Geometry & Topology

Volume 28 (2024)

Configuration spaces of disks in a strip, twisted algebras, persistence, and other stories

Hannah Alpert
Fedor Manin

msp

Geometry & Topology **28**:2 (2024) 641–699
DOI: 10.2140/gt.2024.28.641
Published: 13 March 2024

Configuration spaces of disks in a strip, twisted algebras, persistence, and other stories

HANNAH ALPERT
FEDOR MANIN

We give $\mathbb{Z}$–bases for the homology and cohomology of the configuration space of n unit disks in an infinite strip of width w, first studied by Alpert, Kahle and MacPherson. We also study the way these spaces evolve both as n increases (using the framework of representation stability) and as w increases (using the framework of persistent homology). Finally, we include some results about the cup product in the cohomology and about the configuration space of unordered disks.

55R80; 16S15, 18A25, 55N31, 57Q70

1. Introduction 641
2. Combinatorial and algebraic setup 648
3. Homology of weighted no-$(k+1)$-equal spaces 655
4. Decomposing cell(n, w) into layers 660
5. Betti number growth function 668
6. Cohomology ring 670
7. Persistent homology 677
8. Relations in the twisted algebra and FI_d–modules 680
9. Configuration spaces of unordered disks 686
10. Open questions and further directions 693
Appendix. Computer calculations for small n 694
References 697

1 Introduction

The configuration space of n labeled unit-diameter disks in an infinite strip of width w is denoted by $\mathrm{config}(n, w)$; Figure 1 depicts an example configuration. Specifically, parametrizing the configurations in terms of the centers of the disks, $\mathrm{config}(n, w)$ is the set of points $(x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}^{2n}$ such that $(x_i - x_j)^2 + (y_i - y_j)^2 \geq 1$ for all i and j, and such that $\frac{1}{2} \leq y_i \leq w - \frac{1}{2}$ for all i. We would like to describe the topology of $\mathrm{config}(n, w)$.

© 2024 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via Subscribe to Open.

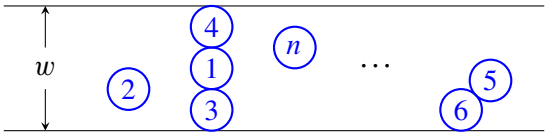

Figure 1: The configuration space $\mathrm{config}(n, w)$ is the set of ways to arrange n disjoint labeled disks of width 1 in $\mathbb{R} \times [0, w]$.

The topology of the configuration space of n points in the plane has been well understood since the work of Arnold [1969] and F Cohen [1976]; see [Sinha 2013] for an overview. Recently, there has been interest in more "physical" models in which the points have thickness and are constrained to lie in a bounded region, drawing inspiration from both statistical physics (as explained by Diaconis [2009]) and robotics (as explained by Farber [2008]). In the first, one imagines molecules of a substance as hard balls and extracts information about states of matter from the way they move past each other. In the latter, one can imagine a number of robots coordinating their movements so that they can travel to different points in a constrained region without bumping against each other; this amounts to a motion planning problem in a disk configuration space.

The topology of disk configuration spaces was first studied mathematically by Baryshnikov, Bubenik and Kahle [Baryshnikov et al. 2014] and experimentally by Carlsson et al. [2012]. While these papers represent real progress, trying to fully understand even the connected components of these spaces seems daunting; see [Kahle 2012]. Further work has taken two approaches to simplifying the question. One is to replace the disks by polygons, such as squares or hexagons, as in [Alpert 2020; Alpert et al. 2023]. The topology of the resulting configuration spaces is closely related to that of disk configuration spaces; in particular, it captures any of their topology that "survives for a long time" as disks grow or shrink. This observation can be formalized using persistent homology.

While it is simpler in some respects than that of disk configuration spaces, the topology of polygon configuration spaces still seems very difficult to understand. A more radical simplification, it turns out, is to remove the side walls of the rectangle, replacing it with the infinite strip, an idea introduced by Alpert, Kahle and MacPherson [Alpert et al. 2021]. That paper defines the spaces $\mathrm{config}(n, w)$ and computes the asymptotic growth of the rank of $H_j(\mathrm{config}(n, w))$ as n increases, up to a constant factor depending on j and w. This turns out to be exponential unless the strip is wide compared to j.

In this paper, we present a number of results about the topology of $\mathrm{config}(n, w)$ and related spaces. The majority of the paper seeks to understand $H_*(\mathrm{config}(n, w); \mathbb{Z})$, and its dependence on n and w, in greater resolution and from a more algebraic perspective. For fixed n and w, we find a geometrically motivated basis for this homology. All the classes in this basis can be assembled out of a small number of classes involving a small number of disks, depending only on w; we make this idea precise using some algebraic machinery due to Sam and Snowden [2017]. We also track the appearance and disappearance of homology classes as the width of the strip changes, using the machinery of persistent homology.

The paper includes several additional results proven using similar methods. We give a basis for the cohomology of $\mathrm{config}(n, w)$ and make progress in understanding its cup product. We explore the possibility of FI_d–module structures on the homology of hard disk configuration spaces and no-$(k+1)$-equal spaces. Finally, we discuss the topology of the configuration space of unordered disks in a strip.

1.1 Main results

We now discuss our main results in greater detail, starting with a description of the homology of $\mathrm{config}(n, w)$ for fixed w and all n.

Theorem A *For fixed w, $H_*(\mathrm{config}(-, w); \mathbb{Z})$ forms a finitely generated, noncommutative twisted algebra whose generators live in $H_{\leq 3w/2-2}(\mathrm{config}(\leq \frac{3}{2}w, w); \mathbb{Z})$.*

This contrasts with the classical family of configuration spaces of points in the plane, which forms a commutative but infinitely generated twisted algebra.

Informally, Theorem A means that $H_j(\mathrm{config}(n, w); \mathbb{Z})$ is spanned by cycles built as follows:

(1) Separate the n disks into groups of at most $\frac{3}{2}w$.

(2) Place the groups in some order along the strip.

(3) Label the disks in some way using the numbers 1 through n.

(4) Let each group do its own thing, without interacting with the others.

The things a group can do — *elementary cycles* — come in two types: 1 to w disks can form a *wheel*, and $w+1$ to $\frac{3}{2}w$ disks can form a *filter*. If z_1 and z_2 are elementary cycles, we refer to the act of placing them next to each other as the *concatenation product*, denoted by $z_1 \mid z_2$.

A wheel of k disks has $k-1$ circular degrees of freedom generated by the rotation of concentric disks, making for a cycle represented by a T^{k-1}.

A filter consists of $r \geq 3$ wheels with $k > w$ disks total such that each wheel is made of at least $k-w$ disks; the wheels are ordered. The filter can move as follows. Every wheel can perform its rotations

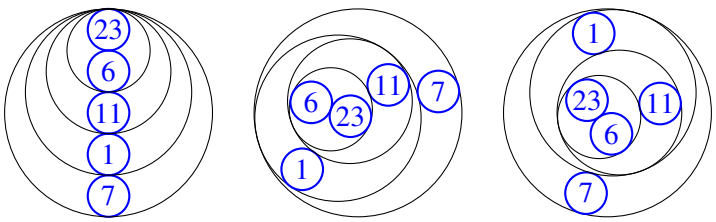

Figure 2: Some configurations of a wheel with five disks. The first configuration gives a canonical (up to switching the first two) ordering of the disks.

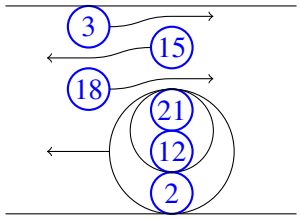

Figure 3: The wheels in a filter always cross over and under each other in the same order. This figure shows a filter with three wheels of size 1 and one wheel of size 3. The resulting cycle is an $S^2 \times T^2$.

independently, for a total of $k-r$ degrees of freedom; each wheel can also move back and forth along the strip, crossing over each other in order. Any $r-1$ wheels can have the same x–coordinate (since they contain at most w disks total) but all r cannot. This creates an S^{r-2} inside the configuration space; thus, the whole $(k-2)$–cycle is represented by an $S^{r-2} \times T^{k-r}$.

For some purposes, it makes sense to consider filters with $r = 2$: such a filter consists of two wheels b_1 and b_2 that don't commute, and can be written as $b_2 \mid b_1 - b_1 \mid b_2$.

Our next theorem gives a basis from among the cycles generated in this way.

Theorem B *$H_*(\mathrm{config}(n,w);\mathbb{Z})$ is free abelian and has a basis consisting of concatenations of wheels and filters with $r \geq 2$. We say one wheel* **ranks above** *another if it has more disks, or has the same number of disks and its largest disk label is greater. A cycle is in the basis if and only if:*

(i) *Each wheel is ordered so that the largest label comes first.*

(ii) *The wheels inside each filter are in ascending order by largest label (regardless of the number of disks).*

(iii) *Adjacent wheels not inside a filter are ordered from higher to lower rank.*

(iv) *Every wheel immediately to the left of a filter ranks above the least wheel in the filter.*

This combinatorial structure admits a natural interpretation via homotopical algebra. Notice that $\mathrm{config}(n,w)$ naturally embeds in $\mathrm{config}(n,w+1)$, forming a filtration. The union of this filtration is the classical configuration space $\mathrm{config}(n)$ of n points in the plane, which turns out to be homotopy equivalent to $\mathrm{config}(n,n)$. The algebraic structure of $H_*(\mathrm{config}(-))$ was considered by Arnold [1969] and Cohen [1976], who showed (in our terms) that it is a *commutative* but infinitely generated twisted algebra whose generators are wheels of all degrees.

For any two wheels in $\mathrm{config}(n)$, we have a choice of commuting them "over" or "under". However, if we order all the wheels (for example, in ascending order by largest label) and make them cross over each other in order, then they commute in a *homotopy coherent* way: informally, this means that concatenated cycles can be permuted in any sequence, or all at the same time, and that all such paths in the space of

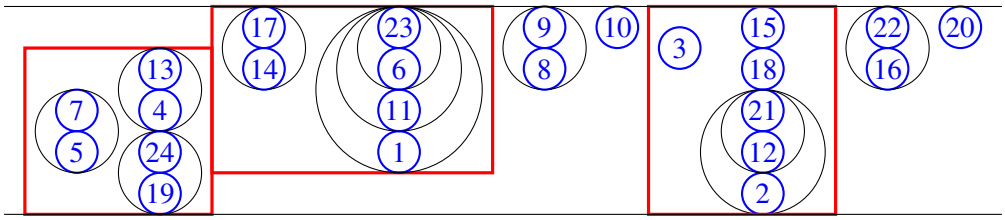

Figure 4: A basic 14–cycle in config$(24,5)$ represented by an $S^1 \times S^0 \times S^2 \times T^{12}$: the three red boxes are filters giving an S^1, an S^0 and an S^2 respectively, and there are 11 additional circular degrees of freedom from spinning the wheels (in black). We also remark: (a) If the disks 10 and 3 switched places, this would no longer represent a basic cycle, by property (iv) in Theorem B. (b) The single disk 10 can move freely past all the disks to its left. So, in a basic cycle, it has to appear all the way on the right.

cycles in config(n) are homotopic.[1] Just as filters with two wheels are commutators of wheels, other filters can then be thought of as nontrivial "higher commutators" that obstruct this homotopy coherence.

These higher commutators appear and disappear as we move up the filtration, increasing w, while wheels of size k are born when $w = k$ and stay forever. This observation can be made precise by considering the filtration's persistent homology.

Theorem C *The basic cycles listed in Theorem B form a $\mathbb{Z}[t]$–basis for the module* $\mathrm{PH}_*(\mathrm{config}(n,*);\mathbb{Z})$. *Every bar born at time w is either infinite or dies by time $2w$.*

1.2 Proof ideas

We analyze config(n,w) by relating it to a class of simpler spaces. The *no-$(w+1)$-equal space* of n points in $\mathbb{R}$, which we denote by $\mathrm{no}_{w+1}(n,\mathbb{R})$, is the subspace of $\mathbb{R}^n$ in which at most w of the coordinates are the same. The topology of this space is fairly easy to understand, although it is related to more complicated questions about hyperplane arrangements; see [Björner and Welker 1995]. The space config(n,w) projects onto $\mathrm{no}_{w+1}(n,\mathbb{R})$ by forgetting the y–coordinates of the disks. Conversely, for every ordering on the numbered disks, there is an injective map from the no-$(w+1)$-equal space to the configuration space of the strip in which the disks, when they meet, go around each other "in order" from top to bottom. Each subspace generated in this way is actually a retract of config(n,w).

However, in many places, two such injections coincide. For example, if we transpose two neighboring disks a and b, then the two injections coincide everywhere except for an open neighborhood of the codimension-1 subspace of the no-$(w+1)$-equal space where a and b coincide. Abstractly, we can think of this subspace as a "weighted" no-$(w+1)$-equal space with $n-1$ symbols, of which one has weight 2. Write $\mathrm{no}_{w+1}(n-1,\mathcal{W})$ for this space; $\mathcal{W}$ represents the set of weights of different points.

[1] Unfortunately, these choices cannot be made equivariantly with respect to relabeling, which underlies the seemingly unavoidable nonequivariance of many of our results.

In the example, the subspace has codimension 1, so its neighborhood looks like $\mathrm{no}_{w+1}(n-1,\mathcal{W})\times(0,1)$; we can decompose the union of the images of the two injections into a copy of $\mathrm{no}_{w+1}(n-1,\mathcal{W})\times[0,1]$ glued onto $\mathrm{no}_{w+1}(n)$. We will see that we can model $\mathrm{config}(n,w)$ by breaking it up into layers that similarly look like thickened weighted no-$(w+1)$-equal spaces.

To compute the homology of $\mathrm{config}(n,w)$, we write it as a direct sum of homology groups of weighted no-$(w+1)$-equal spaces. We use combinatorics (specifically, discrete Morse theory) to compute the homology of these spaces.

1.3 Additional results

Besides Theorems A, B and C, the paper includes a number of other results about $\mathrm{config}(n,w)$ and related spaces.

Counting the basis elements Theorem B gives us a way to compute formulas for the Betti numbers of $\mathrm{config}(n,w)$. We describe a finite computation for each j and w that gives a formula for the rank of $H_j(\mathrm{config}(n,w);\mathbb{Z})$ as a function of n, and we show that this function is a sum of products of polynomial and exponential functions.

The cohomology of config(n,w) We can represent cohomology classes in $H^j(\mathrm{config}(n,w);\mathbb{Z})$ via Poincaré–Lefschetz duality as $(2n-j)$–dimensional compact submanifolds of the configuration space. We give a basis for the cohomology, showing that it is a basis by exploiting the pairing with homology. The main goal of the section is to gain some understanding of the cup product in the cohomology ring, which is fairly complicated and has many indecomposable elements — in contrast with $H^*(\mathrm{config}(n);\mathbb{Z})$, which is generated as a ring by one-dimensional classes whose pairing with homology measures the winding of two points around each other; see [Sinha 2013]. All higher-dimensional classes are linear combinations of cup products of these.

On the other hand, in $\mathrm{config}(n,w)$, pairing with cup products often cannot distinguish between $h\mid h'$ and $h'\mid h$, where h and h' are homology classes. In such cases, a cohomology class which pairs nontrivially with the commutator is perforce indecomposable. This observation allows us to prove Theorem 6.4:

Theorem *The ring $H^*(\mathrm{config}(n,w);\mathbb{Z})$ has indecomposables*

(a) *only in degree 1 when $w=2$;*

(b) *in every degree between 1 and $\lfloor\frac{1}{2}n\rfloor$ and no others when $w=3$;*

(c) *in degree 1 and in every degree between $w-1$ and $\lfloor\frac{1}{2}(n+w-3)\rfloor$ and no degree greater than $n-\lceil n/(w-1)\rceil$ when $w\geq 4$.*

Other algebraic structures One possible operation which takes j–cycles in $\mathrm{config}(n,w)$ to j–cycles in $\mathrm{config}(n+1,w)$ is inserting a singleton disk. If this operation were well defined, then $\mathrm{config}(-,w)$ would be an FI–module, one of the central objects of study in representation stability. However, since some cycles do not commute with singletons, there may be several nonequivalent ways of inserting a singleton,

separated by "barriers". In some cases, if these several ways are in turn well defined, this can be formalized via an FI_d–module structure. We show that $H_j(\mathrm{no}_{k+1}(-,\mathbb{R});\mathbb{Z})$ and $H_j(\mathrm{config}(-,2);\mathbb{Z})$ have natural FI_d–module structures for appropriate d. On the other hand, for $w=3$ we give an example which suggests that the notion of a "barrier" is not well defined and therefore there is no natural FI_d–module structure on $H_j(\mathrm{config}(-,w);\mathbb{Z})$ for $w\geq 3$.

Another strategy to pin down explicitly the algebraic structure of $H_*(\mathrm{config}(-,w);\mathbb{Z})$ is to write down a presentation for the twisted algebra using generators and relations. Category theory dictates that generators and relators in this situation are not just elements but S_n–representations for various n. We write down the generators and relations for $H_*(\mathrm{config}(-,2);\mathbb{Z})$, but already there is some extra difficulty because of the failure of Maschke's theorem integrally. A potential avenue for future research is to write down relations for $H_*(\mathrm{config}(-,w);\mathbb{Q})$ and explore the implications for the multiplicity of various S_n–representations in $H_j(\mathrm{config}(n,w);\mathbb{Q})$.

Unordered disks We also discuss the homology of the configuration space of unordered disks, that is, the quotient space $\mathrm{config}(n,w)/S_n$. Taking this quotient produces a large amount of torsion in the homology, so, instead of trying to write down all the torsion, we compute homology with coefficients in $\mathbb{F}_p$ and $\mathbb{Q}$. (In the bulk of the paper, we use coefficients in $\mathbb{Z}$.) The concatenation product can still be defined in this case and gives $H_*(\mathrm{config}(-,w)/S_-)$ the structure of a bigraded algebra over the base field. We compute a basis for the homology and give generators and relations for this algebra.

Structure of the paper

Section 2 contains preliminaries, including descriptions of the cell complexes we study and the algebraic framework we use in Theorem A. Sections 3 and 4 prove Theorem B, with Theorem A as a consequence; Section 3 addresses the homology of weighted no-$(k+1)$-equal spaces, and Section 4 proves its relationship to homology of configuration spaces of disks in a strip. This is the core technical content of the paper.

The rest of the sections are largely independent of each other and may appeal to different audiences (eg Section 6 to those interested in motion planning, and Section 8 to experts in representation stability). Section 5 concerns finding a formula for the Betti numbers, by counting the basis elements from Theorem B. Section 6 describes aspects of the cup product structure of the cohomology of our configuration spaces. Section 7 proves Theorem C about how the homology changes with w, the width of the strip. Section 8 concerns additional algebraic properties, in particular the question of whether our homology groups form FI_d–modules. Section 9 characterizes the homology with field coefficients in the case of unordered disks. Finally, Section 10 lists some questions for further study.

