

On Rectangle-Decomposable 2-Parameter Persistence Modules

Magnus Bakke Botnan

Vrije Universiteit Amsterdam, The Netherlands
m.b.botnan@vu.nl

Vadim Lebovici

École Normale Supérieure, Paris, France
vadim.lebovici@ens.fr

Steve Oudot

Inria Saclay, Palaiseau, France
steve.oudot@inria.fr

Abstract

This paper addresses two questions: (1) can we identify a sensible class of 2-parameter persistence modules on which the rank invariant is complete? (2) can we determine efficiently whether a given 2-parameter persistence module belongs to this class? We provide positive answers to both questions, and our class of interest is that of rectangle-decomposable modules. Our contributions include: (a) a proof that the rank invariant is complete on rectangle-decomposable modules, together with an inclusion-exclusion formula for counting the multiplicities of the summands; (b) algorithms to check whether a module induced in homology by a bifiltration is rectangle-decomposable, and to decompose it in the affirmative, with a better complexity than state-of-the-art decomposition methods for general 2-parameter persistence modules. Our algorithms are backed up by a new structure theorem, whereby a 2-parameter persistence module is rectangle-decomposable if, and only if, its restrictions to squares are. This local condition is key to the efficiency of our algorithms, and it generalizes previous conditions from the class of block-decomposable modules to the larger one of rectangle-decomposable modules. It also admits an algebraic formulation that turns out to be a weaker version of the one for block-decomposability. Our analysis focuses on the case of modules indexed over finite grids, the more general cases are left as future work.

2012 ACM Subject Classification Mathematics of computing → Algebraic topology

Keywords and phrases topological data analysis, multiparameter persistence, rank invariant

Digital Object Identifier 10.4230/LIPIcs.SoCG.2020.22

Related Version A full version of the paper is available at <https://arxiv.org/abs/2002.08894>.

1 Introduction

A *persistence module* M over a subset $U \subseteq \mathbb{R}^d$ is a collection of vector spaces $\{M_t\}_{t \in U}$ and linear maps $\rho_s^t := M(s \leq t): M_s \rightarrow M_t$ with the property that ρ_s^s is the identity map and $\rho_t^u \circ \rho_s^t = \rho_s^u$ for all $s \leq t \leq u \in U$. Here $s \leq t$ if and only if $s_i \leq t_i$ for all $i \in \{1, 2, \dots, d\}$. In the language of category theory, a persistence module M is a functor $M: U \rightarrow \mathbf{vec}$ where \mathbf{vec} is the category of vector spaces and the partially ordered set U is considered as a category in the obvious way. In this setting, *morphisms* between persistence modules are natural transformations $M \Rightarrow N$ between functors, defined by collections of linear maps $\{\varphi_t: M_t \rightarrow N_t\}_{t \in U}$ such that $\varphi_t \circ M(s \leq t) = N(s \leq t) \circ \varphi_s$ for all $s \leq t \in U$. Their kernels, images and cokernels, as well as products, direct sums and quotients of persistence modules, are defined pointwise at each index $t \in U$. Similarly, an *isomorphism* between two persistence modules is a natural isomorphism between them. We will refer to the case $d = 1$ as *single-parameter persistence*, and for $d \geq 2$ we will use the term *multi-parameter persistence*.



© Magnus Bakke Botnan, Vadim Lebovici, and Steve Oudot;
licensed under Creative Commons License CC-BY


36th International Symposium on Computational Geometry (SoCG 2020).

Editors: Sergio Cabello and Danny Z. Chen; Article No. 22; pp. 22:1–22:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



22:2 On Rectangle-Decomposable 2-Parameter Persistence Modules

► **Remark 1.** Throughout this paper we will work exclusively with finite-dimensional vector spaces over a fixed field \mathbf{k} . When finite-dimensionality is emphasized we will refer to the persistence module as being pointwise finite-dimensional (pfd).

Single-parameter persistence modules are typically obtained through the application of homology to a filtered topological space. This process is known as *persistent homology* and has found a wide range of applications to the sciences, as well as to other parts of mathematics such as symplectic geometry. See [14, 20] for an introduction to persistent homology. What makes such persistence modules particularly amenable to data analysis is that they can be completely described by multisets of intervals in \mathbb{R} called *barcodes* [12]. Such a collection of intervals can then in turn be used to extract topological information from the data at hand, and further utilized in statistics and machine learning. We now give an example of this structure theorem in the simple case of $U = \{1, 2, 3\} \subseteq \mathbb{R}$.

► **Example 2.** Consider the following sequence of vector spaces and linear maps

$$\mathbf{k}^2 \xrightarrow{\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}} \mathbf{k}^2 \xrightarrow{[1 \ -1]} \mathbf{k}.$$

By replacing the basis $\{e_1, e_2\}$ of the middle vector space \mathbf{k}^2 with the basis $\{e_1, e_1 + e_2\}$ we get the following matrix representations of the linear maps

$$\mathbf{k}^2 \xrightarrow{\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}} \mathbf{k}^2 \xrightarrow{[1 \ 0]} \mathbf{k} = \left(\mathbf{k} \xrightarrow{1} \mathbf{k} \xrightarrow{1} \mathbf{k} \right) \oplus \left(\mathbf{k} \xrightarrow{1} \mathbf{k} \rightarrow 0 \right).$$

The two persistence modules on the right-hand side are uniquely specified by their supports $\{1, 2, 3\}$ and $\{1, 2\}$, respectively. Their supports give rise to the barcode which in this case is given by $\{\{1, 2, 3\}, \{1, 2\}\}$.

As illustrated by Example 2, a persistence module can be recovered from its barcode thanks to the notion of *indicator modules*: for $X \times Y \subseteq \mathbb{R}^2$ and a subset $Q \subseteq X \times Y$, the indicator module of Q , denoted \mathbf{k}_Q , is defined by

$$\mathbf{k}_{Q,t} = \begin{cases} \mathbf{k} & (t \in Q) \\ 0 & (t \notin Q) \end{cases} \quad \mathbf{k}_Q(s \leq t) = \begin{cases} \text{Id}_{\mathbf{k}} & \text{if } s \text{ and } t \in Q, \\ 0 & \text{else.} \end{cases}$$

By convention, we set $\mathbf{k}_\emptyset = 0$. A persistence module is an *interval module* if it is the indicator module of an interval¹. Note that, just as choosing a basis for a vector space is not canonical, there may be many ways of decomposing a single-parameter persistence module into a direct sum of such interval modules. However, just as for the dimension of a finite-dimensional vector space, the associated barcode given by the multiset of interval supports of the summands is independent of the chosen decomposition [1].

Another desirable property of single-parameter persistence modules M is that they are completely described up to isomorphism by the *rank invariant*, i.e. the collection of ranks $r(s, t) = \text{rank}(M(s \leq t))$ for all $s \leq t$. This can easily be verified in the previous example, and more generally, for any pfd persistence module M indexed over a finite set $\llbracket 1, n \rrbracket$, the following inclusion-exclusion formula (also known as the *persistence measure* [9, 11]) gives the multiplicity $m(s, t)$ of any interval $\llbracket s, t \rrbracket$ in the barcode of M :

$$m(s, t) = r(s, t) - r(s-1, t) - r(s, t+1) + r(s-1, t+1). \quad (1)$$

¹ In the poset $X \times Y$, we say that Q is an interval if it is convex and zigzag path-connected, i.e. if between any two points $p, q \in Q$, there exists a zigzag path $p \leq p_1 \geq p_2 \leq \dots \geq p_n \leq q$ with p_i 's in Q .

Many applications do however naturally come equipped with multiple parameters, and for such applications it is natural to consider multi-parameter persistence modules, see e.g. the introduction of [2] for an example of how multi-parameter persistence connects to hierarchical clustering. Let us first consider the simplest instantiation of 2-parameter persistence modules, namely modules indexed by the square $S = \{a = (0, 0), b = (1, 0), c = (0, 1), d = (1, 1)\} \subseteq \mathbb{R}^2$.

