module with $M = M_1 \oplus M_2$. Consider $\ell \in \mathcal{A}_n(M)$ for a certain $n \in \mathbb{N}$. Then

$$\mathrm{Dgm}(M^{(\ell)}) = \mathrm{Dgm}(M_1^{(\ell)}) \cup \mathrm{Dgm}(M_2^{(\ell)})$$

and

$$G(M, \ell) \text{ is a subgroup of } G(M_1, \ell) \times G(M_2, \ell).$$

This means the monodromy group of M only permutates points from M_1 between them and points from M_2 between them. Before proving the theorem, we need the following lemma.

Lemma 5.3

Let M be a 2-parameter persistence module. The function

$$\begin{cases}]0;1[\times\mathbb{R} \to \mathbb{N} \cup \{+\infty\} \\ \ell \longmapsto |\mathrm{Dgm}(M^{(\ell)}) \setminus \Delta| \end{cases}$$

is lower semi-continuous in all points where it is finite.

Proof. By Proposition 4.7, all we need is to check that the map $D \mapsto |D \setminus \Delta|$ is lower semi-continuous whenever it is finite.

Consider a diagram D with a finite number of non-diagonal points: $D = \{z_1, \ldots, z_n\} \cup \Delta$, and let $\varepsilon = \min_i d_\infty(z_i, \Delta)$. For any diagram D' with $d_B(D, D') < \varepsilon$, we can consider a matching between D and D' with $\delta < \epsilon$. For each i, z_i cannot be matched to the diagonal because $\delta < d_\infty(z_i, \Delta)$. Hence D' must contain at least n distinct non-diagonal points: D is a local minimum, and so the map is lower semi-continuous in D.

■

We now prove Theorem 5.2.

Proof. We clearly have, for any $\ell \in]0;1[\times\mathbb{R}$, that $M^{(\ell)} = M_1^{(\ell)} \oplus M_2^{(\ell)}$. This gives us that $\mathrm{Dgm}(M^{(\ell)}) = \mathrm{Dgm}(M_1^{(\ell)}) \cup \mathrm{Dgm}(M_2^{(\ell)})$. This implies that if a line is n-admissible for M, then it is m-admissible for M_1 and $n-m$-admissible for M_2 for some $m \leqslant n$.

Let's fix $\ell \in \mathcal{A}_n(M)$ and $\gamma : [0;1] \to \mathcal{A}_n(M)$ a continuous loop based in ℓ. We would like to show that the m previously mentioned remains the same along γ.

If we define $m(t) = |\mathrm{Dgm}(M_1^{(\gamma(t))}) \setminus \Delta|$, then m is lower semi-continuous by the previous Lemma. Similarly, $n - m(t) = |\mathrm{Dgm}(M_2^{(\gamma(t))}) \setminus \Delta|$ is lower semi-continuous, which implies that $m(t)$ is upper semi-continuous, hence continuous, hence constant.

We now know that there exists a single m such that γ is a continuous loop in $\mathcal{A}_m(M_1)$ and $\mathcal{A}_{n-m}(M_2)$. Consider an ordering of $\mathrm{Dgm}(M^{(\ell)}) \setminus \Delta$ as $(z_1, \ldots, z_m, z_{m+1}, \ldots, z_n)$, with the first m elements being in the diagram of M_1 and the $n-m$ others if that of M_2. We can lift the path $\gamma : [0;1] \to \mathcal{A}_m(M_1)$ into $(\tilde{\gamma}_1, \ldots, \tilde{\gamma}_m)$ a path in $(\Delta^+)^m$. Similarly, we can lift the path $\gamma : [0;1] \to \mathcal{A}_{n-m}(M_2)$ into $(\tilde{\gamma}_{m+1}, \ldots, \tilde{\gamma}_n)$ a path in $(\Delta^+)^{n-m}$. Both those paths verify $\tilde{\gamma}_i(0) = z_i$. Consider then $\tilde{\gamma} = (\tilde{\gamma}_1, \ldots, \tilde{\gamma}_m, \tilde{\gamma}_{m+1}, \ldots, \tilde{\gamma}_n)$. It is a continuous path in $(\Delta^+)^n$ that verifies $\tilde{\gamma}_i(0) = z_i$ and $p \circ \tilde{\gamma} = \mathrm{Dgm}(M_1^{(\gamma)}) \cup \mathrm{Dgm}(M_2^{(\gamma)}) = \mathrm{Dgm}(M^{(\gamma)})$. Hence it is the unique lifting of γ.