Acknowledgements

Alpert and Manin were supported by the National Science Foundation under awards DMS-1802914 and DMS-2001042, respectively. We thank Matt Kahle, Jenny Wilson and Nick Wawrykow for many

helpful conversations about this material, as well as the referee for comments that pushed us to make many sections easier to follow. Manin also thanks Shmuel Weinberger for the suggestion to look at cup products.

2 Combinatorial and algebraic setup

For the purpose of computation, Alpert et al. [2021] replace the configuration space $\mathrm{config}(n, w)$ by a homotopy-equivalent cell complex $\mathrm{cell}(n, w)$. We use the same complex; we also use the same method to find a cell complex $P(n, \mathcal{W}, k)$ that is homotopy equivalent to the weighted no-$(k+1)$-equal space $\mathrm{no}_{k+1}(n, \mathcal{W})$. In the case of ordinary, unweighted no-$(k+1)$-equal spaces this recovers a result of [Björner 2008].

In this section, we define the complexes $\mathrm{cell}(n, w)$ and $P(n, \mathcal{W}, k)$ and describe various algebraic structures on their cells which induce a similar structure on their homology. In particular, we introduce twisted algebras and show that, for each w, $H_*(\mathrm{config}(-, w))$ and $H_*(\mathrm{no}_{w+1}(-, \mathbb{R}))$ are examples. In the remainder of the paper, unless otherwise specified, homology and cohomology are always computed with coefficients in $\mathbb{Z}$.

2.1 The complex cell(n)

The cell complex $\mathrm{cell}(n, w)$ is defined as a subcomplex of a cell complex $\mathrm{cell}(n)$ described in [Blagojević and Ziegler 2014]. The complex $\mathrm{cell}(n)$ is defined in terms of the permutohedron $P(n)$, which is the $(n-1)$–dimensional polytope in $\mathbb{R}^n$ equal to the convex hull of the $n!$ points with coordinates $1, 2, \dots, n$ in some order. The faces of $P(n)$ can be labeled by partitions of $\{1, 2, \dots, n\}$, where the sets of the partition are ordered but the elements of each set are unordered. We refer to each set of the partition as a *block*. For instance, each vertex is labeled by a sequence of n singleton blocks, and the top-dimensional face is labeled by the one block $\{1, 2, \dots, n\}$. Given any permutation $\sigma \in S_n$, thought of as an ordering on $\{1, 2, \dots, n\}$, we can order the elements of each block according to σ. Then, to write out the label of a given face, we can write out the elements of each block in order, with vertical bars between blocks. We refer to a label of this form as a *symbol*. For example, the following is a symbol with four blocks that could come from the ordering $4 \prec 5 \prec 6 \prec 7 \prec 8 \prec 1 \prec 2 \prec 3$:

$$(7\,2 \mid 6 \mid 4\,5\,1 \mid 8\,3).$$

To define $\mathrm{cell}(n)$, we start with $n!$ copies of $P(n)$, one for each $\sigma \in S_n$. For the copy of $P(n)$ associated with σ, we use σ to label each face of that copy of $P(n)$ by a symbol. Then, whenever faces from multiple different copies of $P(n)$ have the same symbol, we identify those faces. For instance, all copies of $P(n)$ have the same vertices, but each of them has its own distinct top-dimensional face. In this way, $\mathrm{cell}(n)$ has exactly one cell for every possible symbol on $\{1, 2, \dots, n\}$.

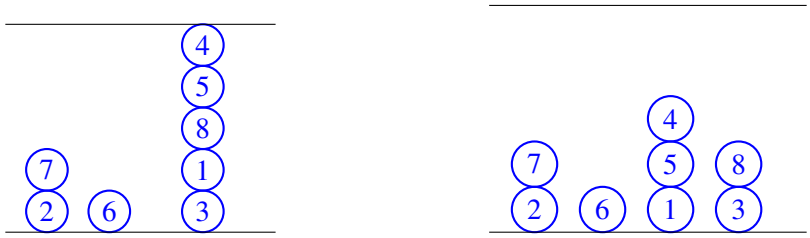

Figure 5: We can imagine each symbol of $\mathrm{cell}(n, w)$ as a configuration in $\mathrm{config}(n, w)$, where the numbers in each block are the labels in a column of disks. Pictured are configurations representing the symbol $(7\,2 \mid 6 \mid 4\,5\,8\,1\,3)$ and its face $(7\,2 \mid 6 \mid 4\,5\,1 \mid 8\,3)$.

The incidence relation on cells of $\mathrm{cell}(n)$ can be deduced from the geometry, or can be described explicitly on symbols as follows. A cell f in $\mathrm{cell}(n)$ is a face of the boundary of a cell g if g can be obtained from f by deleting a bar and *shuffling* the entries of the two neighboring blocks, preserving the ordering of the entries in each block. For example, one shuffle of $4\,6\,1 \mid 7\,3\,2$ would be $7\,4\,6\,3\,1\,2$. A d–dimensional cell consists of $n-d$ blocks.

Informally, we think of the elements of each block of a given symbol as the labels of disks in a vertical stack in $\mathrm{config}(n, w)$, as in Figure 5. Accordingly, as a way to specify that no more than w disks should be in each vertical stack, we define $\mathrm{cell}(n, w)$ to be the subcomplex of $\mathrm{cell}(n)$ consisting of all cells for which every block has at most w elements.

To describe the relationship between $\mathrm{cell}(n, w)$ and $\mathrm{config}(n, w)$, we start by defining $\mathrm{config}(n, w)$ more precisely as the set of configurations of n ordered open disks of diameter $1/w$ in the strip $\mathbb{R} \times (0, 1)$. This is a subspace of the set of configurations of n ordered distinct points in $\mathbb{R} \times (0, 1)$ such that no more than w points are on any vertical line, and [Alpert et al. 2021, Theorem 3.3] constructs a deformation retraction between the two spaces. Thus, we abuse notation and use $\mathrm{config}(n, w)$ to mean configurations of points with no more than w vertically aligned. We use $\mathrm{config}(n)$ to mean configurations of points either in $\mathbb{R} \times (0, 1)$ or in $\mathbb{R}^2$, not distinguishing between these homotopy equivalent spaces.

Theorem 2.1 *There is an affine embedding of the barycentric subdivision of* $\mathrm{cell}(n)$ *into* $\mathrm{config}(n)$ *such that, for each* w*, the restriction to* $\mathrm{cell}(n, w)$ *maps into* $\mathrm{config}(n, w)$ *and is a homotopy equivalence.*

Proof We use the following description of $\mathrm{config}(n)$ in coordinates:

$$\mathrm{config}(n) = \{(x_1, \dots, x_n, y_1, \dots, y_n) \in \mathbb{R}^n \times (0, 1)^n \mid (x_i, y_i) \neq (x_j, y_j) \text{ if } i \neq j\}.$$

To define the map on a point p in $\mathrm{cell}(n)$, we need to specify x–coordinates and y–coordinates, as well as check the condition $(x_i, y_i) \neq (x_j, y_j)$.

Let p be an arbitrary point in $\mathrm{cell}(n)$. It is in at least one of the permutohedra $P(n)$ that constitute $\mathrm{cell}(n)$, and $P(n)$ is embedded in $\mathbb{R}^n$, so in this sense p has coordinates in $\mathbb{R}^n$. We set the x–coordinates of the image of p to be these coordinates of p in $\mathbb{R}^n$.

For the y–coordinates, we start with the case where p is the barycenter of a cell in cell(n). To set each y_i, we find the location of the number i in the symbol of the cell of p. If i appears as the lth element of a block of size k, we set y_i to be $(k-l+1)/(n+1)$. In other words, for each block of size k we set the y–coordinates of the points in the block, in order, to be $k/(n+1), (k-1)/(n+1), \dots, 1/(n+1)$. To assign y–coordinates when p is not a barycenter, we extend the map so that its restriction to each simplex of the barycentric subdivision of cell(n) is affine.

Note that, in any given cell, away from the barycenter, the y–coordinates of the points in each block remain in order. This is because, in the closure of our cell, blocks may merge but may not separate, and, when merging, the elements from each smaller block remain in order. This implies that our map is injective, because different points with the same x–coordinates are distinguished by their y–coordinates, which in each block appear in the same order as in the corresponding symbol from cell(n).

To check that the resulting map lands in config(n), suppose that we have a point p for which $x_i = x_j$. The permutohedron coordinates imply that in the symbol of the cell of p, the numbers i and j are in the same block, and so the note in the previous paragraph implies that $y_i \neq y_j$. Thus, the image of p is in config(n).

To show that the map is a homotopy equivalence, for each symbol α from cell(n), as in [Alpert et al. 2021] we define the corresponding open set $U_\alpha \subseteq$ config(n) to be the set of points $(x_1, \dots, x_n, y_1, \dots, y_n)$ in $\mathbb{R}^n \times (0,1)^n$ such that:

- Whenever i appears before j in the same block, we have $y_i > y_j$.
- Whenever i appears before j in different blocks, we have $x_i < x_j$.
- If k and l are in the same block, and k' and l' are in different blocks (with k' and l' not necessarily distinct from k and l), then we have

$$|x_k - x_l| < |x_{k'} - x_{l'}|.$$

We claim that the union of U_α where α ranges over the symbols from cell(n,w) is config(n,w). If α is a symbol of cell(n,w), then U_α is contained in config(n,w) because points from different blocks have different x–coordinates, so no more than w points can be vertically aligned. For the reverse inclusion, every element of config(n,w) is in some U_α, because we can construct α by taking the blocks to be the sets of points with the same y–coordinate, ordering the blocks from left to right, and ordering the elements of each block from top to bottom.

Each U_α is an open convex set, and [Alpert et al. 2021, Theorem 3.4] proves that the nerve of this open cover of config(n,w) is the barycentric subdivision of cell(n,w). Thus, because our map sends the barycenter of each cell α into a point of the corresponding open set U_α, our map is a homotopy equivalence for each w. □

2.2 Signs of the boundary operator

In order to study the integral cellular chain complex of $\mathrm{cell}(n, w)$, we need to specify orientations on cells and signs for the boundary operator. To describe these, we first generalize slightly. Observe that the entries of a symbol don't have to be the numbers 1 through n, but can be any n–element set. So, for any finite set A, we have a complex $\mathrm{cell}(A)$. (Similarly, we will sometimes write $P(A)$ for the permutohedron whose coordinates are indexed by elements of A.) Moreover, there is a cellular map

$$|\colon \mathrm{cell}(A) \times \mathrm{cell}(B) \to \mathrm{cell}(A \sqcup B)$$

which takes a pair of symbols to the symbol obtained by putting them next to each other with a bar in between. We call this the *concatenation product*.

Any cell in $\mathrm{cell}(n)$ is either top-dimensional or a concatenation product. We define the boundary operator on a top-dimensional cell g by taking the coefficient of a cell $f = a \mid b$ in ∂g to be

$$(-1)^{\mathrm{length}(a)} \cdot \mathrm{sign}(\text{permutation } g \mapsto ab).$$

On a cell $g_1 \mid g_2$, we define ∂ via a Leibniz rule:

$$\partial(g_1 \mid g_2) = \partial g_1 \mid g_2 + (-1)^{\dim(g_1)} g_1 \mid \partial g_2. \tag{2.2}$$

This defines an injective chain complex homomorphism on cellular chains,

$$|\colon C_*(\mathrm{cell}(A)) \otimes C_*(\mathrm{cell}(B)) \to C_*(\mathrm{cell}(A \sqcup B)),$$

using the standard tensor product on chain complexes, where the differential is defined by

$$\partial(a \otimes b) = \partial a \otimes b + (-1)^{\deg a} a \otimes \partial b.$$

Proposition 2.3 (a) *The boundary operator defined above satisfies* $\partial^2 = 0$.

(b) *The* S_n*–action on* $\mathrm{cell}(n)$ *induced by permutations of* $[n]$ *preserves orientations of cells.*

These two features will allow us to define a twisted algebra structure on $H_*(\mathrm{cell}(-))$.

Proof Let g be a cell in $\mathrm{cell}(n)$.

First, suppose that g is top-dimensional, and let e be a codimension-2 face of g. Then $e = e_1 \mid e_2 \mid e_3$, where e_1, e_2 and e_3 are blocks. There are two intermediate faces between e and g, which we denote by $f = b \mid e_3$ and $f' = e_1 \mid b'$. We compare

$$\mathrm{sign}(e \text{ in } \partial f) \cdot \mathrm{sign}(f \text{ in } \partial g) \quad \text{and} \quad \mathrm{sign}(e \text{ in } \partial f') \cdot \mathrm{sign}(f' \text{ in } \partial g),$$

and we show that these two products are opposite signs. This will show that $\partial^2 g = 0$.

If we consider just the contribution from the signs of the permutations, both products give the sign of the permutation relating e and g, so those contributions are equal. For the contribution from the Leibniz rule,

only the incidence between e and f' involves splitting a block that is not the first, so that incidence has a sign contribution of

$$(-1)^{\dim(e_1)} = (-1)^{\mathrm{length}(e_1)-1}$$

from the Leibniz rule, and all the other incidences have a sign contribution of 1 from the Leibniz rule. Finally, the contribution from the length of the first block gives

$$(-1)^{\mathrm{length}(b)+\mathrm{length}(e_1)} = (-1)^{\mathrm{length}(e_2)}$$

for the path through f and $(-1)^{\mathrm{length}(e_1)+\mathrm{length}(e_2)}$ for the path through f'. Taking the product of all of these, we see that the two paths give opposite signs.

If g is not top-dimensional, then g is a concatenation product $g_1 \mid g_2$, in which case we use a standard argument. The Leibniz rule gives

$$\partial^2(g_1 \mid g_2) = \partial^2 g_1 \mid g_2 + [(-1)^{\dim(g_1)} + (-1)^{\dim(\partial g_1)}] \cdot \partial g_1 \mid \partial g_2 + g_1 \mid \partial^2 g_2,$$

which is zero by induction on the number of blocks of g. This proves (a).

For (b), notice that the definitions of the signs of the boundary operator do not use any particular ordering on the numbers 1 through n. Therefore, they are invariant with respect to permutations. Since the S_n–action preserves signs of 0–cells, it preserves signs of all cells. □

2.3 Twisted algebra structure

Let FB be the category of finite sets and bijective maps. Then cell$(-)$ is a functor from FB to the category of cell complexes and cellwise maps (an FB–*complex*, for short). (This is just a categorical way of saying that there is a cellular S_n–action on cell(n).) In particular, the cellular chains $C_*(\mathrm{cell}(-))$ form an FB–chain complex. Moreover, we can define a tensor product (the *Day convolution*) on FB–objects in a monoidal category by

$$(F \otimes G)(S) = \bigoplus_{A \sqcup B = S} F(A) \otimes F(B).$$

A unital monoid object with respect to this tensor product is called a *twisted algebra* (see [Sam and Snowden 2012] for a detailed discussion of twisted *commutative* algebras). In plain English, a twisted algebra is a family of objects A_n for $n = 0, 1, \dots$ (eg abelian groups) equipped with the following structure:

- Each A_n is equipped with an S_n–action.
- For every partition of $\{1, \dots, n\}$ into subsets of size i and j, there is a "multiplication" $A_i \otimes A_j \to A_n$. For different partitions, these multiplications commute with the S_n–action on A_n.
- There is a unit in A_0 such that multiplying by it induces the identity map on A_n.

The observations about the differential in Section 2.2 show that the concatenation product

$$|: C_i(\mathrm{cell}(A)) \otimes C_j(\mathrm{cell}(B)) \to C_{i+j}(\mathrm{cell}(A \sqcup B))$$

makes $C_*(\mathrm{cell}(-))$ into a (noncommutative) differential graded twisted algebra, or *dgta*, whose unit is the unique 0–cell $()$ in $C_*(\mathrm{cell}(\varnothing))$. (By definition, a dgta is simply a twisted algebra in the category of chain complexes. The fact that the multiplication maps $A_i \otimes A_j \to A_n$ are chain maps forces the differential to satisfy a Leibniz rule such as (2.2).) The homology of a dgta naturally forms a graded twisted algebra. The graded twisted algebra $H_*(\mathrm{cell}(-))$ is well understood since the work of Cohen in the 1970s and is in fact commutative; see eg [Miller and Wilson 2019, Theorem 3.4].

We are most interested in the subcomplex $\mathrm{cell}(n, w)$ consisting of cells whose blocks each have size at most w. Since the concatenation product of two such cells again has the same property, $C_*(\mathrm{cell}(-, w))$ is a sub-dgta of $C_*(\mathrm{cell}(-))$. Our goal in this paper is to understand the graded twisted algebra $H_*(\mathrm{cell}(-, w))$, and in particular to show that it is finitely generated.

2.4 The permutohedron and no-$(k+1)$-equal spaces

The difference between $\mathrm{cell}(n)$ and $P(n)$ is that in $\mathrm{cell}(n)$ the numbers within a block are ordered and in $P(n)$ they are not. This gives a natural projection $\mathrm{cell}(n) \to P(n)$ (forget the ordering of entries inside each block) and, for every global ordering of $1, \dots, n$, an injective map $P(n) \to \mathrm{cell}(n)$ (arrange the entries in each block in the given order).

$P(n)$ is a polytope, and is in particular contractible. As with $\mathrm{cell}(n)$, we can filter $P(n)$ by the largest size of a block, producing a sequence of complexes $P(n, k)$. These are homotopy equivalent to the no-$(k+1)$-equal space of n points in $\mathbb{R}$, as pointed out by Björner [2008, Theorem 2.4]; their homology was computed first by Björner and Welker [1995]. As with $\mathrm{cell}(n, w)$, $C_*(P(-, k))$ naturally has a dgta structure which induces a graded twisted algebra structure on $H_*(P(-, k))$. From the results of [Björner and Welker 1995], one sees that this is finitely generated, and in fact just has two generators: one in degree 0 (a point) and one in degree $k-1$ (the boundary of a $P(k+1)$). We recover this, together with a set of relations, in Section 8.

However, we are interested in a somewhat more complicated structure. Let FBW be the category of *weighted* finite sets and weight-preserving bijections. That is, every element is associated with a natural number, which is its weight. Then, given a weighted set $(A, \mathcal{W} \in \mathbb{N}^A)$, there is a complex $P(A, \mathcal{W}, k)$ which consists of all the cells for which the sum total weight of every block is at most k. The following generalization of [Björner 2008, Theorem 2.4] follows by the same argument:

Theorem 2.4 *The complex $P(A, \mathcal{W}, k)$ is homotopy equivalent to the* ***weighted no-$(k+1)$-equal space*** $\mathrm{no}_{k+1}(A, \mathcal{W})$ *of $|A|$ points in $\mathbb{R}$ with weights $\mathcal{W}$, that is, the space of configurations of $|A|$ points in $\mathbb{R}$ such that no set of coincident points has total weight greater than k.*

Remark 2.5 The functor $C_*(P(-,k))$ is an FBW–chain complex, and in fact an FBW–dga. That is, we can define a tensor product on FBW–objects in a monoidal category by

$$(F \otimes G)(S, \mathcal{W}) = \bigoplus_{(A,\mathcal{W}_A)\sqcup(B,\mathcal{W}_B)=(S,\mathcal{W})} F(A, \mathcal{W}_A) \otimes F(B, \mathcal{W}_B),$$

and the concatenation product on $C_*(P(-,k))$ then makes it into a monoid object. This in turn makes $H_*(P(-,k))$ into a graded FBW–algebra. Once one makes all this precise, Theorem 3.4 can be interpreted as showing that this algebra is finitely generated for every k, analogously to our results about $\mathrm{cell}(-,w)$.