► **Example 3.** The persistence module on the left-hand side below can be transformed into the one on the right-hand side via a change of basis at the vertices:

$$\begin{array}{ccc}
 \mathbf{k}^2 & \xrightarrow{\begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}} & \mathbf{k}^2 \\
 \uparrow \begin{pmatrix} 1 \\ 1 \end{pmatrix} & & \uparrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\
 \mathbf{k} & \xrightarrow{1} & \mathbf{k}
 \end{array}
 \qquad
 \begin{array}{ccc}
 \mathbf{k}^2 & \xrightarrow{\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}} & \mathbf{k}^2 \\
 \uparrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} & & \uparrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\
 \mathbf{k} & \xrightarrow{1} & \mathbf{k}
 \end{array}$$

In turn, the persistence module on the right-hand side is the direct sum

$$\begin{array}{ccc}
 \mathbf{k} & \xrightarrow{1} & \mathbf{k} \\
 \uparrow 0 & & \uparrow 0 \\
 0 & \xrightarrow{0} & 0
 \end{array}
 \oplus
 \begin{array}{ccc}
 \mathbf{k} & \xrightarrow{1} & \mathbf{k} \\
 \uparrow 1 & & \uparrow 1 \\
 \mathbf{k} & \xrightarrow{1} & \mathbf{k}
 \end{array}$$

Just as in Example 2, these persistence modules are completely defined by their support. We define the barcode of the aforementioned persistence module to be the (multi-)set of supports of its summands, namely $\{\{c, d\}, \{a, b, c, d\}\}$.

Although commutative diagrams like the one in the previous example may appear unwieldy at first glance, such persistence modules can – just as in the single-parameter case – be completely described (up to isomorphism) by a multiset of elements from

$$\mathcal{I} := \{\{a\}, \{b\}, \{c\}, \{d\}, \{a, b\}, \{a, c\}, \{b, d\}, \{c, d\}, \{a, b, c\}, \{b, c, d\}, \{a, b, c, d\}\}, \tag{2}$$

called *intervals* in the 2×2 grid. See e.g. Figure 13 in [15]. However, in contrast to the single-parameter case, the rank invariant on persistence modules indexed by S is no longer a *complete* invariant, i.e. it does not fully determine the isomorphism type of such modules. For instance, two persistence modules with barcodes $\{\{a, b, c\}, \{a\}\}$ and $\{\{a, b\}, \{a, c\}\}$ are non-isomorphic yet exhibit the same rank invariant.

► **Example 4.** Consider the following two persistence modules:

$$\begin{array}{ccccc}
 \mathbf{k} & \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} & \mathbf{k}^2 & \xrightarrow{\begin{pmatrix} 1 & 0 \end{pmatrix}} & \mathbf{k} \\
 \uparrow 0 & & \uparrow \begin{pmatrix} 1 \\ 0 \end{pmatrix} & & \uparrow 1 \\
 0 & \xrightarrow{0} & \mathbf{k} & \xrightarrow{1} & \mathbf{k}
 \end{array}
 \qquad
 \begin{array}{ccccc}
 \mathbf{k} & \xrightarrow{\begin{pmatrix} 1 \\ 1 \end{pmatrix}} & \mathbf{k}^2 & \xrightarrow{\begin{pmatrix} 1 & 0 \end{pmatrix}} & \mathbf{k} \\
 \uparrow 0 & & \uparrow \begin{pmatrix} 1 \\ 0 \end{pmatrix} & & \uparrow 1 \\
 0 & \xrightarrow{0} & \mathbf{k} & \xrightarrow{1} & \mathbf{k}
 \end{array}$$

The diagram to the left can easily be seen to be composed of two interval summands in the 3×2 grid. By contrast, the diagram to the right is indecomposable: there exists no change of basis for which this persistence module can be written as a direct sum of persistence modules in a non-trivial way. Again, the two modules have the same rank invariant.

22:4 On Rectangle-Decomposable 2-Parameter Persistence Modules

In the setting of no more than four columns and two rows, results from the field of representation theory of quivers show that there exists a finite set of building blocks (indecomposable modules) from which every persistence module can be built (via direct sums, and up to isomorphism). Based on this, one can associate a well-defined barcode-like structure to such a module by counting the multiplicity of every summand in the decomposition. The inclusion of such grids into topological data analysis was inspired by a problem in materials science [15]. For five or more columns the theory becomes increasingly complex. In particular, for six or more columns there is no way to parametrize a set of building blocks in any reasonable way². This is a major obstacle to the development of the theory of multi-parameter persistence.

A natural question to consider then is whether one can endow multi-parameter persistence modules with additional structure in order to enforce nice decomposition theorems akin to that of single-parameter persistence. One such setting coming from computational topology was identified in [3, 8], and further generalized in [10], where it is shown that the so-called *strongly exact* 2-parameter persistence modules indexed over \mathbb{R}^2 are determined (up to isomorphism) by a multiset of particularly simple planar rectangular regions called *blocks*. Basically, a block is either an upper-right or lower-left quadrant, or a horizontal or vertical infinite band. The great advantage of this condition is that it can be checked locally: a 2-parameter persistence module (called a *bimodule* for short) is block-decomposable if, and only if, its restriction to any square as in Example 3 is block-decomposable.

Contributions. In this paper we address two important follow-up questions:

- Can we work out conditions such as above for larger classes of bimodules?
- Can we identify classes of bimodules for which the rank invariant is complete?

Our answers to both questions are positive, and the two classes of bimodules turn out to be the same, namely that of *rectangle-decomposable* bimodules, which by definition are determined (up to isomorphism) by a multiset of *rectangles*, i.e subsets R of the form $R = I \times J \subseteq \mathbb{R}^2$ where I and J are intervals in \mathbb{R} . Specifically, a bimodule is rectangle-decomposable if it decomposes into a direct sum of *rectangle modules*, i.e. indicator modules of rectangles.

Our local condition for rectangle decomposability, called *weak exactness*, is a weaker version of the condition for block decomposability, in that it allows all types of rectangular shapes in the local squares' decompositions, as opposed to just blocks. More precisely, calling \mathcal{R} the following subset of \mathcal{I} :

$$\mathcal{R} = \{\{a\}, \{b\}, \{c\}, \{d\}, \{a, b\}, \{a, c\}, \{b, d\}, \{c, d\}, \{a, b, c, d\}\}, \quad (3)$$

► **Definition 5** (Weak exactness). *Given subsets X, Y of \mathbb{R} , a persistence module $M: X \times Y \subseteq \mathbb{R} \times \mathbb{R} \rightarrow \text{vec}$ is weakly exact if the barcode of the following square*

$$\begin{array}{ccc} M_{(s_x, t_y)} & \xrightarrow{\rho_{(s_x, t_y)}^t} & M_t \\ \rho_s^{(s_x, t_y)} \uparrow & & \uparrow \rho_{(t_x, s_y)}^t \\ M_s & \xrightarrow{\rho_s^{(t_x, s_y)}} & M_{(t_x, s_y)} \end{array} \quad (4)$$

consists of elements from \mathcal{R} for all indices $s \leq t$ in $X \times Y$.

By comparison, the strong exactness condition replaces \mathcal{R} by $\mathcal{B} = \mathcal{R} \setminus \{\{b\}, \{c\}\}$.

² The underlying graph, called a quiver, is known to be of *wild* representation type.