$\tilde{\gamma}(1)$ is a permutation of $\tilde{\gamma}(0)$ which only swaps the first m elements between them. By construction it is the juxtaposition of a permutation in $G(M_1, \ell)$ and another in $G(M_2, \ell)$. ■

The proof showed that we have the following commutative diagram:

$$\begin{array}{ccc} \Pi_1(\mathcal{A}_n(M), \ell) & \twoheadrightarrow & G(M, \ell) \\ \downarrow & & \downarrow \\ \Pi_1(\mathcal{A}_m(M_1), \ell) \times \Pi_1(\mathcal{A}_{n-m}(M_2), \ell) & \twoheadrightarrow & G(M_1, \ell) \times G(M_2, \ell) \end{array}$$

Theorem 5.2 has an immediate consequence on interval decomposable modules, a subclass of 2-parameter persistence modules that can be decomposed in particularly simple elements (See [DX21]).

Corollary 5.4

If M is an interval decomposable module, then its monodromy group in any admissible line is trivial.

5.3 Description of monodromy around a single line

We give in this section a complete overview of the monodromy that can happen around a single non-admissible line. The first result shows that a loop around a single line can only transpose select pairs of points. Further results will show when this happens.

Note that this section does not give a complete characterisation of monodromy as it only studies loops around a single non-admissible lines. There can be "holes" around which to loop in the set of admissible lines that are larger than a single point. We will however not study them as they are both harder to work with and much less likely to appear.

5.3.1 Possible permutations

We want to study the monodromy induced by a loop around a single non-admissible line. Note that given such a line ℓ isolated in $]0;1[\times\mathbb{R} \setminus \mathcal{A}_n$, $|\mathrm{Dgm}(M^{(\ell)}) \setminus \Delta| \leqslant n$ since Lemma 5.3 states it is a local minimum.

It can actually be proven that $|\mathrm{Dgm}(M^{(\ell)}) \setminus \Delta| = n$, although we will not give the proof as it is very similar in nature to other proofs presented here.

As a result the only reason why ℓ is not admissible is because the diagram $\mathrm{Dgm}(M^{(\ell)})$ contains non-diagonal point with multiplicities higher than one. These points are obviously linked to the monodromy as without them, the line would be admissible and a loop around it would be homotopically trival. And so it is natural to think that if two points in the diagrams cross in ℓ, making the line non-admissible, then these two points are the most susceptible to being permutated by the monodromy group. The following result confirms this intuition.

Proposition 5.5

Let M be a discrete 2-parameter persistence module and ℓ be an isolated element of $]0;1[\times\mathbb{R}\setminus\mathcal{A}_n(M)$. Then $\mathrm{Dgm}(M^{(\ell)}) \setminus \Delta$ contains exactly n points counted with multiplicity and they can only have multiplicity 1 or 2:

$$\mathrm{Dgm}(M^{(\ell)}) = \{z_1, \dots, z_k, p_1, p_1, \dots, p_l, p_l\} \cup \Delta$$

with $k + 2l = n$. Consider $\ell' \in \mathcal{A}_n$ close enough to ℓ, so that we have

$$\mathrm{Dgm}(M^{(\ell')}) = \{z'_1, \dots, z'_k, p'_1, p''_1, \dots, p'_l, p''_l\} \cup \Delta$$

with z_i and z'_i close, p_j, p'_j and p''_j close. Then the image of the loop around ℓ in $G(M, \ell')$ is generated by the transpositions (p'_j, p''_j).

Proof. We first need to prove that the diagram does not contain any element of multiplicity 3 or higher.

Suppose that there exists a point $z \in \mathrm{Dgm}(M^{(\ell)})$ with multiplicity at least 3. We will show that there are other lines arbitrarily close to ℓ that are not admissible either, which would go against the hypothesis.

Write $z = (z_b, z_d)$ for the points of birth and death of this point. Consider the point P in the plane such that corresponding to z_d on the line ℓ.

Consider now a non constant path of lines γ with $\gamma(0) = \ell$ and such that for all t, the line $\gamma(t)$ goes through the point P.

See Figure 7a: Let $\varepsilon > 0$ smaller than half the minimum distance between two points in the grid. By construction of the grid presented in Proposition 4.3, we know that for t close enough to 0, the grid of $M^{(\gamma(t))}$ contains at most 2 points in $]z_b - \varepsilon; z_b + \varepsilon[$ and only one in $]z_d - \varepsilon; z_d + \varepsilon[$ (which corresponds to P).