2.5 Generators and relations for twisted algebras

The above discussion deduces the following facts:

Theorem 2.6 *The sequences of graded abelian groups $H_*(\mathrm{cell}(-))$, $H_*(\mathrm{cell}(-,w))$ and $H_*(P(-,k))$ admit the structure of graded twisted algebras.*

To demonstrate Theorem A, we will show that wheels and filters form a finite generating set for $H_*(\mathrm{cell}(-,w))$. Later we will also give presentations of $H_*(\mathrm{cell}(-,2))$ and $H_*(P(-,k))$ by generators and relations. To make this precise, we define a *free twisted algebra* functor F_τ from FB–modules to twisted algebras as the left adjoint to the forgetful functor U_τ from twisted algebras to FB–modules; such a functor always exists for monoids in a reasonable monoidal category [Mac Lane 1998, VII.3, Theorem 2]. Informally, a basis element for a free twisted algebra on a set $\{V_i\}$ of representations of various S_{n_i} is specified by a list of basis vectors of the various V_i, each labeled by a set of n_i labels; one takes products by concatenating these lists and retaining the labels.

If A is a twisted algebra, we say an FB–submodule $G \subset A$ *generates* A if the induced morphism $F_\tau G \to A$ is surjective; A is *finitely generated* if it has a finite-dimensional generating module G. A *presentation* of a twisted algebra A by generators and relations formally consists of a coequalizer diagram

$$F_\tau R \underset{0}{\overset{r}{\rightrightarrows}} F_\tau G \to A,$$

where G and R are FB–modules of generators and relations, respectively, and $r\colon F_\tau R \to F_\tau G$ is an FB–module homomorphism [Riehl 2017, Section 5.4]. By the adjunction, it suffices to provide a homomorphism $R \to U_\tau F_\tau G$, ie to describe the relators of A as linear combinations of words in G.

Informally, to prove finite generation, it's enough to provide a finite number of elements whose closure under multiplication and S_n–action is all of A. However, to provide a full description of a presentation, one must also understand the S_n–action on the generators. To show that A is presented by a generating set G with relations R, it is enough to show that A is generated by G and that every product of elements of G can be reduced to a basis element of A via the relators.

3 Homology of weighted no-$(k+1)$-equal spaces

In the previous section, we showed that the space $\mathrm{no}_{k+1}(n, \mathcal{W})$ retracts to a subcomplex $P(n, \mathcal{W}, k)$ of the permutohedron $P(n)$. To compute the homology of $P(n, \mathcal{W}, k)$, we use a discrete gradient vector field on $P(n)$.

3.1 Discrete Morse theory

In any cell complex, the cellular homology comes from a chain complex generated by the cells; very broadly, discrete Morse theory gives a way to decompose the chain complex as a direct sum of a chain complex that has no homology (which we discard) and a chain complex generated by a smaller subset of cells, the critical cells. To compute the homology exactly, we need to

(1) reduce to the smaller chain complex;

(2) show that the differentials in the smaller chain complex are all zero, so that the homology has a $\mathbb{Z}$–basis in bijection with the set of critical cells.

The basic definitions in discrete Morse theory are as follows. In any polyhedral cell complex, we say that cell f is a *face* of cell g if f is in the boundary of g and $\dim f = \dim g - 1$, and we say that g is a *coface* of f if f is a face of g. A *discrete vector field* on a polyhedral cell complex is a set V of pairs of cells $[f, g]$ such that f is a face of g and each cell can be in at most one pair; an example is shown in Figure 6. A discrete vector field V is *gradient* if there are no closed V–walks. A V–walk is a sequence of pairs $[f_1, g_1], \dots, [f_r, g_r]$ with $[f_i, g_i] \in V$ such that each f_{i+1} is a face of g_i other than f_i. The V–walk is closed if $f_r = f_1$.

A cell is *critical* with respect to a discrete gradient vector field V if the cell is not in any pair in V. The fundamental theorem of discrete Morse theory [Forman 2002] states that there is a cell complex that is a

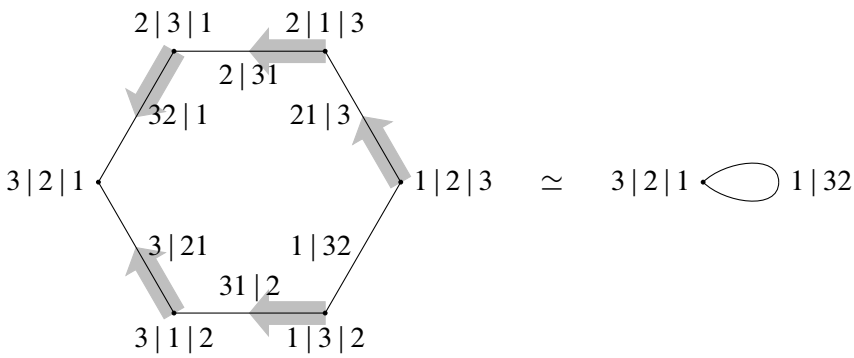

Figure 6: A *discrete gradient vector field* consists of a set of disjoint pairs of cells, each pair incident and of consecutive dimensions. The complex is homotopy equivalent to one in which the paired cells are collapsed, and only the *critical* (unpaired) cells remain.

strong deformation retraction of the original cell complex in which there is one cell per critical cell of V. Thus, we can compute the homology groups $H_j(P(n, \mathcal{W}, k))$ by defining discrete gradient vector fields and computing the homology of the collapsed chain complexes generated by the critical cells.

Our cell complexes $\mathrm{cell}(n, w)$ are not polyhedral cell complexes because, for instance, there are distinct cells with the same boundary. However, discrete Morse theory still gives an isomorphism on homology between the original chain complex and the collapsed chain complex as long as the following property holds: for each pair $[f, g]$ of the discrete gradient vector field, the coefficient of f in the boundary of g is a unit in our coefficient ring, in this case ± 1 because we are using coefficients in $\mathbb{Z}$. In $\mathrm{cell}(n, w)$, every coefficient in every boundary map is ± 1, so this property holds automatically. Alternatively, the barycentric subdivision of $\mathrm{cell}(n, w)$ is a polyhedral cell complex, so our arguments could be adapted to work with the barycentric subdivision instead. Thus, we proceed with the discrete Morse theory as if $\mathrm{cell}(n, w)$ were a polyhedral cell complex.

One way to define a discrete gradient vector field on a cell complex is by defining a total ordering on all the cells. Given a total ordering, the resulting vector field contains a pair $[f, g]$ if and only if both f is the greatest face of g and g is the least coface of f (and, for nonpolyhedral complexes, we also require the coefficient of f in the boundary of g to be a unit). One can prove that this vector field is gradient (see [Bauer 2021, Remark 3.7]).

One advantage of doing the construction in this way is that there is a simple criterion guaranteeing that the discrete gradient vector field forms a perfect Morse function. The term "perfect Morse function" refers to the case where, on the collapsed chain complex generated by the critical cells, the differential is zero, so the critical cells form a basis for the homology of the complex. The following general lemma shows that it suffices to construct, for each critical cell e, a cycle $z(e)$ such that e is its maximum cell and has coefficient equal to a unit in the coefficient ring. It turns out that such cycles automatically represent linearly independent homology classes. In our main theorems we use the lemma with $\mathbb{Z}$ coefficients, and in Section 9 we use it with field coefficients $\mathbb{Q}$ and $\mathbb{F}_p$.

Lemma 3.1 *Let X be any finite cell complex with a total ordering on the cells, giving a discrete gradient vector field. Let R be a ring of coefficients. Suppose there is a collection of cellular cycles with the following properties:*

- *The cycles in our collection are in bijection with the critical cells of the discrete gradient vector field. For each critical cell e, we denote the corresponding cycle by $z(e)$.*
- *Under the total ordering, the greatest of all the cells appearing with nonzero coefficient in the cellular chain $z(e)$ is the cell e.*
- *The coefficient of e in the chain $z(e)$ is a unit in R.*

Then the homology classes of the cycles $z(e)$ form an R–basis for $H_*(X; R)$.

Proof For any pair $[f, g]$ in the discrete vector field, we refer to f as a "match-up cell" and refer to g as a "match-down cell". We also define $z'(f)$ to be the boundary of g; we know that f is the greatest cell appearing in $z'(f)$, and that it has unit coefficient because of how we have defined the discrete vector field from the ordering.

First, we show that every j–cycle z is an R–linear combination of cycles $z(e)$ and $z'(f)$, where e ranges over the critical j–cells and f ranges over the match-up j–cells. This follows from the following observation: if a match-down cell g is the greatest cell in a j–chain, then, in the boundary of that chain, the corresponding match-up cell f appears with nonzero coefficient, because g is the least coface of f, so no other cell in the chain has f as a face. Thus, for any j–cycle z, the greatest cell of z cannot be a match-down cell. It is either a critical cell e or a match-up cell f, so we subtract the appropriate multiple of $z(e)$ or $z'(f)$ to get a new cycle with lesser maximum. Repeating this process gives us z as a linear combination of cycles $z(e)$ and $z'(f)$, so, because each $z'(f)$ is a boundary, this implies that z is homologous to a linear combination of the cycles $z(e)$ only.

To show the uniqueness, we need to show that no nontrivial linear combination of cycles $z(e)$ is null-homologous. Because the cycles $z(e)$ and $z'(f)$ have distinct maxima, they are linearly independent. Thus, it suffices to show that, if a j–cycle z is a boundary, it is a linear combination of the boundaries $z'(f)$. To see this, we look at the set of all $(j+1)$–chains. The chains $z(e)$, $z'(f)$ and g (as e ranges over all critical $(j+1)$–cells, f ranges over all match-up $(j+1)$–cells, and g ranges over all match-up $(j+1)$–cells) form an R–basis for the set of all $(j+1)$–chains, because they have distinct maxima equal to the set of all j–cells. When we apply the boundary map to this basis, the cycles $z(e)$ and $z'(f)$ map to zero, and the match-down cells g map to the j–dimensional boundaries $z'(f)$. Thus, indeed, every j–dimensional boundary is a linear combination of these boundaries $z'(f)$.

Thus, every homology class in $H_*(X)$ can be written as an R–linear combination of the homology classes of the cycles $z(e)$, and the combination is unique. □

3.2 Discrete gradient vector fields on permutohedra

In what follows, we define a total ordering on all of $P(n)$, the polyhedral complex that contains $P(n, \mathcal{W}, k)$ as a subcomplex. We use the resulting discrete gradient vector fields to compute the homology, by analyzing the critical cells and constructing cycles dual to each one.

Recall that the cells of $P(n)$ are in bijection with ordered partitions of $[n]$ into blocks. We say the *weight* of a block is the sum of the weights of its elements. A block is a *singleton* if it only has one element. We assign some blocks to *leader–follower pairs* by walking left to right. A block is a *follower* if it follows a singleton (its *leader*) whose element is smaller than any of the follower's elements and which is not itself a follower. Then a total ordering $\prec$ on cells with symbols f and g is given by looking at the first block at which they differ. Let's call the two blocks f_i and g_i. Then the ordering is given as follows:

(i) If f_i is a follower and g_i is not, then $f \prec g$.

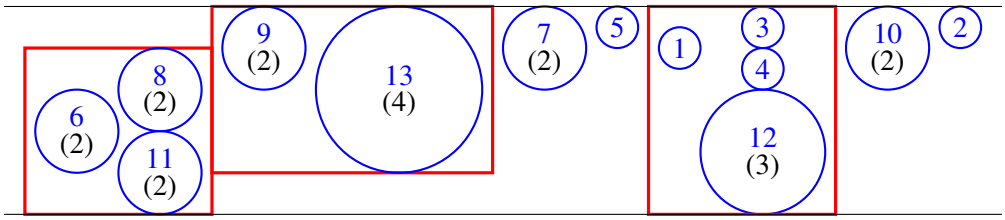

Figure 7: The 3–cycle in $\mathrm{no}_6(13, \mathcal{W})$, with weights indicated by circle size and in parentheses, corresponding to the critical cell $(6 \mid 8\ 11 \mid 9 \mid 13 \mid 7 \mid 5 \mid 1 \mid 3\ 4\ 12 \mid 10 \mid 2)$. The red boxes enclose leader–follower pairs and indicate the boundaries of a $P(3)$, a $P(2)$, and a $P(4)$, respectively; thus, the cycle as a whole is represented by an $S^1 \times S^0 \times S^2$.

(ii) If f_i and g_i are both followers, then they are ordered by number of elements, and then arbitrarily (eg in lexicographic order) if they have the same number of elements.

(iii) If f_i and g_i are not followers, then we order their elements from largest to smallest; they are then ordered lexicographically, but according to a backwards order on the "alphabet". That is,

$$3 \prec 3\,2 \prec 3\,2\,1 \prec 3\,1 \prec 2 \prec 2\,1 \prec 1.$$

Lemma 3.2 *A cell in $P(n, \mathcal{W}, k)$ is critical if and only if every block is either*

(a) *a singleton which is not a follower, or*

(b) *a follower such that its weight and that of its leader add up to at least $k+1$.*

Proof We claim that the matching of the other cells is as follows. We look at the first block where a given cell does not match the characterization of critical given above. There are two possibilities:

- The block is a follower and, together with its leader, has total weight at most k. In that case, the cell matches *up* to the cell in which the follower and its leader are combined.
- The block is not a follower or a singleton; if preceded by a nonfollower singleton, it has at least one element smaller than that singleton. Then the cell matches *down* to the cell where the least element of the block is split off into its own block which comes before the rest. This is now a leader–follower pair, and its leader is not also a follower.

If f matches up to g, then g matches down to f, because the new combined block cannot be a follower (otherwise the leader in f would also be a follower, contradicting the definition of leader). Similarly, if g matches down to f, then f matches up to g. Thus, the match-up and match-down cells pair to form a discrete vector field.

It remains to show that this discrete vector field is induced by our ordering. Let f be a match-down cell and g be the corresponding match-up cell. We need to show f is the greatest face of g and g is the least coface of f.

To show that f is the greatest face of g, consider the result of splitting any earlier block of g. Because g looks like a critical cell at that stage, that block is a follower; after splitting, it is shorter and is still a follower, so it is smaller by property (ii) of the ordering. In contrast, among ways to split the kth block another way, or to split a later block, f is the greatest because it is the only one for which the kth block begins with the least element of the kth block of g (here we apply properties (i) and (iii)).

To show that g is the least coface of f, consider the result of combining any earlier blocks of f. Because f looks like a critical cell at that stage, the two blocks would be nonfollower singletons in decreasing order, so the combined block would be larger, not be a follower, and have the same first element as the first of the two singletons; therefore, the new block is larger by property (iii). In contrast, among ways to combine later blocks of f, g is the least because it is the only one that increases the first element of the kth block (again applying property (iii)). □

3.3 Basis for homology

We now give a basis for the homology of $P(A, \mathcal{W}, k)$ for an n–element set A, thereby demonstrating that it is free abelian. By Lemma 3.1, it is sufficient to exhibit a cycle for every critical cell such that the critical cell is the largest cell in the cycle. The construction implicitly relies on an ordering on A; for it to work, we must fix an ordering so that $\mathcal{W} = (w_1, \dots, w_n)$ satisfies

$$w_1 \le w_2 \le \cdots \le w_n. \tag{3.3}$$

Every critical cell e is a concatenation product of nonfollower singletons and leader–follower pairs $(b_i \mid b_{i+1})$. We now build a corresponding cycle $z(e) \in P(A, \mathcal{W}, k)$. For a cell consisting of one singleton, the cycle will simply be the corresponding 0–cell. Now suppose e consists of a single leader–follower pair $(b_1 \mid b_2)$. Then our cycle $z(e)$ will be the boundary of the top cell $(b_1 b_2)$ of $P(A)$. Since by (3.3) each element of b_2 has greater or equal weight to the singleton b_1, every cell in the boundary is a cell of $P(A, \mathcal{W}, k)$. Moreover, $(b_1 \mid b_2)$ is the highest cell of $z(e)$ according to our total ordering, since b_1 is the highest-ranked block.

For a general critical cell, we define z by requiring that

$$z(e) = z(e_1 \mid e_2) = z(e_1) \mid z(e_2)$$

for any splitting $e = e_1 \mid e_2$ which does not split a leader–follower pair. Here $z(e_i)$ is defined based on the ordering on $A_i \subset A$ inherited from A. To see that e is the highest cell in the resulting cycle, note that the ordering on cells depends on the first block in which two cells differ. For any two cells in $z(e)$, this block will be a leader and the above argument will apply.

It is clear also that all the nonzero coefficients of $z(e)$ are ± 1. We have now proved:

Theorem 3.4 *Suppose that $\mathcal{W} = (w_1, \dots, w_n)$ is a nondecreasing sequence of weights. Then the homology $H_*(P(n, \mathcal{W}, k))$ is free abelian, with a basis given by the classes of the cycles $z(e)$, where e ranges over all cells whose blocks are of the following two types*:

(a) *a singleton which is not a follower*;

(b) *a follower such that its weight and that of its leader add up to at least* $k+1$.

Translating these cellular cycles into cycles in $\mathrm{no}_{k+1}(n, \mathcal{W})$, we get the following picture:

- Singletons and leader–follower pairs correspond to points and groups of points arranged in order along the line.
- Every leader–follower pair corresponds to a set of r points moving back and forth of which any $r-1$ can coincide, but all r cannot.

4 Decomposing cell(n, w) into layers

In this section we will prove Theorems A and B by expressing $H_*(\mathrm{cell}(n, w))$ in terms of the homology of weighted no-$(w+1)$-equal spaces. To this end, we assign an ordering to the top-dimensional cells of $\mathrm{cell}(n)$. For a lower-dimensional cell, its *layer* will be indexed by the first top-dimensional cell containing it; this is when that cell appears in the complex when we think of the complex as glued, in order, out of its top-dimensional cells. Recall that these top cells are identified with permutations in S_n.

The intersections of layers of $\mathrm{cell}(n)$ with $\mathrm{cell}(n, w)$ form the layers of $\mathrm{cell}(n, w)$. We will show that $H_*(\mathrm{cell}(n, w))$ is a direct sum of pieces which appear once each subsequent layer is glued on. Topologically, the layer associated to a permutation $\sigma \in S_n$ is a copy of $P(n-r, \mathcal{W}, w) \times [0,1]^r$, where r and $\mathcal{W}$ depend on σ, which is glued on along $P(n-r, \mathcal{W}, w) \times \partial[0,1]^r$; the combinatorial structure is rather more complicated. It follows (once homological triviality of the gluing is established) that the added summand of $H_j(\mathrm{cell}(n, w))$ is in bijection with elements of $H_{j-r}(P(n-r, \mathcal{W}, w))$. The main technical result of this section states:

Theorem 4.1 *The homology of* $\mathrm{cell}(n, w)$ *decomposes as*

$$H_*(\mathrm{cell}(n, w)) = \bigoplus_{\sigma \in S_n} H_{*-\#\sigma}\big(P(n-\#\sigma, \mathcal{W}(\sigma), w)\big).$$

We will define $\#\sigma$ and $\mathcal{W}(\sigma)$ combinatorially.