► **Example 6.** The persistence module to the left below is strongly exact, while the one to the right is only weakly exact, and the persistence modules in Example 4 are not even weakly exact (each time the weak or strong exactness condition fails on the outermost rectangle):

$$\begin{array}{ccccc}
 \mathbf{k} & \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} & \mathbf{k}^2 & \xrightarrow{\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}} & \mathbf{k}^2 \\
 \uparrow 0 & & \uparrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} & & \uparrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\
 0 & \xrightarrow{0} & \mathbf{k} & \xrightarrow{1} & \mathbf{k}
 \end{array}
 \qquad
 \begin{array}{ccccc}
 \mathbf{k} & \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} & \mathbf{k}^2 & \xrightarrow{\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}} & \mathbf{k}^2 \\
 \uparrow 0 & & \uparrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} & & \uparrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\
 0 & \xrightarrow{0} & \mathbf{k} & \xrightarrow{1} & \mathbf{k}
 \end{array}$$

Our analysis focuses on the case of modules indexed over finite grids³, the more general cases are left as future work. Our contributions summarize as follows:

- In Section 2 we prove that the rank invariant is complete on the class of rectangle-decomposable bimodules (Theorem 8). To this end, we generalize the inclusion-exclusion formula (1) to our setting. Note that our result also follows indirectly from an inclusion-exclusion formula for a generalization of the rank invariant for interval-decomposable modules [17, Prop. 7.13], but that we provide an explicit statement together with a simple and direct proof.
- In Section 3 we show that the rank invariant of a (1-critical) simplicial bifiltration with a total of n simplices can be computed in $O(n^4)$ time (Theorem 9). In conjunction with our inclusion-exclusion formula, this yields an $O(n^4)$ time algorithm for computing the barcode of a persistence bimodule that is known to be rectangle-decomposable (Corollary 10). This is an improvement over merely applying some state-of-the-art algorithm for computing decompositions of general 2-parameter persistence modules, which would take $O(n^{2\omega+1})$ time where $2 \leq \omega < 2.373$ is the exponent for matrix multiplication [13].
- In Section 4 we propose an algebraic formulation of our weak exactness condition (Definition 11). This formulation turns out to be equivalent to Definition 5 and to global rectangle decomposability (the central mathematical result in the paper), specifically:
 - **Theorem 7.** *Let M be a pfd persistence module indexed over $X \times Y$, where X, Y are finite subsets of \mathbb{R} . Then, M is rectangle-decomposable if and only if M is weakly exact.*
- In Section 5 we leverage this result to derive an $O(n^{2+\omega})$ -time algorithm for checking the rectangle-decomposability of persistence bimodules induced in homology from (1-critical) simplicial bifiltrations with at most n simplices (Theorem 19). Once again, this is an improvement over applying some state-of-the-art algorithm for computing decompositions of general 2-parameter persistence modules and then checking the summands one by one.
- Finally, in Section 6 we show how rectangle-decomposable modules arise from (sufficiently tame) real-valued functions on a topological space. This is then used to give a new proof of the *pyramid basis theorem* of [3].

2 Completeness of the rank invariant

Suppose in this section that X, Y are subsets of \mathbb{Z} .

► **Theorem 8.** *The isomorphism type of any pfd rectangle-decomposable persistence module M over $X \times Y$ is fully determined by the rank invariant of M .*

³ A *finite grid* is the product of two finite subsets of \mathbb{R} . Note that any finite grid can be identified with a grid of the form $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$ for appropriate choices of $n, m \in \mathbb{N}$.

The proof consists in showing that the multiplicity $m(s, t)$ of each individual rectangle module $\mathbf{k}_{\llbracket s_x, t_x \rrbracket \times \llbracket s_y, t_y \rrbracket}$ in the decomposition of M is given by the inclusion-exclusion formula (7) below, which involves only the rank invariant $r : X \times Y \rightarrow \mathbb{N}$ of M . This formula is the analogue, in the category of rectangle-decomposable pfd bimodules, of the inclusion-exclusion formula (1) for counting the multiplicities of interval summands in one-parameter persistence.

Fix arbitrary indices $s \leq t \in X \times Y$. Recall that the rank of $(A \oplus B)(s \leq t)$ is equal to the sum of the ranks of $A(s \leq t)$ and $B(s \leq t)$. Meanwhile, for any summand \mathbf{k}_R of M , the rank of $\mathbf{k}_R(s \leq t)$ is 1 if $s, t \in R$ and 0 otherwise. Therefore, $r(s, t)$ counts (with multiplicity) the number of summands of M whose rectangle support contains both s and t . Then, denoting by $m(s, t^+)$ the number of (rectangle) summands whose support contains t and has s as lower-left corner, we have the following inclusion-exclusion formula:

$$m(s, t^+) = r(s, t) - r((s_x - 1, s_y), t) - r((s_x, s_y - 1), t) + r((s_x - 1, s_y - 1), t). \quad (5)$$

This formula can be interpreted as follows: a rectangle containing t has s as lower-left corner if and only if it contains s but neither $(s_x - 1, s_y)$ nor $(s_x, s_y - 1)$; and it contains both $(s_x - 1, s_y)$ and $(s_x, s_y - 1)$ if and only if it contains $(s_x - 1, s_y - 1)$.

Using the same approach at t , we can now compute the number $m(s, t)$ of summands of M whose support has s as lower-left corner and t as upper-right corner (i.e. is the rectangle $\llbracket s_x, t_x \rrbracket \times \llbracket s_y, t_y \rrbracket$). The corresponding inclusion-exclusion formula is:

$$m(s, t) = m(s, t^+) - m(s, (t_x + 1, t_y)^+) - m(s, (t_x, t_y + 1)^+) + m(s, (t_x + 1, t_y + 1)^+). \quad (6)$$

Combining (5) and (6) together gives the desired inclusion-exclusion formula for the multiplicity $m(s, t)$ of the summand $\mathbf{k}_{\llbracket s_x, t_x \rrbracket \times \llbracket s_y, t_y \rrbracket}$ in the decomposition of M from the rank invariant, hence completing the proof of Theorem 8, namely:

$$\begin{aligned} m(s, t) &= r(s, t) - r((s_x - 1, s_y), t) \\ &\quad - r((s_x, s_y - 1), t) + r((s_x - 1, s_y - 1), t) \\ &\quad - r(s, (t_x + 1, t_y)) + r((s_x - 1, s_y), (t_x + 1, t_y)) \\ &\quad + r((s_x, s_y - 1), (t_x + 1, t_y)) - r((s_x - 1, s_y - 1), (t_x + 1, t_y)) \\ &\quad - r(s, (t_x, t_y + 1)) + r((s_x - 1, s_y), (t_x, t_y + 1)) \\ &\quad + r((s_x, s_y - 1), (t_x, t_y + 1)) - r((s_x - 1, s_y - 1), (t_x, t_y + 1)) \\ &\quad + r(s, (t_x + 1, t_y + 1)) - r((s_x - 1, s_y), (t_x + 1, t_y + 1)) \\ &\quad - r((s_x, s_y - 1), (t_x + 1, t_y + 1)) + r((s_x - 1, s_y - 1), (t_x + 1, t_y + 1)). \end{aligned} \quad (7)$$

3 Computing the rank invariant and rectangle decompositions

Let F be a simplicial bifiltration with n simplices in total. Assume without loss of generality that F is indexed over the grid $G = \llbracket 1, n \rrbracket \times \llbracket 1, n \rrbracket$, for any larger indexing grid must contain arrows with identity maps that can be pre- or post-composed, and any smaller grid can be enlarged by inserting arrows with identity maps. Assume further that F is 1-critical, meaning that the set $\{t \in G \mid \sigma \in F_t\}$ has a unique minimal element for any simplex σ entering the filtration. We also fix a homology degree p .