See Figure 7b: For t small enough to have $d_B(\mathrm{Dgm}(M^{(\gamma(t))}), \mathrm{Dgm}(M^{(\ell)})) < \varepsilon$, we know that $\mathrm{Dgm}(M^{(\gamma(t))})$ contains at least 3 points in $]z_b - \varepsilon; z_b + \varepsilon[\times]z_d - \varepsilon; z_d + \varepsilon[$. But we know that the grid only has two points in the first interval and one in the second, hence they are only two possible positions in this square. This means that one of the points has multiplicity at least 2.

Hence for t small enough, $M^{(\gamma(t))}$ is not admissible, which is absurd as ℓ is isolated among non-admissible lines. This concludes the proof that the maximum multiplicity is 2, and so we can write

$$\mathrm{Dgm}(M^{(\ell)}) = \{z_1, \dots, z_k, p_1, p_1, \dots, p_l, p_l\} \cup \Delta$$

See Figure 7c: Keep the same ε as before (it is smaller than half the minimum distance between two points in the diagram). For ℓ' close enough to ℓ to have $d_B(\mathrm{Dgm}(M^{(\ell')}), \mathrm{Dgm}(M^{(\ell)})) < \varepsilon$, we can indeed write

$$\mathrm{Dgm}(M^{(\ell')}) = \{z'_1, \dots, z'_k, p'_1, p''_1, \dots, p'_l, p''_l\} \cup \Delta$$

with $d_\infty(z_i, z'_i) < \varepsilon$ for all i, $d_\infty(p_j, p'_j) < \varepsilon$ and $d_\infty(p_j, p''_j) < \varepsilon$ for all j. Note that the squares with centers z_i and p_j and with side 2ε are all disjoint from one another.

Hence if we consider γ a loop around ℓ that remains within the neighbourhood, we can lift γ into $(\tilde{\gamma}_1, \dots, \tilde{\gamma}_k, \tilde{\gamma}_{k+1}, \tilde{\gamma}_{k+2}, \dots, \tilde{\gamma}_{k+2l-1}, \tilde{\gamma}_{k+2l})$.

Since the squares previously discussed are open and disjoint and the components of $\tilde{\gamma}$ are continuous functions with values in their union, we know that for all $i \in [\![1; n]\!]$, $\tilde{\gamma}_i(0)$ and $\tilde{\gamma}_i(1)$ are in the same square. This means that the permutation $\tilde{\gamma}(0) \mapsto \tilde{\gamma}(1)$ leaves unchanged the z'_i (which are alone in their square) and potentially swaps the p'_j and p''_j. ∎

Figure 7: Illustration of the proof of Proposition 5.5

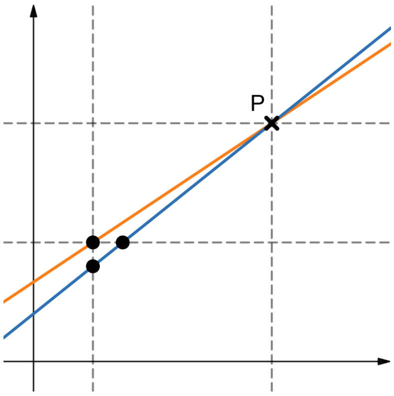

(a) Line ℓ (orange) and a line close to it that still goes through P (blue)

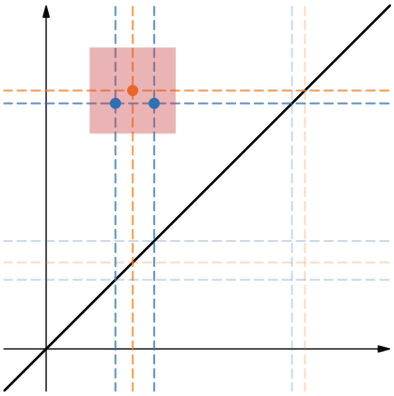

(b) The grids along both lines. There are at most two possible blue points around z

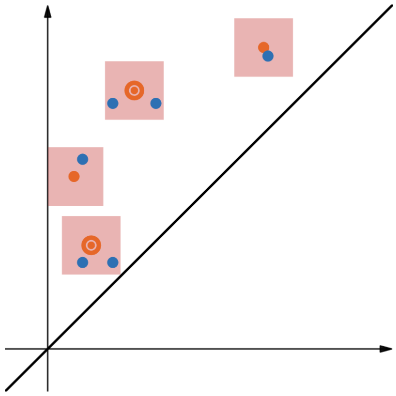

(c) In the neighborhood of $\mathrm{Dgm}(M^{(\ell)})$ (orange), points are limited to disjoint squares

5.3.2 Characterisation of transpositions

As a result, we want to study when pairs of points do swap and when they don't when circling around a single line.