From the configuration space point of view, the new cycles in layer σ are those basic cycles from Theorem B that have a particular collection of wheels (irrespective of how those wheels are grouped into filters). The bijection above is given by replacing wheels of k disks by points of weight k. Thus, for example, this bijection in an appropriate layer takes the 14–cycle depicted in Figure 4 to the 3–cycle depicted in Figure 7.

The decomposition depends in a crucial way on the ordering of labels of disks; it is not at all equivariant with respect to the S_n–action on $\mathrm{cell}(n, w)$. Therefore, the methods of this section will tell us little about the S_n–module structure on the homology.

4.1 Combinatorial description of layers

Given a cell in cell(n) broken up into blocks, we further (deterministically) break up each block into *wheels*: each entry of a block is the *axle* of a wheel if it is the largest entry of the block up to that point, and the wheel consists of the axle and all the following smaller entries before the next axle.

Proposition 4.2 *Given a symbol f, the following represent the same permutation $\sigma(f)$ of $[n]$:*

(i) *The lexicographically least* ***shuffle*** *of f. A shuffle is a permutation in which the order of every block is preserved.*

(ii) *The permutation obtained by arranging all the wheels, regardless of block, in ascending order by axle.*

Proof The first number in the lexicographically least shuffle is an axle, because the first element of every block is an axle. Thus, it must be the least axle. The remaining elements of the wheel of that axle are less than that axle and thus are also less than all the other axles. Thus, the next numbers in the lexicographically least shuffle are the remaining elements of the first wheel. Repeating the same argument for each wheel, we deduce inductively from left to right that the two permutations are identical. □

For example, the symbol

$$f = (7\,2 \mid 6 \mid 4\,5\,8\,1\,3)$$

has five wheels: 7 2, 6, 4, 5 and 8 1 3; and therefore $\sigma(f) = 4\,5\,6\,7\,2\,8\,1\,3$. We say that f is in the *layer* $L(\sigma(f))$ of cell(n), or of cell(n, w). By Proposition 4.2, being in a given layer is equivalent to having a given set of wheels. We also write

$$\#\sigma = n - \text{the number of wheels of } \sigma.$$

Notice that $\#\sigma$ is the dimension of the lowest-dimensional cells of $L(\sigma)$, in which every wheel is its own block.

Let g be a boundary cell of f. Then one of the following holds:

(1) The splitting of a block to make g respects the wheels of f: each wheel goes completely into one of the new blocks. Then $g \in L(\sigma(f))$.

(2) At least one wheel of f is decomposed into two or more wheels of g. Then $\sigma(g) \prec \sigma(f)$ lexicographically.

Thus, we can build up cell(n) or cell(n, w) by gluing each subsequent layer, in lexicographical order, onto the union of the previous ones. In other words, the layers define a filtration by subcomplexes

$$L(\leq \sigma) = \bigcup_{\tau \leq \sigma} L(\tau).$$

Now, for each cell of $L(\sigma)$, we can obtain a cell of $P(\{\text{wheels of } \sigma\})$ with the same block structure, replacing each wheel by a single label.

Proposition 4.3 *The k–cells of the layer $L(\sigma)$ of $\mathrm{cell}(n,w)$ are in incidence-preserving bijection with the $(k-\#\sigma)$–cells of the complex $P(n-\#\sigma,\mathcal{W}(\sigma),w)$, where $\mathcal{W}(\sigma)$ consists of the cardinalities of the wheels of σ. In particular, the cells of the layer $L(\sigma)$ of $\mathrm{cell}(n)$ are in incidence-preserving bijection with those of $P(n-\#\sigma)$.*

Proof All cells of $L(\sigma)$ have exactly the same wheels, and the elements of each wheel always appear consecutively. Thus, given a symbol in $L(\sigma)$, we can view it as a sequence of these wheels, separated by vertical bars between blocks. The resulting new symbol is a symbol of $P(n-\#\sigma,\mathcal{W}(\sigma),w)$, because instead of n numbers in each symbol we have $n-\#\sigma$ wheels. If the cell of the original symbol has dimension k, it has $n-1-k$ bars, so the cell of the new symbol also has $n-1-k$ bars, and thus has dimension $n-\#\sigma-1-(n-1-k)=k-\#\sigma$.

As in option (1) above, cell g in $L(\sigma)$ is a boundary cell of f in $L(\sigma)$ if and only if a block in f is split to form g such that each wheel in the block is assigned entirely to the left or entirely to the right. Reinterpreting the symbols to be sequences of wheels, this is the same as the criterion for incidence in $P(n-\#\sigma,\mathcal{W}(\sigma),w)$. □

In fact, this bijection is not just combinatorial, but can be understood from several points of view: via cellular maps, cellular chains or configurations. In the next subsection, we will construct an injective cellular map

$$[0,1]^{\#\sigma}\times P(n-\#\sigma,\mathcal{W}(\sigma),w)\to \mathrm{cell}(n,w)$$

which matches each product of $[0,1]^{\#\sigma}$ with an $(k-\#\sigma)$–cell to a k–cell in $L(\sigma)$. In particular, the image of $\{\frac{1}{2}\}^{\#\sigma}\times P(n-\#\sigma,\mathcal{W}(\sigma),w)$ forms a "core" or "spine" inside the layer, as illustrated in Figure 8. In

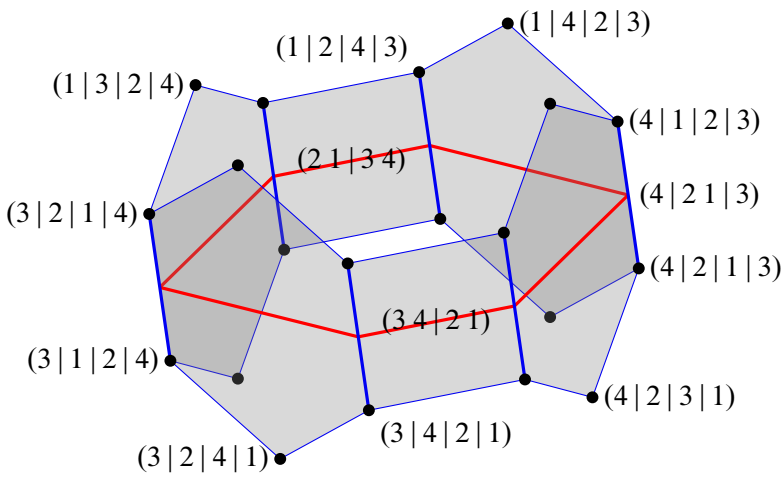

Figure 8: The layer of $\mathrm{cell}(4,3)$ corresponding to the permutation $\sigma = 2\ 1\ 3\ 4$. Selected cells are labeled, and the core (a copy of $P(3,\mathcal{W}(\sigma),3)$, which is the boundary of a $P(3)$) is highlighted in red. The image of $\mathrm{spin}_{\sigma,3}$ is combinatorially the double of the layer along the top and bottom: in the additional cells, 2 1 is replaced by 1 2.

fact, this will be the restriction of a (likewise injective and cellular) map

$$\mathrm{spin}_{\sigma,w}\colon T^{\#\sigma}\times P(n-\#\sigma,\mathcal{W}(\sigma),w)\to L(\leq\sigma)$$

to one top cell of the torus. Here $T^{\#\sigma}$ is given a product cell structure induced by a cell structure on S^1 with two vertices and two edges. In turn, $\mathrm{spin}_{\sigma,w}$ is the restriction of a map

$$\mathrm{spin}_{\sigma}\colon T^{\#\sigma}\times P(n-\#\sigma)\to \mathrm{cell}(n).$$

From the point of view of cellular chains, $\mathrm{spin}_{\sigma,w}$ induces a chain map

$$i_\sigma\colon C_{*-\#\sigma}\big(P(n-\#\sigma,\mathcal{W}(\sigma),w)\big)\to C_*(L(\leq\sigma))$$

via

$$i_\sigma(z)=\mathrm{spin}_{\sigma,w}([T^{\#\sigma}]\times z).$$

In the other direction, we get a surjective chain map

$$p_\sigma\colon C_*(L(\leq\sigma))\to C_{*-\#\sigma}\big(P(n-\#\sigma,\mathcal{W}(\sigma),w)\big)$$

in which cells of $L(<\sigma)$ (defined as $L(<\sigma)=\bigcup_{\tau\prec\sigma}L(\tau)$) are sent to 0 and cells of $L(\sigma)$ are sent to cells of $P(n-\#\sigma,\mathcal{W}(\sigma),w)$, up to sign. From the above discussion, one sees that $p_\sigma\circ i_\sigma=\mathrm{id}$. Together p_σ and i_σ give a splitting

$$H_*(L(\leq\sigma))\cong H_*(L(<\sigma))\oplus H_{*-\#\sigma}\big(P(n-\#\sigma,\mathcal{W}(\sigma),w)\big). \tag{4.4}$$

This implies Theorem 4.1, once we produce the map spin_σ.

Finally, from the configuration point of view, the map $\mathrm{spin}_{\sigma,w}$ associates a configuration in the weighted no-$(w+1)$-equal space to a torus of configurations obtained by replacing points of weight r by wheels of size r and spinning those wheels. This also describes the action of i_σ on cycles. For example, when

$$\sigma=3\ 7\ 5\ 9\ 8\ 10\ 13\ 4\ 15\ 17\ 14\ 18\ 20\ 21\ 12\ 2\ 22\ 16\ 23\ 6\ 11\ 1\ 24\ 19,$$

i_σ sends the 3–cycle in Figure 7 to the 14–cycle in Figure 4. For points in the core, the disks in the wheel are in the "standard" vertically ordered position.

Remark 4.5 Instead of using the splitting, one can prove Theorem 4.1 directly by constructing a discrete gradient vector field and applying Lemma 3.1. Put a total ordering $\prec$ on the cells of $\mathrm{cell}(n,w)$ as follows:

(i) If $\sigma(g)\prec\sigma(f)$, then $g\prec f$.

(ii) If $\sigma(g)=\sigma(f)$, then use the previously defined ordering on the set of cells of $P(n-\#\sigma,\mathcal{W}(\sigma),w)$. This ordering is based on an ordering on the wheels of σ; for this we order first by number of elements, then by axle.

This induces a discrete gradient vector field on $\mathrm{cell}(n, w)$ which restricts to that on $P(n-\#\sigma, \mathcal{W}(\sigma), w)$ on each layer. The images of the bases for $H_*\big(P(n-\#\sigma, \mathcal{W}(\sigma), w)\big)$ under i_σ for all σ give us a set of cycles to which we can apply Lemma 3.1.

For later reference, we describe the set of critical cells of this discrete gradient vector field in a self-contained way. Order wheels in a layer first according to their number of elements and then according to their largest element (*axle*). A block is a *unicycle* if it consists of a single wheel, that is, if its largest element comes first. We assign some blocks to leader–follower pairs by walking left to right: a block is a *follower* if it follows a unicycle (its *leader*) which is not itself a follower and whose wheel is smaller than any of the follower's wheels. A cell of $\mathrm{cell}(n, w)$ is critical if every block is either

(i) a unicycle which is not a follower, or

(ii) a follower such that it and its leader have at least $w+1$ elements in total.

4.2 Maps between permutohedra

To complete the proof of Theorem 4.1, we must still describe the map

$$\mathrm{spin}_\sigma : T^{\#\sigma} \times P(n-\#\sigma) \to \mathrm{cell}(n).$$

Informally, points in $P(n-\#\sigma)$ encode positions of the wheels of σ relative to each other, whereas points in the torus encode positions of disks inside the wheels, which can vary as in Figure 2. In particular, the image of spin_σ will be the union of the $2^{\#\sigma}$ top-dimensional cells obtained from σ by "spinning its wheels", that is, by applying permutations such as those depicted in Figure 9.

To make this precise, we start with the following lemma:

Lemma 4.6 *Let A be a finite set and $\boldsymbol{a}$, $\boldsymbol{b}$ and $\boldsymbol{c}$ additional symbols not in A. Then there is a map*

$$\iota : [0,1] \times P(A \cup \{\boldsymbol{a}\}) \to P(A \cup \{\boldsymbol{b}, \boldsymbol{c}\})$$

with the following properties:

(i) *It is a homeomorphism and a cellular map. That is, it sends every k–face to a disk which is a union of k–faces.*

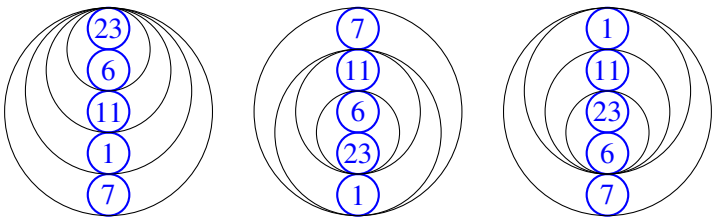

Figure 9: Some configurations of a wheel with five disks which give different orderings, the first of which is the "name" of the wheel.

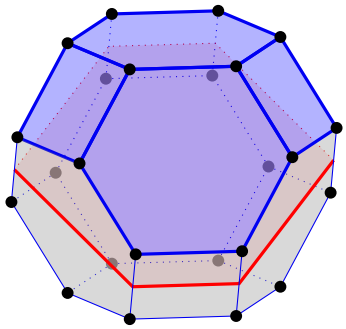

Figure 10: The permutohedron $P(4)$ and the homeomorphism $\iota\colon [0,1]\times P(3)\to P(4)$. The images of $\{\frac{1}{2}\}\times P(3)$ (red) and $\{1\}\times P(3)$ (blue) are highlighted.

(ii) *It is equivariant with respect to the $\mathbb{Z}/2\mathbb{Z}$–actions given by $t\mapsto -t$ on the domain and $\boldsymbol{b}\leftrightarrow\boldsymbol{c}$ on the codomain.*

(iii) *For each cell σ of $P(A\cup\{\boldsymbol{a}\})$, $\iota([0,1]\times\sigma)$ is the cell of $P(A\cup\{\boldsymbol{b},\boldsymbol{c}\})$ with the same blocks except that $\boldsymbol{a}$ is replaced in its block by $\boldsymbol{b}$ and $\boldsymbol{c}$.*

(iv) *It takes $\{0\}\times P(A\cup\{\boldsymbol{a}\})$ to the union of those cells in which $\boldsymbol{b}$ and $\boldsymbol{c}$ are contained in separate blocks with $\boldsymbol{b}$ preceding $\boldsymbol{c}$. Similarly, it takes $\{1\}\times P(A\cup\{\boldsymbol{a}\})$ to the union of those cells in which $\boldsymbol{c}$ precedes $\boldsymbol{b}$.*

Before proving the lemma, we use it to construct spin_σ. First, notice that, for any ordering of $A\cup\{\boldsymbol{b},\boldsymbol{c}\}$, the lemma gives us a well-defined map to the corresponding top cell of $\mathrm{cell}(A\cup\{\boldsymbol{b},\boldsymbol{c}\})$. In particular, if we are instead given an ordering of $A\cup\{\boldsymbol{a}\}$ we have two choices: we can replace $\boldsymbol{a}$ by $\boldsymbol{b}\,\boldsymbol{c}$ or by $\boldsymbol{c}\,\boldsymbol{b}$. Moreover, these choices coincide on $\{0,1\}\times P(A\cup\{\boldsymbol{a}\})$. Identifying the two subspaces gives a map

$$S^1\times P(A\cup\{\boldsymbol{a}\})\to\mathrm{cell}(A\cup\{\boldsymbol{b},\boldsymbol{c}\}).$$

Since top cells of $\mathrm{cell}(A\cup\{\boldsymbol{a}\})$ correspond to orderings of $A\cup\{\boldsymbol{a}\}$, this gives us a map

$$\mathrm{spin}\colon S^1\times\mathrm{cell}(A\cup\{\boldsymbol{a}\})\to\mathrm{cell}(A\cup\{\boldsymbol{b},\boldsymbol{c}\}).$$

The equivariance of ι means that spin is well defined.

We build spin_σ by iterating the map spin. Start by thinking of $P(n-\#\sigma)$ as a top cell of $\mathrm{cell}(\{\text{wheels of }\sigma\})$. At each step, we replace a wheel with k elements by a singleton and a wheel with $k-1$ elements, applying spin to this splitting.

From this description, it is evident that

$$\mathrm{spin}_\sigma\big(P(n-\#\sigma,\mathcal{W}(\sigma),w)\big)=\mathrm{spin}_\sigma(P(n))\cap\mathrm{cell}(n,w).$$

Proof of Lemma 4.6 In the proof, we will use a slightly different notation than in the statement of the lemma: we will replace the finite set A by $\{1,\dots,n-2\}$ and the symbols $\boldsymbol{a}$, $\boldsymbol{b}$ and $\boldsymbol{c}$ by $n-1$, $n-1$ and n, respectively. This lets us write coordinates in $\mathbb{R}^n$ more easily.

The permutohedron $P(n)$ is a *zonotope*, that is, the Minkowski sum of a set of line segments. We refer to [Ziegler 1995, Section 7.3] for known facts about zonotopes.

The standard permutohedron $P(n)$ is the zonotope in $\mathbb{R}^n$ which is the Minkowski sum of the $\binom{n}{2}$ segments connecting every pair of standard unit basis vectors. That is,

$$P(n) = \left\{ \sum_{1 \le i < j \le n} a_{ij} \boldsymbol{e}_i + (1 - a_{ij}) \boldsymbol{e}_j \;\middle|\; 0 \le a_{ij} \le 1 \right\}.$$

Evidently, choosing a different set of n linearly independent vectors in any Euclidean space gives a linearly equivalent polytope. Less obviously, the combinatorial structure of a zonotope only depends on the oriented matroid associated to the set of line segments. In particular, changing the lengths of the line segments without changing their direction does not change the combinatorial structure.