► **Theorem 9.** *Given the above input, the rank invariant of the persistence bimodule M induced by F in p -th homology can be computed in $O(n^4)$ time.*

As we have not seen any proof of this result in the literature, below we provide an algorithm with the desired complexity. But before this, let us point out that this theorem, combined with the inclusion-exclusion formula (7), gives an $O(n^4)$ -time algorithm to compute the barcode of F assuming that M is rectangle-decomposable: once the rank invariant of M has been computed, iterate over all the pairs (s, t) with $s \leq t \in G$ and, for each one of them, apply the formula in constant time to get the multiplicity of the rectangle module $\mathbf{k}_{\llbracket s_x, t_x \rrbracket \times \llbracket s_y, t_y \rrbracket}$ in the decomposition of M . Thus,

► **Corollary 10.** *Computing the decomposition of a rectangle-decomposable module over $X \times Y$ induced in homology by a 1-critical bifiltration with n simplices in total can be done in $O(n^4)$ time.*

This complexity compares favorably to that of the currently best known algorithm for computing direct-sum decompositions of general persistence bimodules⁴, which is $O(n^{2\omega+1})$ where $2 \leq \omega < 2.373$ is the exponent for matrix multiplication [13].

Let us now provide the algorithm for Theorem 9. First, we compute a free resolution of M of size $O(n^2)$ from F in $O(n^3)$ time using the algorithm of [18]. This free resolution takes the form of an exact sequence as follows⁵:

$$0 \longrightarrow M_\zeta \xrightarrow{\psi} M_\eta \xrightarrow{\varphi} M_\gamma \longrightarrow M \longrightarrow 0,$$

where M_γ , M_η and M_ζ are free bigraded modules, equipped with bases $(\gamma_1, \dots, \gamma_k)$, (η_1, \dots, η_l) , $(\zeta_1, \dots, \zeta_m)$ respectively, of sizes $k, l, m \leq n$. The elements γ_i are called *generators*, while the η_j are called *relations* and the ζ_r are called *relations on relations*. Each one of them is homogeneous and thus assigned a unique grade, denoted by $\text{gr}(\cdot)$. The morphisms of free bigraded modules φ, ψ are given as $k \times l$ and $l \times m$ matrices respectively, with coefficients in \mathbf{k} . Exactness implies that ψ is injective and $M \cong \text{Coker } \varphi$.

Ignoring first the relations and the relations on relations (i.e. assuming that $l = m = 0$ hence $M_\eta = M_\zeta = 0$), the rank invariant $r : \llbracket 1, n \rrbracket^2 \times \llbracket 1, n \rrbracket^2 \rightarrow \mathbb{N}$ is given by:

$$\forall s \leq t \in G, \quad r(s, t) = \#\{i \mid \text{gr}(\gamma_i) \leq s\}.$$

Computing the numbers on the right-hand side for all $s \leq t \in G$ is easily done in $O(n^4)$ time by dynamic programming: first we iterate over the $k \leq n$ generators to fill in an $n \times n$ table storing at each index s the number of generators having s as their grade; then we build a $n^2 \times n^2$ lookup table by iterating over the indices (s, t) in lexicographic order and using the following recurrence formula⁶:


$$\begin{aligned} \#\{i \mid \text{gr}(\gamma_i) \leq s\} &= \#\{i \mid \text{gr}(\gamma_i) < s\} + \#\{i \mid \text{gr}(\gamma_i) = s\} \\ &= \#\{i \mid \text{gr}(\gamma_i) \leq (s_x - 1, s_y)\} + \#\{i \mid \text{gr}(\gamma_i) \leq (s_x, s_y - 1)\} \\ &\quad - \#\{i \mid \text{gr}(\gamma_i) \leq (s_x - 1, s_y - 1)\} + \#\{i \mid \text{gr}(\gamma_i) = s\}. \end{aligned}$$

Once this is done, we can take the relations into account (still ignoring relations on relations, i.e. assuming that $m = 0$ hence $M_\zeta = 0$). Each relation η_j gives a \mathbf{k} -linear constraint on a subset Υ of the generators, encoded in the j -th column of the matrix of φ : $\sum_{u \in \Upsilon} \alpha_u \gamma_u = 0$,

⁴ Let us also point out that our approach does not suffer from the limitation of the algorithm of [13], which is that no two generators or relations in a minimal presentation of M can have the same grade.

⁵ Such resolutions exist by Hilbert’s Syzygy theorem. As mentioned in [18], although their algorithm only computes a free presentation, it adapts readily to compute a free resolution with the same time complexity – simply reapply their kernel basis computation procedure to the algorithm’s output matrix.

⁶ The formula adapts to the cases where $s_x = 1$ or $s_y = 1$ by merely removing the invalid terms.



■ **Figure 1** A relation η between two generators γ_1, γ_2 , and its effect on $r(s, t)$ for various $s \leq t$.

where each α_u belongs to $\mathbf{k} \setminus \{0\}$. Call $\text{lub}(\eta_j)$ the least upper bound of the set $\{\text{gr}(\gamma_u) \mid u \in \Upsilon\}$ in the product order \leq on G . The effect of η_j on the rank invariant is to decrement it at indices $s \leq t$ such that $s \geq \text{lub}(\eta_j)$ and $t \geq \text{gr}(\eta_j)$, as illustrated in Figure 1. Hence the following formula for the update of $r(s, t)$:

$$r(s, t) \longleftarrow r(s, t) - \#\{j \mid \text{lub}(\eta_j) \leq s \text{ and } \text{gr}(\eta_j) \leq t\}.$$

The numbers on the right-hand side can be computed in $O(n^4)$ time again using dynamic programming: first we build a $n^2 \times n^2$ table storing at each index (s, t) the number of relations η_j such that $\text{gr}(\eta_j) = t$ and $\text{lub}(\eta_j) = s$; then for each index $t \in G$ we build an intermediate $n \times n$ lookup table by iterating over the indices s in lexicographic order and using the following recurrence formula⁶:

$$\begin{aligned} \#\{j \mid \text{lub}(\eta_j) \leq s \text{ and } \text{gr}(\eta_j) = t\} &= \#\{j \mid \text{lub}(\eta_j) \leq (s_x - 1, s_y) \text{ and } \text{gr}(\eta_j) = t\} \\ &\quad + \#\{j \mid \text{lub}(\eta_j) \leq (s_x, s_y - 1) \text{ and } \text{gr}(\eta_j) = t\} \\ &\quad - \#\{j \mid \text{lub}(\eta_j) \leq (s_x - 1, s_y - 1) \text{ and } \text{gr}(\eta_j) = t\} \\ &\quad + \#\{j \mid \text{lub}(\eta_j) = s \text{ and } \text{gr}(\eta_j) = t\}; \end{aligned}$$

finally, we build the lookup table for the rank invariant r by iterating over the indices (s, t) in lexicographic order and using the following recurrence formula⁶:

$$\begin{aligned} \#\{j \mid \text{lub}(\eta_j) \leq s \text{ and } \text{gr}(\eta_j) \leq t\} &= \#\{j \mid \text{lub}(\eta_j) \leq s \text{ and } \text{gr}(\eta_j) \leq (t_x - 1, t_y)\} \\ &\quad + \#\{j \mid \text{lub}(\eta_j) \leq s \text{ and } \text{gr}(\eta_j) \leq (t_x, t_y - 1)\} \\ &\quad - \#\{j \mid \text{lub}(\eta_j) \leq s \text{ and } \text{gr}(\eta_j) \leq (t_x - 1, t_y - 1)\} \\ &\quad + \#\{j \mid \text{lub}(\eta_j) \leq s \text{ and } \text{gr}(\eta_j) = t\}. \end{aligned}$$

Once this is done, we can take the relations on relations into account. Each one of them, say ζ_r , gives a \mathbf{k} -linear constraint on a subset Ξ of the relations, encoded in the r -th column of the matrix of ψ : $\sum_{v \in \Xi} \beta_v \eta_v = 0$, where each β_v belongs to $\mathbf{k} \setminus \{0\}$. Calling $\text{lub}(\zeta_r)$ the least upper bound of the set $\{\text{lub}(\eta_v) \mid v \in \Xi\}$, the effect of ζ_r on the rank invariant is to compensate for one of the relations by incrementing $r(s, t)$ at indices $s \leq t$ such that $s \geq \text{lub}(\zeta_r)$ and $t \geq \text{gr}(\zeta_r)$. Hence the following formula for the update of $r(s, t)$:

$$r(s, t) \longleftarrow r(s, t) + \#\{r \mid \text{lub}(\zeta_r) \leq s \text{ and } \text{gr}(\zeta_r) \leq t\}.$$

The numbers on the right-hand side can be computed in $O(n^4)$ time using the same two-stage dynamic programming scheme as introduced for relations.