For the rest of this section we will consider M a discrete 2-parameter persistence module, ℓ an isolated element of $]0;1[\times\mathbb{R}\setminus\mathcal{A}_n(M)$ and $p=(b,d)\in\mathrm{Dgm}(M^{(\ell)})$ of multiplicity 2. We will denote by B and D the points in the parameter plane that correspond to b and d respectively on line ℓ (see figure 4a).

We will also consider a convex neighbourhood N of ℓ such that for $\ell'\in N$, the grid of $M^{(\ell')}$ contains at most 2 points in $]b-\varepsilon,b+\varepsilon[$ and at most 2 points in $]d-\varepsilon,d+\varepsilon[$ with epsilon small as previously described. Because p is isolated, we know that B and D are intersections in the grid of M and so a line in N has one point near b in it's grid if it goes through B, and 2 otherwise.

Lemma 5.6

Consider a continuous path of lines γ within N, such that for all t, $\gamma(t)$ does not go through B. Then when lifting the two points close to p along γ, the order of their birth is kept.

Proof. For $t\in[0;1]$, since $\gamma(t)$ does not go through B, there are two elements $b_1(t)<b_2(t)$ in the grid of $M^{(\gamma(t))}$. By Proposition 4.3, b_1 and b_2 are continuous.

Let $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ be the liftings of the two points close to p, and let $\tilde{\gamma}_1^x$ and $\tilde{\gamma}_2^x$ be their first coordinate, i.e. their birth.

We have that for all t, $\tilde{\gamma}_1^x(t)\in\{b_1(t),b_2(t)\}$. This shows that the continuous function $\frac{\tilde{\gamma}_1^x-b_1}{b_2-b_1}$ takes values in $\{0,1\}$ and is hence constant. As a result, we have either $\tilde{\gamma}_1^x=b_1$ or $\tilde{\gamma}_1^x=b_2$ (for all t). The same can be proved for $\tilde{\gamma}_2^x$, and so we have:

$$\tilde{\gamma}_1^x(0)<\tilde{\gamma}_2^x(0)\text{ if and only if }\tilde{\gamma}_1^x(1)<\tilde{\gamma}_2^x(1)$$

Same is true with $=$ or $>$ instead of $<$. ■

An analogous proof shows that the result of Lemma 5.6 also applies to the order of death if the lines avoid D.

We now prove that the diagram of lines in N can only take two distinct forms, which we will call Long-Short and First-Second. This dichotomy will allow us to characterise when monodromy happens. It also serves as a proof that $d<+\infty$, and so D has finite coordinates, which will be necessary later in the characterisation.

Lemma 5.7: Dichotomy

Consider $\ell' \in N$ such that the line does not cross neither B nor D.

Then the grid of $M^{(\ell')}$ contains two points $b_1 < b_2$ close to b and two points $d_1 < d_2$ close to d. Furthermore, $\mathrm{Dgm}(M^{(\ell')})$ contains either (b_1, d_2) and (b_2, d_1) (*Long-Short* configuration) or (b_1, d_1) and (b_2, d_2) (*First-Second* configuration).

Figure 8: The two possibilities from the dichotomy and an impossible one

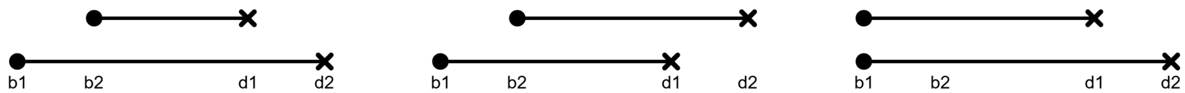

(a) Long-Short configuration (b) First-Second configuration (c) An impossible configuration

Proof. We know that $\mathrm{Dgm}(M^{(\ell')})$ contains two distinct points close to p. These points can take four values: (b_1, d_1), (b_1, d_2), (b_2, d_1) and (b_2, d_2). All we need to show the result is to show that these two points cannot share a coordinate.

By contradiction, suppose that the two points are (b_1, d_1) and (b_1, d_2) (the proof is similar if they share their second coordinate instead).