We use this fact to prove the following:

Lemma 4.7 *The polytope $Z(n) = P(n) \cap \{x_{n-1} = x_n\}$ is combinatorially equivalent to $P(n-1)$.*

Proof We first show that $Z(n)$ can also be written as

$$\left\{ \sum_{1 \le i < j \le n} a_{ij} \boldsymbol{e}_i + (1 - a_{ij}) \boldsymbol{e}_j \;\middle|\; 0 \le a_{ij} \le 1,\, a_{i(n-1)} = a_{in},\, a_{(n-1)n} = \tfrac{1}{2} \right\} \subset P(n). \tag{4.8}$$

Indeed, suppose that $\boldsymbol{z} \in Z(n)$, so that

$$\boldsymbol{z} = \sum_{1 \le i < j \le n} b_{ij} \boldsymbol{e}_i + (1 - b_{ij}) \boldsymbol{e}_j \in P(n)$$

and $z_{n-1} = z_n$. If we switch $\boldsymbol{e}_{n-1}$ and $\boldsymbol{e}_n$ in this formula, we get $\boldsymbol{z}$ back. We can get an expression as in (4.8) by averaging these two expressions for $\boldsymbol{z}$, getting

$$a_{i(n-1)} = a_{in} = \tfrac{1}{2}(b_{i(n-1)} + b_{in}), \quad a_{(n-1)n} = \tfrac{1}{2}.$$

From (4.8), we see that $Z(n)$ is the zonotope generated by the segments

$$\boldsymbol{e}_i \leftrightarrow \boldsymbol{e}_j \quad \text{for } 1 \le i < j \le n-2 \qquad \text{and} \qquad \boldsymbol{e}_{n-1} + \boldsymbol{e}_n \leftrightarrow 2\boldsymbol{e}_i \quad \text{for } 1 \le i \le n-2.$$

This zonotope is combinatorially equivalent to one in which segments $\boldsymbol{e}_{n-1} + \boldsymbol{e}_n \leftrightarrow 2\boldsymbol{e}_i$ are replaced by $\frac{1}{2}(\boldsymbol{e}_{n-1} + \boldsymbol{e}_n) \leftrightarrow \boldsymbol{e}_i$. This in turn is linearly equivalent to $P(n-1)$. □

The lemma gives us an injective map $\iota_{1/2} \colon P(n-1) \to P(n)$ which is a homeomorphism onto its image; it remains to extend it to $\iota \colon [0,1] \times P(n-1) \to P(n)$. We do this by letting

$$\iota(t, x) = \iota_{1/2}(x) + R(x)(2t - 1)(\boldsymbol{e}_n - \boldsymbol{e}_{n-1}),$$

where $R(x)$ is the maximum value for which the image still lies in $P(n)$. Property (ii) of the lemma follows immediately from this definition.

To show that ι is the desired map, we must show it is a homeomorphism. To prove injectivity, it suffices to show that $R(x)$ is always positive. But in fact $R(x) \ge \frac{1}{2}$, since we can always vary $a_{(n-1)n}$. Surjectivity

follows from the fact that $P(n)$ is convex and symmetric about the hyperplane $Z(n)$. Finally, ι can only be discontinuous if $R \circ \iota_{1/2}^{-1} \colon Z(n) \to \mathbb{R}$ is discontinuous, but convexity of $P(n)$ implies that this is a convex function, and therefore continuous. Continuity of the inverse likewise follows from the fact that the reciprocal of this function is continuous.

Next we show that ι is cellular. Given a face σ of $P(n-1)$, we know that $\iota_{1/2}(\sigma)$ is a face of $Z(n)$ cut out by a half-space of the hyperplane $\{x_{n-1} = x_n\}$. Extending this half-space orthogonally to $\mathbb{R}^n$, we obtain the half-space that cuts out the face $\iota([0,1] \times \sigma)$ of $P(n)$. It follows that $\iota(\{0\} \times \sigma)$ and $\iota(\{1\} \times \sigma)$ are also unions of faces of $P(n)$ of the appropriate dimension: specifically, each facet of $\iota([0,1] \times \sigma)$ which is not the image of $[0,1] \times \tau$ for some facet τ of σ is contained in one of those two, depending on whether $x_{n-1} - x_n$ is positive or negative for points in that facet.

To show properties (iii) and (iv), we recall the relationship between symbols of cells and the zonotope structure. Namely, points in a face corresponding to a given symbol are those which can be expressed as

$$\sum_{1 \le i < j \le n} a_{ij} \boldsymbol{e}_i + (1 - a_{ij}) \boldsymbol{e}_j ,$$

where a_{ij} is 0 if j is in a later block than i and 1 if i is in a later block than j. When i and j are in the same block, a_{ij} can be anything in $[0,1]$.

From the proof of Lemma 4.7, it follows that $\iota_{1/2}$ maps a face with a given symbol into a face with the same symbol with n added to the same block as $n-1$. Since ι is cellular, this shows (iii). Property (iv) then also follows from our argument that ι is cellular. □

4.3 Proofs of main theorems

Proof of Theorem B Combining the arguments in Section 3 and the proof of Theorem 4.1, we find a basis for $H_*(\mathrm{cell}(n,w))$: for each permutation $\sigma \in S_n$ and each basic cycle z of $H_j\big(P(n - \#\sigma, \mathcal{W}(\sigma), w)\big)$, it contains the cycle $i_\sigma(z)$. The exact set of cycles we get depends on the correspondence between wheels of σ and elements of $[n - \#\sigma]$; to make this concrete, we order the wheels first by size and then by axle.

Geometrically, the correspondence $z \mapsto i_\sigma(z)$ matches

- points moving in $\mathbb{R}$ to wheels moving in $\mathbb{R} \times [0,w]$;
- points moving through each other to wheels moving over and above each other, with the wheel with the smaller axle on top.

We now verify the numbered conditions of Theorem B:

- Condition (i) follows from the way we split σ into wheels.
- Condition (ii) follows from the definition of the layer $L(\sigma)$.
- Conditions (iii) and (iv) follow from Lemma 3.2 and the ordering of the wheels. □

Proof of Theorem A Every basic cycle is a concatenation product of wheels and filters. It is enough to show that there are finitely many (unlabeled) types of wheels and filters in $\mathrm{config}(-, w)$, and that these have at most $\frac{3}{2}w$ disks and generate cycles of dimension at most $\frac{3}{2}w - 2$.

A wheel consists of at most w disks, and a wheel with k disks generates a $(k-1)$–cycle. Filters with at least three wheels have at most $\frac{3}{2}w$ disks, and a filter with k disks generates a $(k-2)$–cycle. Filters with two wheels b_1 and b_2 can be written as $b_2 \,|\, b_1 - b_1 \,|\, b_2$, and do not need to be counted separately. Therefore, $H_*(\mathrm{config}(-, w))$ is spanned by concatenation products of cycles in $H_{\leq 3/2w-2}\big(\mathrm{config}\big(\leq \frac{3}{2}w, w\big)\big)$.

Every wheel with the same number of disks has the same shape, and there are finitely many shapes of filters of width w. Therefore, $H_*(\mathrm{config}(-, w))$ is finitely generated as a twisted algebra. □

5 Betti number growth function

Alpert et al. [2021] examine the growth of the dimension of $H_j(\mathrm{config}(n, w))$ for fixed j and w and varying n, and show that it is asymptotically polynomial times exponential. Having computed a basis for homology, we can now answer the question, is the dimension exactly equal to a polynomial times an exponential? The answer is that it is a sum of such functions, each with a different base of the exponential. The polynomials are integer-valued, and the theory of finite difference calculus states that every integer-valued polynomial in n is an integer linear combination of binomial coefficient functions $\binom{n}{a}$ for various a. Thus, in this section we prove the following theorem:

Theorem 5.1 *For fixed j and w, as a function of n the dimension of $H_j(\mathrm{config}(n, w))$ is an integer linear combination of terms of the form $\binom{n}{a}b^{n-a}$, where a and b are nonnegative integers.*

Note that the case $b = 0$ is permitted, and represents the function that is equal to 1 when $n = a$, and 0 otherwise.

Our proof counts the critical cells from Remark 4.5, and decomposes them into concatenation products of factors that are easier to count. To describe how the counts transform under concatenation product, we introduce the following terminology. A *cell family* F consists of, for each n, a set F_n of cells from $\mathrm{cell}(n)$. Its *counting function* $f(n)$ is the number of cells in F_n. The concatenation product $F \mid G$ of cell families F and G is the cell family consisting of all concatenation products of a cell in F with a cell in G, taken with all possible disk labelings that preserve the order within each factor.[2] The following lemma shows that, to prove our theorem, it suffices to consider separately the various independent factors of a concatenation product:

Lemma 5.2 *Let F and G be two cell families, such that their counting functions $f(n)$ and $g(n)$ are integer linear combinations of terms of the form $\binom{n}{a}b^{n-a}$, where a and b are nonnegative integers. Suppose that, for every element of their concatenation product $F \mid G$, there is only one way to write it as*

[2]This is closely related to the Day convolution defined in Section 2.3, except that F and G need not be FB–subsets of the FB–set of cells of $\mathrm{cell}(-)$, and in our application will not be S_n–equivariant.

the concatenation of an element of F and an element of G. Then the counting function of $F \mid G$ is also an integer linear combination of terms of the form $\binom{n}{a} b^{n-a}$.

Proof Let $f \mid g$ denote the counting function of $F \mid G$. Then we have

$$(f \mid g)(n) = \sum_{i+j=n} \binom{i+j}{i} f(i) g(j),$$

which we refer to as the *labeled convolution* of f and g. Exponential generating functions are convenient for dealing with these labeled convolutions; given the exponential generating functions $\sum_i f(i) x^i / i!$ for f and $\sum_j g(j) x^j / j!$ for g, their product gives the exponential generating function $\sum_n (f \mid g)(n) x^n / n!$ for $f \mid g$.

We can compute the labeled convolution of two terms $\binom{n}{a_1} b_1^{n-a_1}$ and $\binom{n}{a_2} b_2^{n-a_2}$ by converting them to exponential generating functions, taking the product, and then converting the product back. The exponential generating function of $\binom{n}{a} b^{n-a}$ is

$$\sum_n \binom{n}{a} b^{n-a} \frac{x^n}{n!} = \sum_n \frac{x^a}{a!} \frac{(bx)^{n-a}}{(n-a)!} = \frac{x^a}{a!} e^{bx},$$

so the product of exponential generating functions of $\binom{n}{a_1} b_1^{n-a_1}$ and $\binom{n}{a_2} b_2^{n-a_2}$ is

$$\frac{x^{a_1+a_2}}{a_1! a_2!} e^{(b_1+b_2)x} = \binom{a_1+a_2}{a_1} \frac{x^{a_1+a_2}}{(a_1+a_2)!} e^{(b_1+b_2)x},$$

and so the labeled convolution of $\binom{n}{a_1} b_1^{n-a_1}$ and $\binom{n}{a_2} b_2^{n-a_2}$ is

$$\binom{a_1+a_2}{a_1} \binom{n}{a_1+a_2} (b_1+b_2)^{n-(a_1+a_2)}.$$

(This identity can also be verified by constructing sets counted by each of the counting functions.) □

Using this lemma, we are ready to prove Theorem 5.1.

Proof of Theorem 5.1 Given j and w, the critical cells from Remark 4.5 form a cell family, and we want to find the counting function of this cell family. To do this, it helps to count the ways to distribute the singletons separately from counting the ways to arrange the larger wheels. We define a *skyline* to be a critical cell for which the only singleton blocks are leaders of a filter. Given any critical cell, we can find its associated skyline by deleting all the nonleader singletons and then shifting the disk labels down so they are consecutive. For fixed j and w, there are only finitely many possible skylines among the critical cells. Thus, it suffices to count the cell family consisting of all critical cells with a given skyline.

For each skyline, its cell family is the concatenation product of other cell families. We define a *tadpole* to be a critical cell with exactly one filter, which is at the far right, and we define a *tail* to be a critical cell

with no filters. Every critical cell can be written uniquely as a concatenation of some number of tadpoles and possibly one tail on the right. Thus, it suffices to count the cell family consisting of all tadpoles with a given skyline, and the cell family consisting of all tails with a given skyline.

A tadpole with a singleton as the leader of its filter may have other singletons to the left, but a tadpole with a larger wheel as the leader of its filter may not have any singletons. Thus, given a tadpole skyline, if the leader of its filter is not a singleton, the counting function of its cell family returns 1 when the input is the number of disks in the skyline, and 0 otherwise. If the leader of its filter is a singleton, then, for any tadpole with that skyline, the leader of the filter must be disk 1. If k is the total number of disks in the skyline, then the counting function of the tadpole skyline family is $\binom{n-1}{k-1}$, representing the number of ways to choose which disk labels appear in the skyline. Given a tail skyline, if k is the total number of disks in the skyline (possibly 0), the counting function of its cell family is $\binom{n}{k}$.

Thus, for each skyline with a given j and w, we can compute the counting function of its cell family by taking the labeled convolution of the counting functions of its tadpole skyline and tail skyline factors. Because each factor has the desired form, so does the labeled convolution, by Lemma 5.2. The counting function of the family of all critical cells is the sum of the counting functions of all the skylines, and so it also has the desired form. □

6 Cohomology ring

We now give a basis for $H^j(\mathrm{config}(n,w))$ and describe the cup product structure in terms of that basis. We use a similar strategy to that used in [Alpert et al. 2021] to find a lower bound for the dimensions of homology groups. The main tool is Poincaré–Lefschetz duality: for a compact $2n$–manifold with boundary $(M,\partial M)$,

$$H^j(M) \cong H_{2n-j}(M,\partial M).$$

Moreover, when homology classes are realized by submanifolds (as they will be in our case), then:

(i) The pairing between classes in $H^j(M)$ and $H_j(M)$ is given by the transverse signed intersection number between classes in $H_{2n-j}(M,\partial M)$ and $H_j(M)$.

(ii) The cup product between classes in $H^i(M)$ and $H^j(M)$ is given by the transverse intersection map

$$\pitchfork\colon H_{2n-i}(M,\partial M)\otimes H_{2n-j}(M,\partial M)\to H_{2n-i-j}(M,\partial M).$$

While $\mathrm{config}(n,w)$ as previously described is not a compact $2n$–manifold with boundary, we can define a homotopy equivalent compact manifold with boundary:

Definition Let $M(n,w)\subset\mathbb{R}^{2n}$ be the configuration space of open disks of radius 1 in a strip of any finite length $N>n$ and width $w+\varepsilon$ for any $0<\varepsilon<1$.

The boundary consists of those configurations in which a disk touches either another disk or the boundary of the strip. It has corners, but every point of the boundary has a neighborhood homeomorphic to a half-space. (Without the addition of ε, a boundary configuration with w vertically aligned disks would not have a neighborhood homeomorphic to a half-space.)

Alternatively, in an open manifold without boundary (such as the space of configurations of n points in $\mathbb{R} \times (0,1)$ of which no more than w are vertically aligned, as used in Theorem 2.1), Poincaré duality gives an isomorphism $H^j(M) \cong H^{\mathrm{BM}}_{2n-j}(M)$. Here H^{BM}_* indicates *Borel–Moore homology*, the homology of the complex of locally finite chains; this is isomorphic to the homology of a compactification relative to the added points.

6.1 Basis for cohomology

We will describe a basis for $H_{2n-j}(M(n,w), \partial M(n,w))$ by associating basis elements to critical j–cells of $\mathrm{cell}(n,w)$, as described in Remark 4.5. Given a critical cell with symbol f, we define the submanifold $V(f) \subset M(n,w)$ as the set of configurations such that:

(i) The disks in each block of σ are lined up vertically in order.

(ii) If a block b_1 comes before b_2 in f, they have at least $w+1$ elements combined, and one of them is a follower, then the column of disks labeled by elements of b_1 is to the left of that labeled by elements of b_2.

To see that this is a submanifold, and that moreover $\partial V(f) = V(f) \cap \partial M(n,w)$, notice that two columns of disks which have at least $w+1$ elements combined cannot move past each other while still satisfying condition (i).

To show that this is a basis, we describe the intersection pairing between these submanifolds and our generators of $H_j(\mathrm{config}(n,w))$.

Lemma 6.1 *Let f and g be symbols of critical cells of* $\mathrm{cell}(n,w)$. *Write $Z(g)$ for the basic cycle corresponding to g. Then*:

(a) $V(g) \pitchfork Z(g) = \pm 1$.

(b) *If $V(g) \pitchfork Z(f) \neq 0$, then $g \preceq f$ according to the ordering in Remark 4.5.*

From the lemma, we see that, under the ordering of the critical cells by $\prec$, the intersection pairing is described by a triangular matrix. Therefore, the $V(g)$ form a basis for $H_{2n-j}(M(n,w), \partial M(n,w))$.

Proof We use the embedding of Theorem 2.1 to associate $V(g)$ to a cellular j–cocycle $\nu(g)$ in $\mathrm{cell}(n,w)$. The value of $\nu(g)$ on a j–cell is given by the signed intersection number of the embedded cell with $V(g)$.

Suppose that a symbol h is in the support of $\nu(g)$. By comparing condition (i) of the definition of $V(g)$ to the construction of the embedding in Theorem 2.1, we see that f must have the same set of blocks as g, although possibly in a different order; in particular, h and g are in the same layer. Moreover, the barycenter of h is the unique point of intersection. Condition (ii) restricts the possible orderings.

Now suppose that $h \neq g$. We look at the first block, say h_i and g_i, where they differ. This means that h_i is a block which occurs later in g_i, say g_j. The structure of critical cells gives us the following possibilities:

- The block g_i is a follower in g, and g_j is not. Since g_i is a follower, it and g_{i-1} have at least $w+1$ elements in total. On the other hand, condition (ii) implies that g_i and g_j have at most w elements in total, so g_{i-1} has more elements than g_j. Since g_j is not a follower, it must be a unicycle. Then $g_{i-1} \prec g_j = h_i$ in the ordering on wheels, implying that h_i is not a follower in h. Therefore, $g \prec h$.
- There are no followers between g_i and g_j, inclusive. Then g_j and g_i are both unicycles and $g_j \prec g_i$ in the ordering on wheels. Therefore, h_i is not a follower and $h_i \prec g_i$. Because the ordering on cells uses the reverse ordering on wheels, $g \prec h$.
- There is at least one follower between g_i and g_j, and it is not g_i. But g_j can't be a follower, because if it were, $V(g)$ could not contain configurations with g_j to the left of its leader. Then g_j and the first follower after g_i appear in opposite orders in g and h, so g_j must have fewer elements than the first leader after g_i. Since g_j is not a follower, it is a unicycle, and the ordering implies that it must also have fewer elements than g_i. Thus, $h_i \prec g_i$, and therefore $g \prec h$.

In other words, we always have that $g \preceq h$. On the other hand, we know that, if a cell h is in the support of $Z(f)$, then $h \preceq f$. Therefore, $g \preceq f$. This proves (b).

From this we also know that g is the unique cell which is in the support of both $\nu(g)$ and $Z(g)$. Since the coefficient in both cases is ± 1, this proves (a). □

6.2 Cup product structure

The cup product structure of $H^*(\mathrm{config}(n,w))$ is complicated, with many indecomposables as well as many nontrivial products. We start with some simple observations. First, the cohomology algebra of $H^*(\mathrm{config}(n))$ is well understood: it is generated by one-dimensional classes. A good description is given in [Sinha 2013]. Secondly, the pullback of $H^*(\mathrm{config}(n))$ to $H^*(\mathrm{config}(n,w))$ along the inclusion map is a subalgebra and contains all classes of degree less than $w-1$. That means that all classes in $H^{w-1}(\mathrm{config}(n,w))$ which are not pullbacks from $H^{w-1}(\mathrm{config}(n))$ are indecomposable. These include the basis elements corresponding to critical cells which have one leader–follower pair with $w+1$ total elements and in which all other blocks are singletons.