All in all, the algorithm takes $O(n^4)$ time. The injectivity of ψ means that the relations on relations are linearly independent, so the correctness of the output table representing the rank invariant follows by design. This completes the proof of Theorem 9.

4 Algebraic formulation of weak exactness

As shown in [5, 10], a persistence module $M : X \times Y \subseteq \mathbb{R} \times \mathbb{R} \rightarrow \text{vec}$ is strongly exact if, and only if, the following sequence induced by (4) is exact for all indices $s \leq t \in X \times Y$:

$$M_s \xrightarrow{(\rho_s^{(t_x, s_y)}, \rho_s^{(s_x, t_y)})} M_{(t_x, s_y)} \oplus M_{(s_x, t_y)} \xrightarrow{\rho_{(t_x, s_y)}^t - \rho_{(s_x, t_y)}^t} M_t. \tag{8}$$

Similarly, we can characterize weak exactness (Definition 5) algebraically:

► **Definition 11** (Algebraic weak exactness). *A persistence module $M : X \times Y \subseteq \mathbb{R} \times \mathbb{R} \rightarrow \text{vec}$ is called algebraically weakly exact if the following equalities hold for all $s \leq t \in X \times Y$:*

$$\begin{aligned} \text{Im } \rho_s^t &= \text{Im } \rho_{(t_x, s_y)}^t \cap \text{Im } \rho_{(s_x, t_y)}^t, \\ \text{Ker } \rho_s^t &= \text{Ker } \rho_s^{(t_x, s_y)} + \text{Ker } \rho_s^{(s_x, t_y)}. \end{aligned}$$

This condition holds in particular when the sequence (8) is exact, but not only. Indeed, as can be checked easily, any rectangle (not just block) module is algebraically weakly exact. So is any rectangle-decomposable pfd persistence bimodule, since the property is obviously preserved under taking direct sums of pfd persistence bimodules. The converse holds as well:

► **Theorem 12** (Decomposition of algebraically weakly exact pfd bimodules). *For any algebraically weakly exact pfd persistence module M over a finite grid $(X \times Y, \leq)$, there is a unique multiset \mathcal{RM} of rectangles of $X \times Y$ such that:*

$$M \cong \bigoplus_{R \in \mathcal{RM}} \mathbf{k}_R.$$

Since this result holds in particular for persistence bimodules indexed over squares, it ensures that a pfd persistence module over a square is algebraically weakly exact if, and only if, it is rectangle-decomposable. Hence the equivalence between weak exactness (Definition 5) and algebraic weak exactness (Definition 11), and the correctness of Theorem 7.

The rest of this section is devoted to the proof of Theorem 12. From this point on, and until the end of the section, whenever we talk about *weak exactness* we refer consistently to the algebraic formulation from Definition 11.

4.1 A preliminary remark concerning submodules and summands

A morphism $f : M \rightarrow N$ between two persistence modules over $(X \times Y, \leq)$ is a *monomorphism* (resp. *epimorphism*) if for every $t \in X \times Y$, $f_t : M_t \rightarrow N_t$ is injective (resp. surjective). We say that a monomorphism $f : M \rightarrow N$ between two persistence modules M and N *splits* if there is a morphism $g : N \rightarrow M$ such that $g \circ f = \text{Id}_M$. If every monomorphism with domain M splits, we say that M is an *injective persistence module*.

It is not true that any submodule of a persistence module is a summand. However, if $f : M \rightarrow N$ is a monomorphism between two persistence modules M and N which splits, it is well known that there is a direct sum decomposition $N \cong M \oplus \text{Coker}(f)$. Therefore, an injective submodule of a persistence module is a summand thereof. In our analysis we will often use the following result:

► **Lemma 13.** *For any indices $k \in \llbracket 1, n \rrbracket$ and $l \in \llbracket 1, m \rrbracket$, the indicator module $\mathbf{k}_{\llbracket 1, k \rrbracket \times \llbracket 1, l \rrbracket}$ is an injective persistence module over $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$.*

Proof. This lemma is a consequence of [5, Lem. 2.1] since the subset $\llbracket 1, k \rrbracket \times \llbracket 1, l \rrbracket$ is clearly a directed ideal of the poset $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$, following the definition of [5, Sec. 2.1]. ◀

4.2 Proof of Theorem 12

Uniqueness of the decomposition follows directly from Krull-Schmidt-Remak-Azumaya's theorem [1], since the endomorphism ring of any rectangle module is clearly isomorphic to \mathbf{k} and thus local. We therefore focus on the existence of a decomposition into rectangle summands. Our proof proceeds by induction on the poset of grid dimensions (n, m) , also viewed as a subposet of \mathbb{R}^2 equipped with the product order:

- Our base cases are when $n = 1$ or $m = 1$. The result is then a direct consequence of Gabriel's theorem [16], which asserts that M decomposes as a direct sum of interval modules, each interval being a rectangle of width 1.
- Fix $n > 1$ and $m > 1$, and assume that the result is true for all grids of sizes $n' \times m'$ such that $(n', m') < (n, m)$. Fix a persistence module M over $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$ that is pfd and weakly exact. Observe that M has finite *total dimension* $\sum_{t \in \llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket} \dim M_t$, so we know from a simple induction that M decomposes as a direct sum of indecomposables – see [5, Theorem 1.1] for a more general statement. As any summand of a weakly exact module is again weakly exact, we may restrict our attention to pfd indecomposable modules. For the sake of contradiction, assume that M is pfd, weakly exact, indecomposable, and not isomorphic to a rectangle module. Then:

► **Lemma 14.** *The map $\rho_{(1,1)}^{(n,m)}$ is zero.*

Proof. Suppose the contrary. Then we have $\text{Ker } \rho_{(1,1)}^{(n,m)} \subsetneq M_{(1,1)}$. Let $\alpha \in M_{(1,1)} \setminus \text{Ker } \rho_{(1,1)}^{(n,m)}$. The submodule N of M spanned by the collection of images $(\rho_{(1,1)}^{(i,j)}(\alpha))_{(i,j) \in \llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket}$ is isomorphic to $\mathbf{k}_{\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket}$, an injective persistence module by Lemma 13. Hence, N is a summand of M , contradicting that M is not isomorphic to a rectangle module. ◀

► **Lemma 15.** *The space $M_{(1,1)}$ maps injectively to the nodes of the grid $\llbracket 1, n-1 \rrbracket \times \llbracket 1, m-1 \rrbracket$.*

Proof. Let us restrict M to the grid $\llbracket 1, n-1 \rrbracket \times \llbracket 1, m \rrbracket$. The restriction – denoted by N – may no longer be indecomposable, however it is still pfd and weakly exact, therefore our induction hypothesis asserts that N decomposes as a finite (internal) direct sum where each summand is isomorphic to some rectangle module. Consider any one of these summands, say $N' \cong \mathbf{k}_{R'}$, such that $(1, 1) \in R'$. Then, we claim that $(n-1, 1) \in R'$ as well. Indeed, otherwise, one can extend N' to a persistence module over $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$ by putting zero spaces on the last column n . This yields an injective rectangle submodule of M (Lemma 13), and therefore a rectangle summand of M – a contradiction.