Denote by B' the point in the plane corresponding to b_1 on line ℓ' and consider the line ℓ'' that goes through B' and D. This line is still in the neighbourhood of ℓ and as such should be admissible.

The two points in $\mathrm{Dgm}(M^{(\ell')})$ have the same birth, hence by Lemma 5.6, so do the two points in $\mathrm{Dgm}(M^{(\ell'')})$. Additionally, since ℓ'' goes through D, we know that there is only one point near d in the grid of $M^{(\ell'')}$. This shows that the two points also have the same death: $\mathrm{Dgm}(M^{(\ell'')})$ is not admissible. ∎

The admissible lines in the neighbourhood of ℓ can be split into 8 categories based on whether the line goes above, below or through B and D. These categories are demarcated in the space of lines by the segments corresponding to B and D.

The ones we will be interested in are the ones which do not go through either points. We will call these four sectors $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$ and $\mathcal{D}$ as in Figure 9.

Figure 9: The four line positions

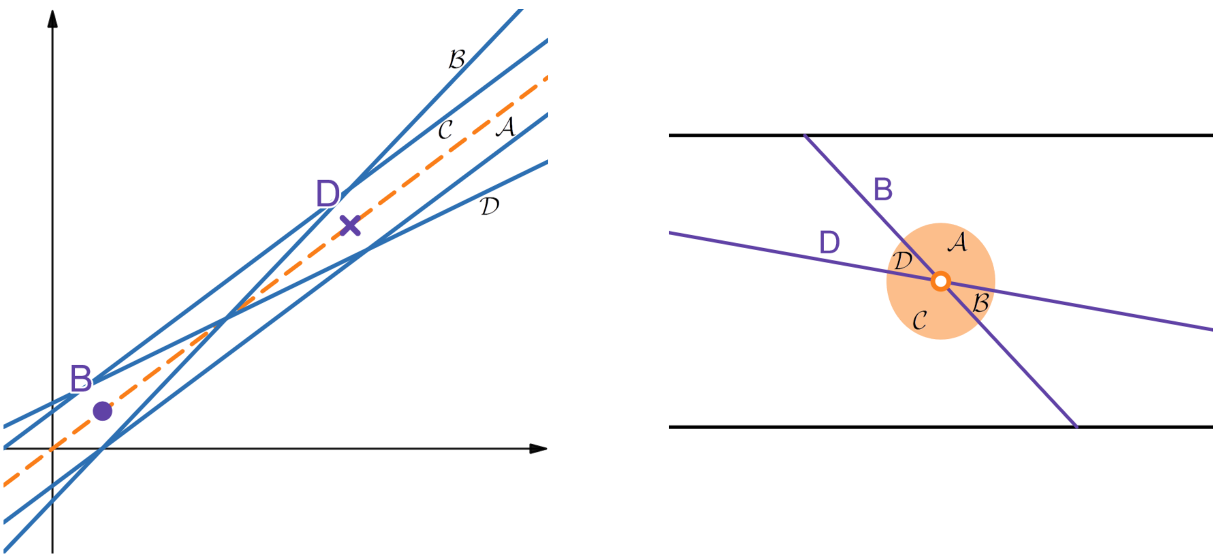

(a) One line from each sector (b) The different sectors in the space of lines

Theorem 5.8: Characterisation

Consider a line from each of the four sectors previously described. The loop around ℓ swaps the two points close to p if and only if an odd number of the four lines are in Long-Short configuration.

Proof. By Lemma 5.6, we know that all lines within a same sector have the same configuration.

We will say that there is a swap from one sector to another if when lifting, the order of birth of the two points are swapped. This happens if the Long/First configuration is sent to the Short/Second configuration.

Consider now a path from a line in $\mathcal{A}$ to one in $\mathcal{B}$ (not going through any other sectors). The lines in this path never cross B, although they do cross D. Hence the order of birth is kept. The same can be said when going from $\mathcal{C}$ to $\mathcal{D}$. Hence:

$$\mathcal{A} \to \mathcal{B} \text{ and } \mathcal{C} \to \mathcal{D} \text{ never swap.}$$

When going from a line in $\mathcal{B}$ to one in $\mathcal{C}$, we do cross B, but not D. Hence the order of death is kept. If the configuration are the same, then there still is no swap, if they are different however then keeping the death order swaps the birth order. Hence:

$$\mathcal{B} \to \mathcal{C} \text{ and } \mathcal{D} \to \mathcal{A} \text{ swap only if the configurations are different.}$$