In higher degrees, the story is more complicated. Recall that our basic cocycles are carved out using three kinds of relations: coincidence between the x–coordinates of two disks, vertical ordering of elements within a block, and horizontal ordering between blocks. When we take a cup product, the coincidences from both factors accumulate; see Table 1 for some examples. Likewise, when two blocks are ordered in a product, the ordering must be "inherited" from one of the factors. If many pairs of blocks are ordered, this may force the cohomology class to be indecomposable.

$$\nu(2\,1\,|\,6\,|\,5\,|\,4\,|\,3) \cup \nu(3\,2\,|\,6\,|\,5\,|\,4\,|\,1) = \nu(3\,2\,1\,|\,6\,|\,5\,|\,4)$$
$$\nu(3\,2\,|\,6\,|\,5\,|\,4\,|\,1) \cup \nu(3\,1\,|\,6\,|\,5\,|\,4\,|\,2) = \nu(3\,2\,1\,|\,6\,|\,5\,|\,4) + \nu(3\,1\,2\,|\,6\,|\,5\,|\,4)$$
$$\nu(3\,2\,1\,|\,6\,|\,5\,|\,4) \cup \nu(5\,4\,|\,6\,|\,3\,|\,2\,|\,1) = \nu(3\,2\,1\,|\,5\,4\,|\,6)$$
$$\nu(5\,4\,|\,6\,2\,3\,|\,1) \cup \nu(4\,1\,|\,6\,|\,5\,|\,3\,|\,2) = \nu(5\,4\,1\,|\,6\,2\,3).$$

Table 1: Some examples of nontrivial cup products in config$(6,4)$. These are easy to deduce using intersections of dual cycles.

The last example in Table 1 illustrates an important class of decomposable cohomology classes: those associated to critical cells consisting of only two blocks, where the total number of elements is at least $w+2$. In other words, a filter with more than $w+1$ elements always pairs with a decomposable cohomology class.

We also describe an important set of cases in which cup products are zero.

Proposition 6.2 *If there is a pair of labels i and j which are contained in the same block in both f and g, then $V(f) \pitchfork V(g) = \varnothing$.*

Proof If the two labels are in opposite orders in the two blocks, then the intersection of the two cycles is empty. If they are in the same order, then the intersection may be nonempty, but we can move it off itself by moving every point of $V(f)$ slightly in the x_i–direction (say, move the point p by a distance of $\varepsilon d(p, \partial M(n,w))$ for some $\varepsilon > 0$ small enough). After this operation, $V(f) \cap V(g)$ lies in the boundary of $M(n,w)$. □

In all other cases, we can compute cup products by looking at the intersections of the associated submanifolds.

Proposition 6.3 *If no two blocks in f and g have two labels in common, then $V(f)$ and $V(g)$ intersect transversely.*

Proof Locally, $V(f)$ and $V(g)$ are linear subspaces of config(n), cut out by the linear equations constraining the x–coordinates in each block to coincide. Thus, the dimension of $V(f)$ is n (for the y–coordinates) plus the number of blocks in f (for the x–coordinates), and similarly for $V(g)$. In the intersection, we imagine starting with the constraints for $V(f)$ and including the constraints for $V(g)$ one block at a time. Each block of size k in g merges k different blocks in f into one, so the net change in the number of blocks is $1-k$. In total, the number of blocks in $V(f) \cap V(g)$ is $\#\text{blocks}(f) + \#\text{blocks}(g) - n$, so the local codimension of $V(f) \cap V(g)$ is $2n - \#\text{blocks}(f) - \#\text{blocks}(g)$, which equals $\operatorname{codim} V(f) + \operatorname{codim} V(g)$. □

Finally, we show that there are many indecomposable elements in $H^*(\text{config}(n,w))$, but that they do not occur in the very highest degrees.

Theorem 6.4 *The ring $H^*(\mathrm{config}(n,w))$ has indecomposables*

(a) *only in degree 1 when $w=2$;*

(b) *in every degree between 1 and $\lfloor \frac{1}{2}n \rfloor$ and no others when $w=3$;*

(c) *in degree 1 and in every degree between $w-1$ and $\lfloor \frac{1}{2}(n+w-3) \rfloor$ and no degree greater than $n-\lceil n/(w-1)\rceil$ when $w \geq 4$.*

When $w \geq 4$, indecomposables also seem to occur in most degrees below $n-\lceil n/(w-1)\rceil$, but perhaps not all.

Proof We first show that every class of degree greater than $n-\lceil n/(w-1)\rceil$ is decomposable. The proof does not depend on w. Every cell with fewer than $\lceil n/(w-1)\rceil$ blocks has a block with w elements, and therefore $V(f) \subset M(n,w)$ satisfies an equation of the form $x_{i_1} = \cdots = x_{i_w}$. Therefore, it is enough to show:

Lemma 6.5 *Suppose that $V \subset M(n,w)$ is a connected compact submanifold of dimension at most $2n-w$, satisfying $\partial V \subset \partial M(n,w)$, which is cut out by relations of the form $x_i = x_j$, $y_i < y_j$ and $x_i < x_j$. Suppose furthermore that V satisfies $x_{i_1} = \cdots = x_{i_w}$ for some set of indices $i_1,\dots,i_w$. Then V is the transverse intersection of two proper compact submanifolds satisfying $\partial V \subset \partial M(n,w)$.*

Proof Define W_1 to be the connected component of $\{x_{i_1} = \cdots = x_{i_w}\}$ which contains V; this exists since V is connected. The defining relations of W_1 are:

- $x_{i_k} = x_{i_l}$ for each $k \neq l$.
- $y_{i_k} < y_{i_l}$ or $y_{i_k} > y_{i_l}$ for each $k \neq l$.
- $x_j < x_{i_k}$ or $x_j > x_{i_k}$ whenever $j \neq i_k$ for any k.

Let W_2 be the submanifold cut out by all defining relations of V which involve pairs of points constrained to be to the same side of $x_{i_1},\dots,x_{i_w}$. Then $V = W_1 \cap W_2$, since every necessary relation defining V is a defining relation of either W_1 or W_2. Moreover, by counting the number of defining equalities, we immediately see that the intersection is transverse. □

Applying this to each connected component of $V(f)$, we get a decomposition of the corresponding cohomology class as a sum of cup products.

We now build indecomposable cocycles when $w \geq 4$; the proof for $w=3$ will be similar, but not identical. We will show that basic cocycles corresponding to critical cells of certain shapes are indecomposable. Specifically, we consider a critical cell f which starts with some number r of blocks with two elements, followed by one block (a follower) with $w-1$ elements, and where the remaining blocks are singletons. The degree of such an f is $r+w-2$, and r can be any number between 1 and $\frac{1}{2}(n-w+1)$, giving us all degrees between $w-1$ and $\lfloor \frac{1}{2}(n+w-3) \rfloor$. It is clear that there are critical cells of any such shape; in particular, we select f to be the cell with all entries in order from greatest to least.

We will show that the corresponding cohomology class $\nu(f)$ is indecomposable by induction on r. We will do this by constructing a cycle with which $\nu(f)$ pairs nontrivially, but any decomposable class pairs trivially.

To construct this cycle, first let R be the set of blocks of f which are left of the follower. For any subset $S \subseteq R$, there is a critical cell f_S in which the blocks in S are moved to the right of the follower, and otherwise the ordering on blocks is the same. Let Z_S be the concatenation product of the toroidal cycles obtained by interpreting each block as a wheel and spinning it. This cycle is represented by a map $g_S\colon T^{r+w-2} \to \mathrm{config}(n, w)$. The Z_S represent linearly independent homology classes, since they generate a 2^r–dimensional subspace of $H_{r+w-2}(\mathrm{config}(n, w))$: for every $S \subseteq R$, the basic cycle $Z(f_S)$ is a linear combination of them, given by Z_S itself if $S = R$, or otherwise the difference between Z_S and $Z_{S'}$, where S' is the union of S with the greatest block not in S. We will show that, for every decomposable class ν,

$$\sum_{S \subseteq R} (-1)^{|S|} \langle \nu, Z_S \rangle = 0. \tag{6.6}$$

In particular, ν cannot pair nontrivially with exactly one of the cycles Z_S. On the other hand, $\nu(f)$ pairs nontrivially with Z_S if and only if $S = \varnothing$.

It is enough to show this equation holds for every pairwise cup product, $\nu = \alpha \cup \beta$. We study the pullbacks of such a pairwise cup product along each g_S. Suppose that α is of degree p, and write

$$T^{r+w-2} = \prod_{i \in R \sqcup R'} S^1_i,$$

where R' is the set of degrees of freedom of the $(w-1)$–element wheel. To understand $g_S^* \alpha$, it is enough to understand how it pairs with $T^P = \prod_{i \in P} S^1_i$ for each p–element set $P \subseteq R \sqcup R'$. Then

$$\langle \nu, Z_S \rangle = \sum_{P \sqcup Q = R \sqcup R'} \langle g_S^* \alpha, [T^P] \rangle \cdot \langle g_S^* \beta, [T^Q] \rangle. \tag{6.7}$$

We now show that these decompositions are not independent for different S.

Lemma 6.8 *If $P \cap R' \neq R'$, then $\langle g_S^* \alpha, [T^P] \rangle$ is the same for every S.*

Proof It suffices to show that the pushforward cycles $(g_S)_*[T^P]$ are homologous regardless of S. In fact, the corresponding maps $T^P \to \mathrm{config}(n, w)$ are homotopic. We can compress the allowed movements of the $(w-1)$–element wheel into a subset of the strip of width $w-2$, letting wheels of width 2 pass by. Clearly, wheels of width 2 can also pass by each other since we are assuming $w \geq 4$. Using this set of motions, we can construct a homotopy between any two such maps. □

Lemma 6.9 *If $P \cap R' = R'$, then $\langle g_S^* \alpha, [T^P] \rangle = \langle g_{S'}^* \alpha, [T^P] \rangle$ whenever $(S \triangle S') \cap P = \varnothing$ (where $\triangle$ indicates symmetric difference).*

Proof Again, it suffices to show that $(g_S)_*[T^P]$ and $(g_{S'})_*[T^P]$ are homologous, and again the corresponding maps $T^P \to \mathrm{config}(n,w)$ are homotopic. We can build the homotopy by moving the individual disks of blocks not in P around the $(w-1)$–element set. □

Together, the lemmas imply that, for any pair of nonempty complementary sets $P, Q \subset R \sqcup R'$, the quantity

$$\langle g_S^*\alpha, [T^P]\rangle \cdot \langle g_S^*\beta, [T^Q]\rangle$$

is independent of whether $i \in S$ for at least one $i \in R$. In particular,

$$\sum_{S\subseteq R} (-1)^{|S|}\langle g_S^*\alpha, [T^P]\rangle \cdot \langle g_S^*\beta, [T^Q]\rangle = 0.$$

From here we get (6.6) by summing over all pairs P and Q and using (6.7).

Finally, we deal with the case $w = 3$. In this case, we consider critical cells composed of r two-element blocks (with the bigger element first) followed by $n-2r$ singletons for some $1 \le r \le \lfloor \frac{1}{2}n \rfloor$. Such a cell is critical if and only if the singletons are in reverse order. In particular, every permutation σ of the set of two-element blocks gives a critical cell f_σ. For each $\sigma \in S_r$, let Z_σ be an r–cycle, represented by a map $g_\sigma\colon T^r \to \mathrm{config}(n,w)$, obtained by arranging the blocks in the order dictated by σ and spinning each of them. These cycles are linearly independent in homology since the basic cycles $Z(f_\sigma)$ are all linear combinations of them.

We now define a function $\mu\colon S_r \to \mathbb{Z}$. Consider the expression $[\cdots[[1,2],3],\cdots r]$. If σ cannot be obtained from this by commuting some of the brackets (equivalently, by spinning the wheel $1\,2\cdots r$, as in Figure 9), then $\mu(\sigma) = 0$. If it can, then $\mu(\sigma) = (-1)^c$, where c is the number of commutations required. We will show that, for every decomposable class ν,

$$\sum_{\sigma\in S_r} \mu(\sigma)\langle \nu, Z_\sigma\rangle = 0. \tag{6.10}$$

In particular, ν cannot pair nontrivially with exactly one of the Z_σ, and therefore there is an indecomposable class of degree r.

It is enough to show the equation for every pairwise cup product, $\nu = \alpha \cup \beta$. Write $T^r = \prod_{i=1}^r S_i^1$; for every $P \subset \{1,\dots,r\}$, write $T^P = \prod_{i\in P} S_i^1$. Then

$$\langle \nu, Z_\sigma\rangle = \sum_{P\sqcup Q=\{1,\dots,r\}} \langle g_\sigma^*\alpha, [T^P]\rangle \cdot \langle g_\sigma^*\beta, [T^Q]\rangle. \tag{6.11}$$

Once again, these decompositions are not independent for different σ:

Lemma 6.12 *Whenever σ and τ impose the same ordering on elements of P,*

$$\langle g_\sigma^*\alpha, [T^P]\rangle = \langle g_\tau^*\alpha, [T^P]\rangle.$$

Proof To show that $(g_\sigma)_*[T^P]$ is homologous to $(g_\tau)_*[T^P]$, we construct a homotopy between the corresponding maps $T^P \to \mathrm{config}(n, w)$. This involves moving around the individual disks of the blocks not in P to make sure they are in the right order. □

The lemma implies that, for any pair of nonempty complementary sets $P, Q \subset \{1, \dots, r\}$,

$$\langle g_\sigma^*\alpha, [T^P]\rangle \cdot \langle g_\sigma^*\beta, [T^Q]\rangle = \langle g_\tau^*\alpha, [T^P]\rangle \cdot \langle g_\tau^*\beta, [T^Q]\rangle$$

if τ and σ differ by commuting one of the brackets of $[\cdots[[1,2],3],\cdots r]$, specifically the innermost bracket in which the right side is in Q if $1 \in P$, or vice versa. In particular,

$$\sum_{\sigma \in S_r} \mu(\sigma)\langle g_\sigma^*\alpha, [T^P]\rangle \cdot \langle g_\sigma^*\beta, [T^Q]\rangle = 0.$$

From here we get (6.10) by summing over all pairs P and Q and using (6.11). □

7 Persistent homology

The majority of this paper has investigated the properties of $\mathrm{config}(n, w)$ as we increase n and keep w fixed. But we can also look at what happens when w grows. Since $\mathrm{config}(n, w)$ naturally injects into $\mathrm{config}(n, w+1)$, forming a filtration, the right framework for understanding this is *persistent homology*, which considers homology for all w at once, together with the maps induced by the inclusions. We give a short introduction to the machinery; for more details, see [Edelsbrunner and Harer 2008; Zomorodian and Carlsson 2005].

Specifically, we can regard $\bigoplus_w H_*(\mathrm{config}(n, w))$ as a $\mathbb{Z}[t]$–module in which multiplication by t corresponds to applying the maps $H_*(\mathrm{config}(n, w)) \to H_*(\mathrm{config}(n, w+1))$ induced by the inclusions $\mathrm{config}(n, w) \hookrightarrow \mathrm{config}(n, w+1)$. We denote this $\mathbb{Z}[t]$–module by $\mathrm{PH}_*(\mathrm{config}(n, *))$. Similarly, $\bigoplus_w H^*(\mathrm{config}(n, w))$ is a $\mathbb{Z}[t]$–module in which multiplication by t corresponds to applying the pullback maps $H^*(\mathrm{config}(n, w)) \to H^*(\mathrm{config}(n, w-1))$, and we denote this $\mathbb{Z}[t]$–module by $\mathrm{PH}^*(\mathrm{config}(n, *))$.

A *cyclic* summand of $\mathrm{PH}_*(\mathrm{config}(n, *))$ or $\mathrm{PH}^*(\mathrm{config}(n, *))$ is generated by a single element that is nonzero for some interval of values of w. It is standard to refer to a cyclic summand as a *bar*, and to the endpoints of the corresponding interval as the values w of its *birth* and *death*. A decomposition of $\mathrm{PH}_*(\mathrm{config}(n, *))$ or $\mathrm{PH}^*(\mathrm{config}(n, *))$ into cyclic summands means selecting $\mathbb{Z}$–bases for the various values of w in a way that agrees with the maps between the values of w.

The fundamental theorem for modules over a PID guarantees that a persistence module with coefficients in a field will always decompose into cyclic summands. No such guarantee exists for integral persistence modules: for example, one could have a single class born at one time and later become divisible by 2, yielding a module isomorphic to the ideal $(2, t) \subset \mathbb{Z}[t]$.

Theorems 7.1 and 7.4, which we state and prove below, together prove Theorem C.

Theorem 7.1 *The homology basis elements described in Theorem B induce a decomposition of* $\mathrm{PH}_*(\mathrm{config}(n,*))$ *as the direct sum of cyclic* $\mathbb{Z}[t]$*–modules. The cohomology basis elements from Section 6.1 give a decomposition of* $\mathrm{PH}^*(\mathrm{config}(n,*))$ *as the direct sum of cyclic* $\mathbb{Z}[t]$*–modules.*

To prove the theorem, it suffices to show that, when a given basis element stops being in the basis, it also becomes zero in homology or cohomology. To verify this for homology, we show the corresponding statement for weighted no-$(w+1)$-equal spaces.

Theorem 7.2 *For the weighted no-*$(w+1)$*-equal spaces, the homology basis elements from Theorem B give a* $\mathbb{Z}[t]$*–basis for* $\mathrm{PH}_*(\mathrm{no}_*(n,\mathcal{W}))$.

Proof Recall that a basic cycle consists of a product of single vertices corresponding to singletons and boundaries of permutohedral cells corresponding to leader–follower pairs.

The key fact is that, as we increase w, a particular cell stays critical as long as, for every leader–follower pair, the weights add up to at least $w+1$. But, once they add up to only w for some leader–follower pair, that means that the corresponding permutohedron boundary is filled in, and so the cycle becomes a boundary. Therefore, every cell of $P(n)$ which is critical in $P(n,\mathcal{W},w)$ for some w corresponds to a direct summand of the persistence module. □

Proof of Theorem 7.1 Theorem 7.1 follows easily from Theorem 7.2. This is because the splitting

$$H_*(\mathrm{cell}(n,w)) = \bigoplus_{\sigma\in S_n} H_{*-\#\sigma}\big(P(n-\#\sigma,\mathcal{W}(\sigma),w)\big)$$

of Theorem 4.1 is natural with respect to increasing w (even on the chain level, as one readily sees). Therefore,

$$\mathrm{PH}_*(\mathrm{cell}(n,*)) = \bigoplus_{\sigma\in S_n} \mathrm{PH}_{*-\#\sigma}\big(P(n-\#\sigma,\mathcal{W}(\sigma),*)\big).$$

In particular, an element leaves the basis exactly when one of its filters has the wheel sizes adding up to at most w, and the correspondence with the no-$(w+1)$-equal homology proves that the element is null-homologous in this case.