By our claim, $M_{(1,1)}$ maps injectively to the nodes $(i, 1)$ for $i \in \llbracket 1, n-1 \rrbracket$. Similarly, by restricting M to the grid $\llbracket 1, n \rrbracket \times \llbracket 1, m-1 \rrbracket$, we deduce that $M_{(1,1)}$ maps injectively to the nodes $(1, j)$ for $j \in \llbracket 1, m-1 \rrbracket$. Then, by weak exactness, we have

$$\forall (i, j) \in \llbracket 1, n-1 \rrbracket \times \llbracket 1, m-1 \rrbracket, \text{Ker } \rho_{(1,1)}^{(i,j)} = \text{Ker } \rho_{(1,1)}^{(i,1)} + \text{Ker } \rho_{(1,1)}^{(1,j)} = 0,$$

so $M_{(1,1)}$ maps injectively to all the nodes of the grid $\llbracket 1, n-1 \rrbracket \times \llbracket 1, m-1 \rrbracket$. ◀

► **Lemma 16.** *The spaces $M_{(1,1)}$ and $M_{(n,m)}$ are zero.*

Proof. By weak exactness and Lemma 14, we have

$$M_{(1,1)} = \text{Ker } \rho_{(1,1)}^{(n,m)} = \text{Ker } \rho_{(1,1)}^{(n,1)} + \text{Ker } \rho_{(1,1)}^{(1,m)}.$$

Assuming for a contradiction that $M_{(1,1)} \neq 0$, we have that at least one of the two terms on the right-hand side of the above equation must be non-zero – say $\text{Ker } \rho_{(1,1)}^{(n,1)} \neq 0$. Let $\alpha \neq 0$ be an element in that kernel. By Lemma 15, its images at the nodes of $\llbracket 1, n-1 \rrbracket \times \llbracket 1, m-1 \rrbracket$ are non-zero. Meanwhile, its images at the nodes of $\{n\} \times \llbracket 1, m \rrbracket$ are zero, by composition. There are two cases:

- Either $\rho_{(1,1)}^{(1,m)}(\alpha) = 0$, in which case the images of α at the nodes of $\llbracket 1, n \rrbracket \times \{m\}$ are also zero, which implies that the persistence submodule of M spanned by the images of α is isomorphic to $\mathbf{k}_{\llbracket 1, n-1 \rrbracket \times \llbracket 1, m-1 \rrbracket}$.
- Or $\rho_{(1,1)}^{(1,m)}(\alpha) \neq 0$, in which case, for all $i \in \llbracket 1, n-1 \rrbracket$, we have

$$\alpha \notin \text{Ker } \rho_{(1,1)}^{(1,m)} \stackrel{\text{(Lemma 15)}}{=} \text{Ker } \rho_{(1,1)}^{(1,m)} + \text{Ker } \rho_{(1,1)}^{(i,1)} = \text{Ker } \rho_{(1,1)}^{(i,m)},$$

which implies that the images of α at the nodes of $\llbracket 1, n-1 \rrbracket \times \{m\}$ are non-zero as well. Hence, the persistence submodule of M spanned by the images of α is isomorphic to $\mathbf{k}_{\llbracket 1, n-1 \rrbracket \times \llbracket 1, m \rrbracket}$.

In both cases, the persistence submodule of M spanned by the images of α is an injective rectangle module (Lemma 13), hence a rectangle summand of M – a contradiction.

By applying vector-space duality pointwise to M , we obtain an indecomposable module M^* of the grid $\llbracket 1, n \rrbracket^{\text{op}} \times \llbracket 1, m \rrbracket^{\text{op}}$ – which is isomorphic to $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$ as a poset. This persistence module is still pfd, and still weakly exact as well since the equations of weak exactness are stable under vector-space duality (kernels become images, sums become intersections, and vice-versa). Hence, by the first part of the proof, $M_{(1,1)}^* = 0$, i.e the space at node (n, m) of M is zero, hence the result. ◀

► **Lemma 17.** *The space $M_{(1,m)}$ is zero.*

Proof. Assume for a contradiction that $M_{(1,m)} \neq 0$. Call N the restriction of M to the grid $\llbracket 1, n-1 \rrbracket \times \llbracket 1, m \rrbracket$. By our induction hypothesis, N decomposes as a finite (internal) direct sum where each summand is isomorphic to some rectangle module. Since $M_{(1,m)} \neq 0$, at least one of these rectangles contains the node $(1, m)$. Among such rectangles, take one – say $R' = \llbracket 1, i \rrbracket \times \llbracket j, m \rrbracket$ – that has lowest lower-left corner, and call N' the corresponding summand of N . Denote by N'' the rest of the internal decomposition of N , i.e. $N = N' \oplus N''$.

First, we claim that $i = n-1$. Indeed, otherwise we can extend N' to a rectangle persistence submodule \bar{N}' of M by putting zero spaces on the last column n , and N'' to another persistence submodule \bar{N}'' by putting the internal spaces of M on the last column, so that $M = \bar{N}' \oplus \bar{N}''$ – a contradiction.

Second, we claim that $j \in \llbracket 2, m-1 \rrbracket$. Indeed, $j \geq 2$ since by Lemma 16 we know that $M_{(1,1)} = 0$. Meanwhile, if j were equal to m , then N' would go to zero on the last column of $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$ since $M_{(n,m)} = 0$ by Lemma 16, and so we could extend N to a rectangle persistence submodule \bar{N}' of M by putting zero spaces on the last column, and N'' to another persistence submodule \bar{N}'' by putting the internal spaces of M on the last column, so that $M = \bar{N}' \oplus \bar{N}''$ – a contradiction.

22:12 On Rectangle-Decomposable 2-Parameter Persistence Modules

Consider now the space $N_{(1,j)} = M_{(1,j)}$, and take a generator α of the subspace $N'_{(1,j)} \cong \mathbf{k}$. By Lemma 16 we know that the map $\rho_{(1,j)}^{(n,m)}$ is zero, so by weak exactness we have $\alpha = \alpha_h + \alpha_v$ for some $\alpha_h \in \text{Ker } \rho_{(1,j)}^{(n,j)}$ and $\alpha_v \in \text{Ker } \rho_{(1,j)}^{(1,m)}$. We claim that $\alpha_h \notin N''_{(1,j)}$. Indeed, otherwise we would have

$$\rho_{(1,j)}^{(1,m)}(\alpha) = \rho_{(1,j)}^{(1,m)}(\alpha_h) + \rho_{(1,j)}^{(1,m)}(\alpha_v) = \rho_{(1,j)}^{(1,m)}(\alpha_h) \in \rho_{(1,j)}^{(1,m)}(N''_{(1,j)}) \subseteq N''_{(1,m)},$$

thus contradicting our assumption that $N = N' \oplus N''$ with the support of N' containing $(1, m)$. Likewise, for any node $t \in R'$ we have $\rho_{(1,j)}^t(\alpha_h) \notin N''_t$, for otherwise we would get a contradiction from

$$\rho_{(1,j)}^{(t_x,m)}(\alpha) = \rho_{(1,j)}^{(t_x,m)}(\alpha_h) = \rho_t^{(t_x,m)}(\rho_{(1,j)}^t(\alpha_h)) \in \rho_t^{(t_x,m)}(N''_t) \subseteq N''_{(t_x,m)}.$$

Thus, the persistence submodule N^h of N generated by α_h is isomorphic⁷ to N' and in direct sum with N'' . We can therefore exchange N' for N^h in the internal decomposition of N . Since N^h is mapped to zero on the last column of $\llbracket 1, n \rrbracket \times \llbracket 1, m \rrbracket$, we can extend it to a rectangle persistence submodule \bar{N}^h of M by putting zero spaces on the last column, meanwhile we can extend N'' to another persistence submodule \bar{N}'' by putting the internal spaces of M on the last column, so that $M = \bar{N}^h \oplus \bar{N}''$ – a contradiction. ◀

► **Lemma 18.** $M_{(1,j)} = 0$ for all $j \in \llbracket 1, m \rrbracket$.