Of course a loop around ℓ can be represented by doing $\mathcal{A} \to \mathcal{B} \to \mathcal{C} \to \mathcal{D} \to \mathcal{A}$. Such a loop swaps the two points if and only if it swaps their birth, that is if and only if there is an odd number of swaps. If we write $\mathcal{A} = 0$ for Long-Short and $\mathcal{A} = 1$ for First-Second, then the number of swaps written mod 2 is:

$$1_{\mathcal{B} \neq \mathcal{C}} + 1_{\mathcal{D} \neq \mathcal{A}} = (\mathcal{B} + \mathcal{C}) + (\mathcal{D} + \mathcal{A}) = \mathcal{A} + \mathcal{B} + \mathcal{C} + \mathcal{D} \mod 2.$$

This concludes the proof. ■

References

[BC20] Magnus Botnan and William Crawley-Boevey. "Decomposition of persistence modules". In: *Proceedings of the American Mathematical Society* 148.11 (2020), pp. 4581–4596.

[Ben+16] Paul Bendich et al. "Persistent Homology Analysis of Brain Artery Trees". In: *The annals of applied statistics* 10.1 (2016), pp. 198–218.

[BGK15] Subhrajit Bhattacharya, Robert Ghrist, and Vijay Kumar. "Persistent Homology for Path Planning in Uncertain Environments". In: *IEEE Transactions on Robotics* 31.3 (2015), pp. 578–590.

[Car09] Gunnar Carlsson. "Topology and data". In: *Bulletin of the American mathematical society* 46.2 (2009), pp. 255–308.

[CEF13] Andrea Cerri, Marc Ethier, and Patrizio Frosini. "A Study of Monodromy in the Computation of Multidimensional Persistence". In: *Discrete Geometry for Computer Imagery*. Springer Berlin Heidelberg, 2013, pp. 192–202.

[CEF16] Andrea Cerri, Marc Ethier, and Patrizio Frosini. "The Coherent Matching Distance in 2D Persistent Homology". In: *Computational Topology in Image Context*. Springer International Publishing, 2016, pp. 216–227.

[CEF19] Andrea Cerri, Marc Ethier, and Patrizio Frosini. "On the geometrical properties of the coherent matching distance in 2D persistent homology". In: *Journal of Applied and Computational Topology* 3 (2019), pp. 381–422.

[CEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. "Stability of Persistence Diagrams". In: *Discrete and Computational Geometry* 37 (2007), pp. 103–120.

[Cha+09] Frédéric Chazal et al. "Proximity of Persistence Modules and their Diagrams". In: *Proceedings of the twenty-fifth annual symposium on Computational geometry* (2009), pp. 237–246.

[CK18] René Corbet and Michael Kerber. "The representation theorem of persistence revisited and generalized". In: *Journal of Applied and Computational Topology* 2 (2018), pp. 1–31.

[DX21] Tamal K. Dey and Cheng Xin. "Rectangular Approximation and Stability of 2-parameter Persistence Modules". In: *arXiv preprint arXiv:2108.07429* (2021).

[ELZ02] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. "Topological Persistence and Simplification". In: *Discrete and Computational Geometry* 28.4 (2002), pp. 511–533.

[HMR21] Felix Hansel, Michael Moor, and Basitan Rieck. "A Survey of Topological Machine Learning Methods". In: *Front. Artif. Intell* 4 (2021), pp. 47–61.

[Les15] Michael Lesnick. "The Theory of the Interleaving Distance on Multidimensional Persistence Modules". In: *Foundations of Computational Mathematics* 15 (2015), pp. 613–650.

[Mun17] Elisabeth Munch. "A User's Guide to Topological Data Analysis". In: *Journal of Learning Analytics* 4.2 (2017), pp. 47–61.

[PHA16] Florian T. Pokorny, Majd Hawasly, and Subramanian Amamoorthy. "Topological trajectory classification with filtrations of simplicial complexes and persistent homology". In: *The International Journal of Robotics Research* 35.1-3 (2016), pp. 204–223.

[XW14] Kelin Xia and Guo-Wei Wei. "Persistent homology analysis of protein structure, flexibility, and folding". In: *Inernational Journal for Numerical Methods in Biomedical Engineering* 30.8 (2014), pp. 814–844.

[XW15] Kelin Xia and Guo-Wei Wei. "Multidimensional persistence in biomolecular data". In: *Journal of Computational Chemistry* 36.20 (2015), pp. 1502–1520.