For cohomology, the restriction map $H^*(\mathrm{config}(n,w)) \to H^*(\mathrm{config}(n,w-1))$ corresponds to intersecting each basis element $V(g)$ of $H_{2n-*}(M(n,w),\partial M(n,w))$ with the smaller space $M(n,w-1)$. When w gets too small for $V(g)$ to be in the basis, it is because some block of g has more than w elements; thus, the intersection of $V(g)$ with this $M(n,w)$ is empty, and the restricted cohomology class is zero. Thus, it is also true for cohomology that, when a given element stops being in the basis, it becomes zero. □

Remark 7.3 We can also explore how the persistence module structure on homology and cohomology interacts with other structures we have discussed:

(1) Cohomological persistence does not play nicely with the cup product structure: frequently a class is born indecomposable at time w and becomes decomposable at time $w-1$. For example, in the notation of the previous section, we have the relation

$$\nu(5\,4 \mid 6\,2\,3 \mid 1) \cup \nu(4\,1 \mid 6 \mid 5 \mid 3 \mid 2) = \nu(5\,4\,1 \mid 6\,2\,3)$$

in $\mathrm{config}(6,3)$ and $\mathrm{config}(6,4)$, but $\nu(5\,4\,1 \mid 6\,2\,3)$ is indecomposable in $\mathrm{config}(6,5)$.

(2) The concatenation product is perfectly well defined on persistence modules. So we can think of $\mathrm{PH}_*(\mathrm{config}(-,*))$ as a graded twisted algebra in $\mathbb{Z}[t]$–modules! The $\mathbb{Z}[t]$–module structure of this algebra is relatively easy to understand, as we have seen above, but it does not interact in a nice way with the symmetric group action. Moreover, unlike the algebras $H_*(\mathrm{config}(-,w))$ for fixed w, it is not finitely generated as an algebra. For these reasons, we do not study this structure further.

7.1 Asymptotics of the persistence module

A main goal of [Alpert et al. 2021] was to understand the growth of Betti numbers of $\mathrm{config}(n,w)$ as n increases. Now that we have described the persistence module of $\mathrm{config}(n,*)$, we can refine this: as the number of disks increases, we keep track of the number of bars of different lengths. It turns out that the number of short bars grows faster and eventually dominates the number of longer bars. We make this precise in the following theorem:

Theorem 7.4 *Each $\mathbb{Z}[t]$–basis element of $\mathrm{PH}_*(\mathrm{config}(n,*))$ born at $w = w_0$ either persists for all $w > w_0$ or dies by $w = 2w_0$. For each j and w_0, either the maps*

$$H_j(\mathrm{config}(n,w_0)) \to H_j(\mathrm{config}(n,w))$$

are isomorphisms for all $w > w_0$ and all n, or the fraction of basis elements of $H_j(\mathrm{config}(n,w_0))$ that persist to $H_j(\mathrm{config}(n,w_0+1))$ approaches 0 as n approaches ∞.

We note that, because each homology basis element is matched to a cohomology basis element in the same degree with the same birth and death, the theorem for homology immediately implies a corresponding statement for cohomology, which we do not include here.

For the second statement of Theorem 7.4, it helps to have an asymptotic estimate of the dimension of $H_j(\mathrm{config}(n,w))$ as n approaches ∞. Examining the basis for homology and estimating the number of elements recovers the following theorem:

Theorem 7.5 [Alpert et al. 2021] *If $w \geq 2$ and $0 \leq j \leq w-2$, then the inclusion of $\mathrm{config}(n,w)$ into the configuration space of points in the plane induces an isomorphism on H_j. If $w \geq 2$ and $j \geq w-1$, then there are positive constants c_1 and c_2, depending on w and j, such that the following is true. Write $j = q(w-1)+r$ with $q \geq 1$ and $0 \leq r < w-1$. Then*

$$c_1 \cdot n^{qw+2r}(q+1)^n \leq \dim H_j(\mathrm{config}(n,w)) \leq c_2 \cdot n^{qw+2r}(q+1)^n.$$

Proof of Theorem 7.4 The basis elements that persist indefinitely are those with no filters. Each filter can be written as a leader–follower pair, and any leader–follower pair that appears at time w_0 has total weight at most $2w_0$, because the leader block has at most the weight of the follower block. Thus, it only remains a filter until at most $w = 2w_0$. This proves the first statement of the theorem.

For the second statement, by Theorem 7.5 it suffices to show that, for $j \geq w-1$, the number of basis elements of $H_j(\mathrm{config}(n,w))$ that persist to $H_j(\mathrm{config}(n,w+1))$ grows polynomially in n. We observe that, if a filter for w contains any wheel of size 1, then its total size is $w+1$, so it does not remain a filter for $w+1$. Thus, in any basis element of $H_j(\mathrm{config}(n,w))$ that persists, every filter must contain only wheels of size at least 2. One rule for being a basis element is "Every wheel immediately to the left of a filter is greater than the least wheel in the filter", so this then implies that, in any basis element of $H_j(\mathrm{config}(n,w))$ that persists, all of the wheels of size 1 appear to the right of all other wheels and filters.

To estimate the number of such basis elements, we simply count the cell symbols that end with (at least) $n-2j$ singleton blocks in descending order. There are $\binom{n}{2j}\cdot(2j)!\cdot 2^{2j-1}$ such symbols, and that function grows polynomially in n. Because the total number of basis elements of $H_j(\mathrm{config}(n,w))$ grows exponentially in n, asymptotically almost all of the basis elements do not persist to $w+1$. □

8 Relations in the twisted algebra and FI_d–modules

The twisted algebra structure of $H_*(\mathrm{cell}(-,w))$ is unusual because many pairs of elements do not commute. In particular, there are some elements that do not commute with the 0–cycles coming from a single disk; we think of these as *barriers* that prevent singleton disks from passing back and forth. This noncommuting with singleton 0–cycles is the main reason for Theorem 7.5: a given homology element in degree j can be written as the concatenation product of up to $\lfloor j/(w-1)\rfloor$ barriers, giving up to $1+\lfloor j/(w-1)\rfloor$ nonequivalent ways to insert a new disk as a singleton.

One algebraic object that exhibits this kind of exponential growth is an FI_d–module. The best example for understanding the idea of an FI_d–module is the jth homology of the configuration space of n disks on the disjoint union of d planes. Each additional disk can be added to any of the d planes. Sam and Snowden [2017] define FI_d–modules for the first time. Ramos [2017] shows that finitely generated FI_d–modules satisfy a notion of generalized representation stability, and [2019] that the homology groups of a certain kind of graph configuration space are finitely generated FI_d–modules.

We show in this section that the homology of unweighted no-$(w+1)$-equal spaces and of $\mathrm{config}(n,2)$ are both FI_d–modules. On the other hand, we give an example to show that the homology of $\mathrm{config}(n,w)$ for $w>2$ is probably not well described via FI_d–modules, since there seems to be no consistent way of decomposing homology classes as products of barriers.

Formally, the category FI_d has one object $[n]=\{1,\dots,n\}$ for each natural number n. The morphisms are pairs (φ,c), where φ is an injection, say from $[n]$ to $[m]$, and c is a d–coloring on the complement of

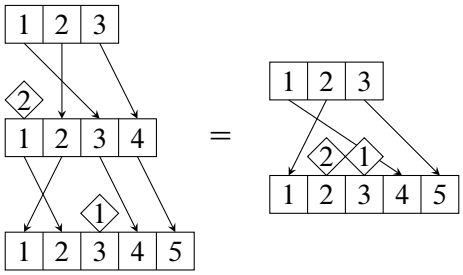

Figure 11: To compose two morphisms in FI_d, we have $(\varphi', c') \circ (\varphi, c) = (\varphi' \circ \varphi, c'')$, where $c''(i)$ is equal to $c'(i)$ if i is not in the image of φ' (for instance, $i = 3$ has color 1 in the example shown) and is equal to $c(\varphi'^{-1}(i))$ if i is in the image of φ' (for instance, $i = 2$ has color 2 in the composition because $c(1) = 2$ and $\varphi'(1) = 2$).

the image of φ; that is, c is a map from $[m] \setminus \varphi([n])$ to a set of size d, which in this paper we choose to be $\{0, 1, \dots, d-1\}$. The morphisms compose as illustrated in Figure 11: for each element colored by the first morphism, in the composition, the image of that element under the second morphism is the one that gets that color. (In the picture, the color of a given element is shown in a diamond just above the element.) More formally, if $(\varphi, c)\colon [n_1] \to [n_2]$ and $(\varphi', c')\colon [n_2] \to [n_3]$ are two morphisms, then we have

$$(\varphi', c') \circ (\varphi, c) = (\varphi' \circ \varphi, c''),$$

where $c''(i)$ is equal to $c'(i)$ if $i \notin \varphi'([n_2])$, and is equal to $c(\varphi'^{-1}(i))$ if $i \in \varphi'([n_2])$.

An FI_d*–module* M over a commutative ring k is defined to be a functor from FI_d to k–modules; that is, we have a k–module M_n for each n, and for each $(\varphi, c)\colon [n] \to [m]$ we have a corresponding k–module map $(\varphi, c)_*\colon M_n \to M_m$. Here we use $k = \mathbb{Z}$. An FI_d–module is *finitely generated* if there exists a finite set of elements $x_1, \dots, x_r \in \bigsqcup_{n=1}^{\infty} M_n$ such that the only FI_d–submodule of M containing $x_1, \dots, x_r$ is M itself. Figure 12 sketches the FI_{j+1}–module structure for $H_j(\mathrm{config}(n, 2))$; the colors of the disks, shown in the picture as the numbers in the diamonds, indicate where to insert the disks between barriers.

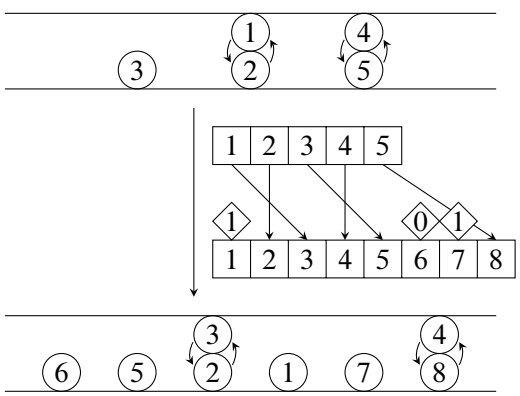

Figure 12: When applying this FI_3–morphism to a class in $H_2(\mathrm{config}(5, 2))$, we insert the disks with color k labels immediately after the kth circling pair.

8.1 No-$(w+1)$-equal spaces

Theorem 8.1 *The graded twisted algebra $H_*(\mathrm{no}_{w+1}(-))$ has a presentation with two generators and two relations. The generators are the representations spanned by a singleton 0–cycle and the $(w-1)$–cycle $\partial(1\,2\cdots w+1)$; these are trivial representations of S_1 and S_{w+1}, respectively. The relations are that two singletons commute and, for each $(w+2)$–block, taking the signed sum of the boundaries of those facets that have a $(w+1)$–block in them gives zero.*

Proof The fact that our proposed generators do generate comes automatically from the description of the basis. To show that the relations are true, we know that $\partial^2 = 0$, in particular when applied to a $(w+2)$–block, so the signed sum of the boundaries of *all* facets of the $(w+2)$–block gives zero. The facets that have no $(w+1)$–block in them are cells in $P(n, w)$, so their boundaries are null-homologous; thus, the sum of the boundaries of the remaining facets gives zero in homology.

To show that the specified relations are sufficient, we take every product of generators that is not in the basis and use the specified relations to write it in terms of the basis. First we may assume that consecutive singletons are always in descending order, using the relation that commutes singletons. Then, if a generator is not in the basis, some boundary of a $(w+1)$–block must be immediately preceded by a singleton that is less than every element of the $(w+1)$–block. If we combine those elements to form a $(w+2)$–block, this substring of our generator is the boundary of one facet of the $(w+2)$–block, and the boundaries of all the other facets are in the basis. Thus, applying the relation replaces a nonbasis substring of our generator by a sum of basis substrings. Applying the relations repeatedly from left to right rewrites our original nonbasis generator in terms of the basis. □

Theorem 8.2 *For each j a multiple of $w-1$, the homology groups $H_j(\mathrm{no}_{w+1}(-))$ form a finitely generated FI_d–module for $d = 1 + j/(w-1)$.*

Proof The FI_d–module structure is as follows. Suppose we have an FI_d–morphism from $[m]$ to $[n]$ with $n \geq m$. For any element of $H_j(\mathrm{no}_{w+1}(m))$, we write it in terms of the generators from Theorem 8.1. We apply the relabeling injection and, for each additional number with color $i \in \{0, 1, \dots, j/(w-1)\}$, we insert an element of that color into each summand between the ith and the $(i+1)^{\text{st}}$ factors of the generator that look like the boundary of a $(w+1)$–block (rather than a singleton). The result is an element of $H_j(\mathrm{no}_{w+1}(n))$.

To show that this map on homology is well defined, we need to show that, if we write an element of $H_j(\mathrm{no}_{w+1}(m))$ in terms of the generators in two different ways, the resulting elements of $H_j(\mathrm{no}_{w+1}(n))$ are the same. To see this, suppose that we apply a relation to an original generator. The same relation can just as easily be applied *after* the relabeling and insertion, just by moving the new singletons past old singletons so that they are out of the way. Thus, applying the relation does not change which homology class we get.

The fact that the maps respect composition of FI_d–morphisms is automatic once we use the fact that consecutive singletons commute. Thus, we have an FI_d–module. It is finitely generated by the basis of $H_j\big(\mathrm{no}_{w+1}(j(w+1))\big)$. □

A proof along similar lines shows that the homology of weighted no-$(w+1)$-equal spaces can be written as a direct sum, with each summand equal to the span of the generators with a particular number of filters. This allows us to insert weight-1 points between the filters in a well-defined way. However, the result is not formally an FI_d–module (or a direct sum of them) in a reasonable way, because relabeling cannot permute points of different weights without leaving the no-$(w+1)$-equal space. Because the statement of the theorem would be cumbersome, we do not include the details.

8.2 Disks in a strip of width $w = 2$

Essentially the same proofs as for the no-$(w+1)$-equal spaces give generators and relations for the configuration spaces for $w = 2$, which then give its homology an FI_d–module structure.

Theorem 8.3 *The graded twisted algebra $H_*(\mathrm{config}(-,2))$ has a presentation with three generators and three relations. The generators are*:

(1) $H_0(\mathrm{config}(1,2)) \cong \mathbb{Z}$, *ie a singleton* 0*–cycle on which* S_1 *acts trivially.*

(2) $H_1(\mathrm{config}(2,2)) \cong \mathbb{Z}$, *ie a* 2*–wheel on which* S_2 *acts trivially. Write this as* $w(1,2) = 2\,1 + 1\,2$.

(3) *The two-dimensional representation of* S_3 *spanned by the cycles*

$z(1,2,3) =$ *and* $z(1,3,2) =$

in $H_1(\mathrm{config}(3,2))$, *where transpositions in* S_3 *act by switching the two basis vectors. This representation is irreducible over* $\mathbb{Z}$, *but over* $\mathbb{Q}$ *it splits into the direct sum of a trivial representation and a sign representation.*

The relations are:

(1) *The singletons commute*: $a \mid b = b \mid a$.

(2) *The relation induced by boundaries of* 4*–cells in* $\mathrm{cell}(4)$.

(3) *The relation*, *in* $H_1(\mathrm{cell}(3,2))$,

$$z(1,2,3) + z(1,3,2) = 1 \mid w(2,3) + w(2,3) \mid 1 + 2 \mid w(3,1) + w(3,1) \mid 2 + 3 \mid w(1,2) + w(1,2) \mid 3.$$

Rationally, one can use the sign representation in $H_1(\mathrm{config}(3,2))$ as a generator, removing the need for the third relation.

Proof We first establish a basis for $H_*(\mathrm{config}(n,2))$ consisting of concatenation products of (1), (2) and (3). We use the same critical cells as in Remark 4.5, but interpret them differently: leader–follower pairs of the form $a \mid b\,c$ where $a < b < c$ become $z(a,b,c)$ factors, while the rest of the blocks individually yield singletons and 2–wheels.

To show that the generators and relations are valid, it suffices to show:

(i) Any concatenation product of generators can be written as a sum of basis elements by rewriting via the relations.

(ii) Every element of the basis from Theorem B can be written as a sum of concatenation products of the new generators.

Together, the two steps imply that the new basis spans $H_*(\mathrm{config}(n, w))$; since it has the same cardinality as the old basis, this shows that it is indeed a basis. The first step then implies that the provided relations are sufficient.

To see (ii), we note that

$$\partial(a\ b\ c) + z(a, b, c) = w(b, c) \mid a + b \mid w(c, a) + w(a, b) \mid c. \tag{8.4}$$

In this fashion we obtain filters of singletons. We already know that every filter of a singleton and a 2–wheel is given by an expression of the form

$$a \mid w(b, c) - w(b, c) \mid a.$$

Every element of the old basis is a concatenation product of these two types of cycles as well as singletons and 2–wheels.

To see (i), note first that, in the description of the basis, the 2–wheels separate the strip into intervals that do not interact with each other or with the 2–wheels; that is, the requirements for being a basis element are the same as the requirements for each of these intervals individually to give a basis element. Then, to rewrite a product of generators in terms of basis elements, we only need to rewrite products of singletons and 3–disk generators.

Relation (3) allows us to eliminate 3–disk generators that are in the "wrong" order. Equation (8.4) lets us turn relation (2) into a way to eliminate subwords of the form $a \mid z(b, c, d)$ where $a < b < c < d$. Thus, we can rewrite every product of singletons and 3–disk generators in normal form using the same method as in Theorem 8.1. □

Theorem 8.5 *For each j, the homology groups $H_j(\mathrm{config}(-, 2))$ form a finitely generated FI_{j+1}–module.*

Proof Every generator of $H_j(\mathrm{config}(n, 2))$ has exactly j factors of types (2) and (3). We define the FI_{j+1}–module structure as follows: to insert a new disk of color $i \in \{0, 1, \dots, j\}$ into a generating cycle, we insert it as a singleton 0–cycle between the ith and $(i+1)^{\text{st}}$ factors of type (2) or (3) of the generator.

As in Theorem 8.2, the relations can be applied either before or after the insertion, so the FI_{j+1}–module structure is well defined. The basis of $H_j(\mathrm{config}(3j, 1))$ gives a finite generating set for the FI_{j+1}–module. □

8.3 Disks in a strip of width $w > 2$

We say that an element of $H_*(\mathrm{config}(n,w))$ is a *barrier* if the two ways to concatenate it with a 1–disk 0–cycle represent different homology classes. The wheels of size w are barriers, as are any filters that contain a wheel of size 1. Because the homology $H_*(\mathrm{config}(n,w))$ is generated by concatenations of wheels and filters, we can count the number of barriers in each generator.

The following proposition implies that counting barriers is not well defined on arbitrary homology classes in $H_*(\mathrm{config}(n,w))$. As mentioned above, this suggests that the structure of $H_*(\mathrm{config}(n,w))$ is not well described by FI_d–modules.