Proof. The result is already proven⁸ for $j = m$ by Lemma 17. Let then $j \in \llbracket 1, m - 1 \rrbracket$. Call N the restriction of M to the grid $\llbracket 1, n \rrbracket \times \llbracket 1, m - 1 \rrbracket$. By our induction hypothesis, N decomposes as a finite (internal) direct sum where each summand is isomorphic to some rectangle module. Assuming for a contradiction that some summand N' has a support R' that intersects the first column, we know from Lemma 17 that N' maps to zero at node $(1, m)$. By composition, N' maps to zero as well at the nodes on the last row m . Therefore, as in the proof of Lemma 17, we can extend N' to a rectangle summand of M by putting zero spaces on row m , thus reaching a contradiction. ◀

It follows from Lemma 18 that M itself is not supported outside the rectangle $R = \llbracket 2, n \rrbracket \times \llbracket 1, m \rrbracket$. The induction hypothesis (applied to the restriction of M to R) implies then that M decomposes as a direct sum of rectangle modules, which raises a contradiction. This concludes the induction step and the proof of Theorem 12.

5 Algorithm for checking rectangle decomposability

As in Section 3, let F be a simplicial bifiltration with n simplices in total, and let us assume without loss of generality that F is indexed over the grid $G = \llbracket 1, n \rrbracket \times \llbracket 1, n \rrbracket$. We further assume that F is 1-critical, and we fix a homology degree p .

Given this input, how fast can we check whether the persistence bimodule M induced in p -th homology decomposes into rectangle summands? An obvious solution is to first decompose M from the data of F , then to check the summands one by one. As explained in Section 3, the currently best known algorithm for decomposition runs in time $O(n^{2\omega+1})$, where $2 \leq \omega < 2.373$ is the exponent for matrix multiplication [13]. The advantage of the algebraic weak exactness condition from Section 4 is that it can be checked locally, which reduces the total running time to $O(n^{2+\omega})$. Below we sketch the algorithm:

⁷ Note that we do not need to check that α_h goes to zero when leaving R' , since by assumption R' reaches row m and, as we saw earlier, $i = n - 1$ so R' reaches column $n - 1$ as well.

⁸ It is also proven for $j = 1$ by Lemma 16, although we do not use this fact in the proof.

1. Compute the rank invariant $r : \llbracket 1, n \rrbracket^2 \times \llbracket 1, n \rrbracket^2 \rightarrow \mathbb{N}$ of M .
2. Compute invariants for kernels and images, denoted by $\kappa : \llbracket 1, n \rrbracket^2 \times \llbracket 1, n \rrbracket^2 \rightarrow \mathbb{N}$ and $\iota : \llbracket 1, n \rrbracket^2 \times \llbracket 1, n \rrbracket^2 \rightarrow \mathbb{N}$ respectively, which return the dimensions of $\text{Ker } \rho_s^{(s_x, t_y)} + \text{Ker } \rho_s^{(t_x, s_y)}$ and of $\text{Im } \rho_{(s_x, t_y)}^t \cap \text{Im } \rho_{(t_x, s_y)}^t$ respectively at indices $s \leq t$, and zero elsewhere.
3. For each pair of indices $s \leq t$, check whether $r(s, t) = \iota(s, t)$ and $r(s, s) - r(s, t) = \kappa(s, t)$. If any such equality fails, then answer that M is not rectangle-decomposable. Otherwise, answer that M is rectangle-decomposable.

We now provide further implementation details and analyze the algorithm on the fly:

Step 1 has already been detailed in Section 3 and runs in $O(n^4)$ time.

Step 3 obviously runs in $O(n^4)$ time, and its correctness comes from the commutativity of the square in (4): indeed, commutativity implies that $\text{Im } \rho_s^t \subseteq \text{Im } \rho_{(s_x, t_y)}^t \cap \text{Im } \rho_{(t_x, s_y)}^t$ and $\text{Ker } \rho_s^{(s_x, t_y)} + \text{Ker } \rho_s^{(t_x, s_y)} \subseteq \text{Ker } \rho_s^t$, so checking weak exactness for this square amounts to checking equality between the dimensions of the various spaces involved, hence the equations.

For Step 2, we first compute, for each $t = (j, l) \in G$, the barcode of the zigzag module⁹ induced in homology by the following zigzag of simplicial complexes:

$$F_{(1,l)} \longrightarrow \cdots \longrightarrow F_{(j-1,l)} \longrightarrow F_t \longleftarrow F_{(j,l-1)} \longleftarrow \cdots \longleftarrow F_{(j,1)}. \tag{9}$$

We then do the same with the following zigzag, for each $s = (i, k) \in G$:

$$F_{(i,n)} \longleftarrow \cdots \longleftarrow F_{(i,k+1)} \longleftarrow F_s \longrightarrow F_{(i+1,k)} \longrightarrow \cdots \longrightarrow F_{(n,k)}. \tag{10}$$

Then, for each indices $(i, k) = s \leq t = (j, l)$, by restriction, the dimension of $\text{Im } \rho_{(i,l)}^t \cap \text{Im } \rho_{(j,k)}^t$ is given by the number of intervals in the barcode of (9) that span the sub-zigzag $F_{(i,l)} \longrightarrow F_t \longleftarrow F_{(j,k)}$, while (dually) the dimension of $\text{Ker } \rho_s^{(i,l)} + \text{Ker } \rho_s^{(j,k)}$ is given by $r(s, s)$ minus the number of intervals in the barcode of (10) that span the sub-zigzag $F_{(i,l)} \longleftarrow F_s \longrightarrow F_{(j,k)}$ (the proof of these simple facts is given in [6, Appendix A]). Regarding the running time: since the zigzags (9)-(10) involve $O(n)$ simplex insertions and deletions each, their barcode computation takes $O(n^\omega)$ using the algorithm based on fast matrix multiplication [19]. Then, each barcode having $O(n)$ intervals, the computation of the dimensions and their storage in tables of integers representing the invariants κ and ι takes $O(n)$. This is true for each choice of indices $s \leq t$, hence a total running time in $O(n^{2+\omega} + n^3) = O(n^{2+\omega})$. As a consequence,

► **Theorem 19.** *Checking the rectangle-decomposability of the bimodule induced in homology by a 1-critical bifiltration with n simplices in total can be done in $O(n^{2+\omega})$ time.*

6 An example of rectangle-decomposable module

In [8] the authors show that a large pyramidal diagram can be associated to a sufficiently tame real valued function $f : X \rightarrow \mathbb{R}$. We briefly recall their construction. Under the assumption that the function is of *Morse type*, there exists a finite set of *critical values* $a_1 < a_2 < \dots < a_n$, and we may choose real values s_i satisfying

$$-\infty < s_0 < a_1 < s_1 < \dots < a_n < s_n < +\infty. \tag{11}$$

⁹ A zigzag module is a persistence module indexed over a poset of the form $\bullet \longleftrightarrow \bullet \longleftrightarrow \dots \longleftrightarrow \bullet$, where double-headed arrows mean that the actual arrows can be oriented either forward or backward. Such modules always decompose into direct sums of interval modules [4, 7].

22:14 On Rectangle-Decomposable 2-Parameter Persistence Modules

Here the idea is that the preimage of $[s_i, s_{i+1}]$ deformation retracts onto the fiber over a_{i+1} , and that the fiber is constant (up to homotopy) between critical values. This gives a way of studying how the topology of the fibers connect across scales.

Denoting $X_i^j = f^{-1}[s_i, s_j]$ and ${}^j_i X = X_0^i \cup X_j^n$, obvious inclusions yield a commutative diagram, such as the following one for $n = 2$:

$$\begin{array}{ccccccccc}
 & & & & (X_0^2, X_0^2) & \longrightarrow & (X_0^2, {}^2_2 X) & & \\
 & & & & \uparrow & & \uparrow & & \\
 & & & & (X_0^1, X_0^1) & \longrightarrow & (X_0^2, X_0^1) & \longrightarrow & (X_0^2, {}^2_1 X) & \longrightarrow & (X_0^2, {}^1_1 X) \\
 & & & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 (X_0^0, X_0^0) & \longrightarrow & (X_0^1, X_0^0) & \longrightarrow & (X_0^2, X_0^0) & \longrightarrow & (X_0^2, {}^2_0 X) & \longrightarrow & (X_0^2, {}^1_0 X) & \longrightarrow & (X_0^2, {}^0_0 X) \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 X_0^0 & \longrightarrow & X_0^1 & \longrightarrow & X_0^2 & \longrightarrow & (X_0^2, X_2^2) & \longrightarrow & (X_0^2, X_1^2) & \longrightarrow & (X_0^2, X_0^2) \\
 & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & & X_1^1 & \longrightarrow & X_1^2 & \longrightarrow & (X_1^2, X_2^2) & \longrightarrow & (X_1^2, X_1^2) & & \\
 & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
 & & X_2^2 & \longrightarrow & (X_2^2, X_2^2) & & & & & &
 \end{array}$$

Building on the work of [8], it is shown in [3] that the above diagram, upon application of homology, decomposes into a direct sum of interval modules, where each interval is the intersection of a rectangle in \mathbb{Z}^2 with the pyramid above. This result is referred to as the *pyramid basis theorem*. We now give a new proof of this fact using Theorem 12. More precisely, we show the following:

► **Theorem 20** (Pyramid basis theorem). *The homology pyramid as constructed in [8] is interval-decomposable, where the intervals are restrictions of rectangles in \mathbb{Z}^2 to the pyramid.*

To simplify the notation we prove the case for $n = 2$ and it will be evident that the argument generalizes. First, extend the homology diagram to a bimodule on a finite grid as follows:

$$\begin{array}{cccccccccccc}
 0 & \longrightarrow & 0 & \longrightarrow & H_p(X_0^2, X_0^2) & \longrightarrow & H_p(X_0^2, {}^2_2 X) & \longrightarrow & \text{PO}_2 & \longrightarrow & \text{PO}_3 \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & H_p(X_0^1, X_0^1) & \longrightarrow & H_p(X_0^2, X_0^1) & \longrightarrow & H_p(X_0^2, {}^2_1 X) & \longrightarrow & H_p(X_0^2, {}^1_1 X) & \longrightarrow & \text{PO}_1 \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 H_p(X_0^0, X_0^0) & \longrightarrow & H_p(X_0^1, X_0^0) & \longrightarrow & H_p(X_0^2, X_0^0) & \longrightarrow & H_p(X_0^2, {}^2_0 X) & \longrightarrow & H_p(X_0^2, {}^1_0 X) & \longrightarrow & H_p(X_0^2, {}^0_0 X) \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 H_p(X_0^0) & \longrightarrow & H_p(X_0^1) & \longrightarrow & H_p(X_0^2) & \longrightarrow & H_p(X_0^2, X_2^2) & \longrightarrow & H_p(X_0^2, X_1^2) & \longrightarrow & H_p(X_0^2, X_0^2) \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 \text{PB}_1 & \longrightarrow & H_p(X_1^1) & \longrightarrow & H_p(X_1^2) & \longrightarrow & H_p(X_1^2, X_2^2) & \longrightarrow & H_p(X_1^2, X_1^2) & \longrightarrow & 0 \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 \text{PB}_3 & \longrightarrow & \text{PB}_2 & \longrightarrow & H_p(X_2^2) & \longrightarrow & H_p(X_2^2, X_2^2) & \longrightarrow & 0 & \longrightarrow & 0
 \end{array}$$

Here PB_i denotes pullback and PO_i denotes pushout. Inductively these are defined (up to canonical isomorphism) by

$$\text{PB}_1 = \ker (H_p(X_0^0) \oplus H_p(X_1^1) \rightarrow H_p(X_0^1))$$

$$\text{PB}_2 = \ker (H_p(X_1^1) \oplus H_p(X_2^2) \rightarrow H_p(X_1^2))$$

$$\text{PB}_3 = \ker (\text{PB}_1 \oplus \text{PB}_2 \rightarrow H_p(X_1^1)).$$

and dually for the pushouts, with kernels replaced by cokernels. By Theorem 12 it suffices to show that the extended diagram is weakly exact. The fact that any square with vertices on the original "pyramid" is strongly exact (i.e. the sequence (8) induced by such a square is exact) follows from the exactness of the relative Mayer–Vietoris sequence. Moreover, as remarked in [5, Section 5.1], the extension of the pyramid to pullbacks and pushouts preserves strong exactness (and thus weak exactness). It remains to consider squares with a 0 vector space as either its top-left or bottom-right corner. The fact that such squares are weakly exact is an easy consequence of commutativity. We conclude that the bimodule shown above is weakly exact and therefore rectangle-decomposable. The restrictions of the rectangle summands to the original homology pyramid give the intervals in the *pyramid basis theorem*.

References

- 1 Gorô Azumaya. Corrections and supplementaries to my paper concerning Krull-Remak-Schmidt's theorem. *Nagoya Mathematical Journal*, 1:117–124, 1950.
- 2 Ulrich Bauer, Magnus B Botnan, Steffen Oppermann, and Johan Steen. Cotorsion torsion triples and the representation theory of filtered hierarchical clustering. *arXiv preprint*, 2019. [arXiv:1904.07322](#).
- 3 Paul Bendich, Herbert Edelsbrunner, Dmitriy Morozov, Amit Patel, et al. Homology and robustness of level and interlevel sets. *Homology, Homotopy and Applications*, 15(1):51–72, 2013.
- 4 Magnus Bakke Botnan. Interval decomposition of infinite zigzag persistence modules. *Proceedings of the American Mathematical Society*, 145(8):3571–3577, 2017.
- 5 Magnus Bakke Botnan and William Crawley-Boevey. Decomposition of persistence modules. *To appear in the Proceedings of the AMS*, 2018. [arXiv:1811.08946](#).
- 6 Magnus Bakke Botnan, Vadim Lebovici, and Steve Oudot. On rectangle-decomposable 2-parameter persistence modules, 2020. [arXiv:2002.08894](#).
- 7 Gunnar Carlsson and Vin De Silva. Zigzag persistence. *Foundations of computational mathematics*, 10(4):367–405, 2010.
- 8 Gunnar Carlsson, Vin De Silva, and Dmitriy Morozov. Zigzag persistent homology and real-valued functions. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 247–256. ACM, 2009.
- 9 Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. Springer, 2016.
- 10 Jérémy Cochoy and Steve Oudot. Decomposition of exact pfd persistence bimodules. *Discrete and Computational Geometry*, 2019. To appear, currently available as [arXiv preprint arXiv:1605.09726](#).
- 11 David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
- 12 William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and its Applications*, 14(05):1550066, 2015.
- 13 Tamal K Dey and Cheng Xin. Generalized persistence algorithm for decomposing multi-parameter persistence modules. *arXiv preprint arXiv:1904.03766*, 2019.
- 14 Herbert Edelsbrunner and John Harer. Persistent homology—a survey. *Contemporary mathematics*, 453:257–282, 2008.
- 15 Emerson G Escolar and Yasuaki Hiraoka. Persistence modules on commutative ladders of finite type. *Discrete & Computational Geometry*, 55(1):100–157, 2016.
- 16 Peter Gabriel. Unzerlegbare Darstellungen I. *manuscripta mathematica*, 6(1):71–103, March 1972. [doi:10.1007/BF01298413](#).
- 17 Woojin Kim and Facundo Memoli. Generalized persistence diagrams for persistence modules over posets. *arXiv preprint arXiv:1810.11517*, 2018.

22:16 On Rectangle-Decomposable 2-Parameter Persistence Modules

- 18 Michael Lesnick and Matthew Wright. Computing minimal presentations and betti numbers of 2-parameter persistent homology. *arXiv preprint arXiv:1902.05708*, 2019.
- 19 Nikola Milosavljević, Dmitriy Morozov, and Primoz Skraba. Zigzag persistent homology in matrix multiplication time. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 216–225. ACM, 2011.
- 20 Steve Y Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209. American Mathematical Society Providence, RI, 2015.