Proposition 8.6 *There is a nontrivial element of $H_4(\mathrm{config}(8,3))$ that is simultaneously a sum of nonbarrier generators, a sum of one-barrier generators, and a sum of two-barrier generators.*

As a warm-up before proving the proposition, we consider the case of $n=4$, $w=3$ and $\sigma = 2\ 1\ 4\ 3$. Then $P(n-\#\sigma, \mathcal{W}(\sigma), w)$ is a no-$(w+1)$-equal space with two points of weight 2, which we denote by [2 1] and [4 3]. Then, because i_σ is a chain map, we have

$$i_\sigma\big(\partial([2\ 1][4\ 3])\big) = \partial\big(i_\sigma([2\ 1][4\ 3])\big).$$

The left-hand side is the sum (with some signs) of the two ways to concatenate the wheels 2 1 and 4 3, whereas the right-hand side is the sum (with some signs) of the four cycles $\partial(2\ 1\ 4\ 3)$, $\partial(2\ 1\ 3\ 4)$, $\partial(1\ 2\ 4\ 3)$, and $\partial(1\ 2\ 3\ 4)$, each of which can be thought of as an image under the S_4–relabeling of $\partial(1\ 2\ 3\ 4)$, which, for width $w=3$, is a filter with four wheels each containing one disk, and is a barrier.

Proof of Proposition 8.6 Let $\sigma = 2\ 1\ 4\ 3\ 6\ 5\ 8\ 7$. Then $P(n-\#\sigma, \mathcal{W}(\sigma), w)$ is a no-$(w+1)$-equal space with four points of weight 2, which we denote by [2 1], [4 3], [6 5] and [8 7]. We can apply any element $\tau \in S_4$ to any chain in $C_*\big(P(n-\#\sigma, \mathcal{W}(\sigma), w)\big)$ by permuting these four labels of points.

To construct our cycle, we take

$$z = \sum_{\tau\in S_4} \mathrm{sign}(\tau)\cdot i_\sigma\big(\tau([2\ 1] \mid [4\ 3] \mid [6\ 5] \mid [8\ 7])\big).$$

That is, we take the signed sum of all of the ways to concatenate the four wheels 2 1, 4 3, 6 5 and 8 7. Each of these summands is already a basis element, although they differ as to which pairs are considered filters of two wheels. Thus, their sum is nonzero in homology, and none of these wheels is a barrier, so it is a sum of nonbarrier generators.

To write our element z as a sum of one-barrier cycles, we pair up all elements τ in S_4 that differ by swapping the middle two wheels. One such pair gives

$$i_\sigma\big([2\ 1] \mid \partial([4\ 3][6\ 5]) \mid [8\ 7]\big),$$

which concatenates the wheels 2 1 (on the left) and 8 7 (on the right) to the cycle

$$i_\sigma\big(\partial([4\ 3][6\ 5])\big) = \partial\big(i_\sigma([4\ 3][6\ 5])\big),$$

which is a sum (with some signs) of relabelings of the barrier filter $\partial(3\ 4\ 5\ 6)$. In this way we can write our element z as a sum of generators, each one a concatenation of a 2–wheel, a barrier filter, and another 2–wheel.

To write our element z as a sum of two-barrier cycles, we group the elements τ in S_4 into quadruples: permutations get grouped together if they differ by swapping the first two and/or the last two wheels. One such quadruple gives

$$i_\sigma\big(\partial([2\ 1][4\ 3]) \mid \partial([6\ 5][8\ 7])\big) = \pm\partial\big(i_\sigma([2\ 1][4\ 3])\big) \mid \partial\big(i_\sigma([6\ 5][8\ 7])\big),$$

which, similar to the computation above, is a sum of generators, each one a concatenation of two barrier filters. □

9 Configuration spaces of unordered disks

The configuration space of n unordered disks of diameter 1 in a strip of width w is the quotient of $\mathrm{config}(n, w)$ by the action of S_n that permutes the disk labels, and it is homotopy equivalent to the quotient of $\mathrm{cell}(n, w)$ by the action of S_n. Because the action is cellular and free, this quotient is a cell complex, which we call $\mathrm{ucell}(n, w)$. In this section, we compute the homology of $\mathrm{ucell}(n, w)$ (and thus of the configuration space of unordered disks in a strip) with field coefficients, using the discrete Morse theory methods from Section 3. In the version of discrete Morse theory that applies here, we do not need to assume that the coefficients of the boundary map are ± 1, as is true for polyhedral cell complexes, but only that these coefficients are units. We use field coefficients throughout this section, so all nonzero coefficients are automatically units.

The concatenation product

$$H_j(\mathrm{ucell}(n, w)) \otimes H_{j'}(\mathrm{ucell}(n', w)) \to H_{j+j'}(\mathrm{ucell}(n + n', w))$$

is well defined since there is no need to choose labels for the disks. Therefore, for any ring R, $H_*(\mathrm{ucell}(*, w); R)$ forms a noncommutative bigraded R–algebra.

The cells of $\mathrm{ucell}(n, w)$ are labeled by symbols as in $\mathrm{cell}(n, w)$, but the numbers in the symbols become indistinguishable; the only remaining information is the sizes of blocks, which we also refer to as weights. We notate them all by $\circ$, and use exponents to denote the weights of the blocks. For instance, in cell(3) we have

$$\partial(1\ 2\ 3) = -1 \mid 2\ 3 + 2 \mid 1\ 3 - 3 \mid 1\ 2 + 2\ 3 \mid 1 - 1\ 3 \mid 2 + 1\ 2 \mid 3,$$

and the corresponding relation in ucell(3) is

$$\partial(\circ^3) = -\circ \mid \circ\circ + \circ \mid \circ\circ - \circ \mid \circ\circ + \circ\circ \mid \circ - \circ\circ \mid \circ + \circ\circ \mid \circ = -\circ \mid \circ\circ + \circ\circ \mid \circ.$$

The following lemma describes all the coefficients of these boundary maps:

Lemma 9.1 *In* ucell(n), *the coefficient of the face* $\circ^k \mid \circ^{n-k}$ *in the boundary of the cell* $\circ^n$ *can be described as follows*:

- *If* k *and* $n-k$ *are both odd*, *the coefficient is* 0.
- *If* $n = 2n'$ *is even and* $k = 2k'$ *is even*, *the coefficient is* $\binom{n'}{k'}$.
- *If* $n = 2n'+1$ *is odd and* $k = 2k'$ *is even*, *the coefficient is* $\binom{n'}{k'}$.
- *If* $n = 2n'+1$ *is odd and* $k = 2k'+1$ *is odd*, *the coefficient is* $-\binom{n'}{k'}$.

Proof Consider the cell $1\,2\cdots n$ in cell(n). In its boundary, there are $\binom{n}{k}$ cells that project to $\circ^k \mid \circ^{n-k}$ in ucell(n), and our task is to add up all of their signs. The sign of each such face is $(-1)^k$ times the sign of the corresponding permutation.

We pair up the numbers 1 and 2, 3 and 4, and so on, pairing up $n-1$ and n if n is even, or leaving only n unpaired if n is odd. We can match and cancel faces of $1\,2\cdots n$ with opposite signs in the following way. Given a face, if 1 and 2 are in different blocks, then swapping them gives another face with opposite sign. Similarly, if 1 and 2 are in the same block, but 3 and 4 are in different blocks, then swapping 3 and 4 gives another face with opposite sign. In this way, for each face for which a pair of numbers is split up, we match and cancel it with another such face by finding the first pair of numbers that is split up, and swapping those numbers.

The remaining faces have 1 and 2 in the same block, 3 and 4 in the same block, and so on, and each one corresponds to an even permutation, so the total sign is $(-1)^k$. If k and $n-k$ are both odd, then there are no such faces. If $k = 2k'$ is even and there are n' pairs, then the faces all have positive sign, and there are $\binom{n'}{k'}$ of them. And, if $k = 2k'+1$ is odd and there are n' pairs, then the faces all have negative sign, and there are $\binom{n'}{k'}$ of them. □

9.1 Discrete Morse theory on ucell(n, w) with $\mathbb{Q}$ coefficients

When ordering the cells to produce a discrete gradient vector field, part of the ordering will be chosen later to be lexicographical. Thus, we need to compute the lexicographically least way to split each block.

Lemma 9.2 *The lexicographically least face of the cell* $\circ^n$ *in* ucell(n) *with nonzero coefficient is given as follows. If* n *is odd*, *the least face is* $\circ^1 \mid \circ^{n-1}$. *If* n *is even and greater than* 2, *the least face is* $\circ^2 \mid \circ^{n-2}$. *In the case* $n = 2$, *the cell* $\circ^2$ *is a cycle.*

Proof If n is odd, then the coefficient of the face $\circ^1 \mid \circ^{n-1}$ is $-\binom{(n-1)/2}{0} = -1$. If $n > 2$ is even, then the coefficient of the face $\circ^1 \mid \circ^{n-1}$ is zero, but the coefficient of the face $\circ^2 \mid \circ^{n-2}$ is $\binom{(n-2)/2}{1} = \frac{1}{2}(n-2) \neq 0$. □

Definition Two consecutive blocks in a symbol in ucell(n, w) form a *leader–follower pair in characteristic* 0 if they have the form $\circ^1 \mid \circ^{2k'}$ or $\circ^2 \mid \circ^{2k'}$ for some k'.

Theorem 9.3 *There is a discrete gradient vector field on* $\mathrm{ucell}(n,w)$ *such that the critical cells are in bijection with a basis for* $H_*(\mathrm{ucell}(n,w);\mathbb{Q})$. *The critical cells are described as follows*:

- *For* $w=2$, *every cell.*
- *For* $w=3$, *the concatenation of zero or more blocks* $\circ^2$, *followed by zero or more singletons* $\circ^1$.
- *For* $w>2$ *even*, *the concatenation of zero or more copies of* $\circ^2 \mid \circ^w$ *or strings ending in* $\circ^1 \mid \circ^w$, *followed by the concatenation of zero or more singletons* $\circ^1$. *For each string ending in* $\circ^1 \mid \circ^w$, *it consists of an optional* $\circ^2$, *followed by zero or more singletons* $\circ^1$, *followed by the pair* $\circ^1 \mid \circ^w$.
- *For* $w>3$ *odd*, *the concatenation of zero or more pairs* $\circ^2 \mid \circ^{w-1}$, *followed by an optional* $\circ^2$, *followed by zero or more singletons* $\circ^1$.

Proof The proof is similar to the proof of Theorem 3.4, which finds a basis for homology of weighted no-$(w+1)$-equal spaces. Informally, we construct a discrete vector field that pairs cells as follows: given a symbol, we read it from left to right, and find the first place where either there is a block of weight greater than 2, in which case we match down by breaking it into a leader–follower pair; or there is a leader–follower pair of total weight at most w, in which case we match up by merging it into one block.

The cells that remain unpaired are those for which every leader–follower pair has total weight greater than w and all other blocks are either $\circ^1$ or $\circ^2$. We observe that, if w is even, the only possibilities for leader–follower pairs in critical cells are $\circ^1 \mid \circ^w$ and $\circ^2 \mid \circ^w$, because the total weight must exceed w while the weight of the follower block must be even and at most w. For the same reason, if w is odd, the only possible leader–follower pair in a critical cell is $\circ^2 \mid \circ^{w-1}$. We also observe that there are no instances of either $\circ^1 \mid \circ^2$ for $w \geq 3$, or $\circ^2 \mid \circ^2$ for $w \geq 4$, because these would be leader–follower pairs. Combining these observations, we deduce that the critical cells must have the form given in the theorem statement.

More formally, to check that this discrete vector field is gradient, we exhibit an ordering that produces it. We order the cells of $\mathrm{ucell}(n,w)$ as follows: if f and g are two cells, we find the first block where they differ, and order according to this first differing block in f and g:

- a follower block is less than a nonfollower;
- two follower blocks are ordered in increasing order of weight; and
- two nonfollower blocks are ordered in decreasing order of weight.

Then we pair up two cells f and g, with f a face of g, if f is the greatest face of g and g is the least coface of f. One can check that this pairing agrees with the pairing described informally above, and thus that the critical cells match the desired description.

To check that the resulting critical cells correspond to a basis, by Lemma 3.1 it suffices to find a cycle $z(e)$ for each critical cell e such that e is the greatest cell in $z(e)$. Because we are using coefficients in $\mathbb{Q}$, we

do not have to worry about whether the coefficient is a unit, as long as it is nonzero. To construct $z(e)$, we take the concatenation product of the following cycles: for each block that is not in a leader–follower pair, it is already a cycle, and, for each leader–follower pair, as our cycle we take the boundary of the block resulting from merging the pair. Note that every cell of this boundary that has a nonzero coefficient has block weight at most w, because the original leader–follower pair is one of the two faces with the most unbalanced block weights in this boundary. Because our leader–follower pair is the lexicographically least face of the merged block, and our ordering is the reverse of lexicographical for nonfollowers, we see that e is the greatest cell in $z(e)$. This implies that the cells $z(e)$ form a basis for $H_*(\mathrm{ucell}(n,w);\mathbb{Q})$. □

Corollary 9.4 *The homology $H_*(\mathrm{ucell}(*,w);\mathbb{Q})$ forms a bigraded algebra over $\mathbb{Q}$ under concatenation product. It has the following generators:*

(1) *the singleton block $\circ^1$;*

(2) *the block $\circ^2$;*

(3) *for $w>3$, the cycle $\partial(\circ^{w+1})$; and*

(4) *for $w>2$ even, the cycle $\partial(\circ^{w+2})$.*

It has the following relations:

(1) *the singleton block $\circ^1$ commutes with $\circ^2$ for all $w\geq 3$;*

(2) *the singleton block $\circ^1$ commutes with $\partial(\circ^{w+1})$ for all odd $w>3$; and*

(3) *the symbol $\circ^2\mid\circ^2$ is null-homologous for all $w\geq 4$.*

Proof The description of the basis shows that the specified generators do generate.

Relation (1) is true because the boundary of the cell $\circ^3$ is $-\circ^1\mid\circ^2+\circ^2\mid\circ^1$. Relation (2) comes from the relation $\partial^2(\circ^{w+2})=0$ on $\mathrm{ucell}(n)$; when w is odd, expanding $\partial(\circ^{w+2})$ gives $-\circ^1\mid\circ^{w+1}+\circ^{w+1}\mid\circ^1$ plus a sum of cells in $\mathrm{ucell}(n,w)$, so applying ∂ again gives our desired homology relation. Relation (3) is true because $\circ^2\mid\circ^2$ is the only face of $\circ^4$ with nonzero coefficient.

The relations are enough to transform an arbitrary product of generators into one of our basis cycles. □

Corollary 9.5 *For fixed j and w, the Betti numbers $\beta_j(\mathrm{ucell}(n,w);\mathbb{Q})$ as a function of n grow with an upper bound of $O(n^q)$, where $q=\lfloor j/(w-1)\rfloor$. If w is odd, the Betti numbers either are 0 for all n or are 1 for all sufficiently large n. If w is even, the Betti numbers either are 0 for all n or grow as $\Theta(n^q)$, and the latter case holds for all $j\geq(w-1)(w-3)$.*

Proof Deleting nonleader singleton blocks $\circ^1$ from a critical cell gives a critical cell with smaller n, and, for each j and w, there are finitely many ways to form one of these "skyline" critical cells that have no nonleader singletons. For each skyline critical cell, the only places to insert singletons are at the end (that

is, on the right side) and immediately before the leader–follower pair $\circ^1 \mid \circ^w$, which exists only if w is even. Thus, if w is odd, then, for all sufficiently large n, the Betti number $\beta_j(\mathrm{ucell}(n,w);\mathbb{Q})$ is constant, equal to the number of skyline critical cells for j and w, which is 1 if j is congruent to 0 or 1 mod $w-1$, and 0 otherwise.

If w is even, for each j we claim that either there is a skyline with $\lfloor j/(w-1)\rfloor$ instances of $\circ^1 \mid \circ^w$, or there is no skyline at all. We write j as $q(w-1)+r$, with $0 \le r \le w-2$. To construct the critical cell, if $r \le q$, we concatenate r instances of $\circ^2 \mid \circ^1 \mid \circ^w$ and then $q-r$ instances of $\circ^1 \mid \circ^w$. If $r = q+1$, we concatenate q instances of $\circ^2 \mid \circ^1 \mid \circ^w$, followed by one instance of $\circ^2$. If $r > q+1$, there is no way to build a critical cell of dimension j, because all blocks that contribute to the dimension are either $\circ^w$, contributing $w-1$, or $\circ^2$, contributing 1, and there can be at most $q+1$ instances of $\circ^2$.

Given a skyline critical cell with n' disks and k instances of $\circ^1 \mid \circ^w$, the number of critical cells with n disks arising from this skyline is $\binom{n-n'+k}{k}$, corresponding to the number of ways to arrange $n-n'$ additional singletons and k dividers. This is a polynomial in n of degree k. If w is even, for each j either there is no skyline, or there is a skyline with $q = \lfloor j/(w-1)\rfloor$ instances of $\circ^1 \mid \circ^w$, which is the largest possible k. Thus, the Betti numbers grow like either 0 or $\Theta(n^q)$.

In the case where w is even and $j \ge (w-1)(w-3)$, the quotient q is at least $w-3$ and the remainder r is at most $w-2$, so the case $r > q+1$ is impossible, and the Betti numbers grow like $\Theta(n^q)$. □

9.2 Discrete Morse theory on ucell(n,w) with $\mathbb{F}_p$ coefficients

Using coefficients mod p, our strategy for computing the homology is the same as with $\mathbb{Q}$ coefficients, but the answer becomes more complicated because we need to account for divisibility of binomial coefficients.

Lemma 9.6 *For any prime p, the lexicographically least face of the cell $\circ^n$ in* ucell(n) *with coefficient not divisible by p is given as follows. If n is odd, the least face is $\circ^1 \mid \circ^{n-1}$. If n is even, then we write n as $2p^k \cdot a$, where a is not divisible by p, and the least face is $\circ^{2p^k} \mid \circ^{2p^k(a-1)}$.*

Proof If n is odd, then the coefficient of the face $\circ^1 \mid \circ^{n-1}$ is $-\binom{(n-1)/2}{0} = -1$, which is a unit in any $\mathbb{F}_p$.

If $n = 2p^k \cdot a$ is even, then the faces of $\circ^n$ have coefficients $\binom{p^k a}{k'}$ for various k'. These are the coefficients of $(x+y)^{p^k a} \equiv (x^{p^k} + y^{p^k})^a$ mod p, using the Frobenius homomorphism. These coefficients are 0 unless p^k divides k', and the coefficient $\binom{p^k a}{p^k}$ is congruent mod p to $\binom{a}{1} = a$, which is a unit in $\mathbb{F}_p$. □

Definition Two consecutive blocks in a symbol in ucell(n,w) form a *leader–follower pair in characteristic p* if they have the form $\circ^1 \mid \circ^{2k'}$ for some k', or the form $\circ^{2p^k} \mid \circ^{2p^k(a-1)}$ for some $k \ge 0$ and a not divisible by p. We assign the pairs disjointly from left to right, so that, once a block is a follower in a pair with the previous block, it cannot also be a leader in a pair with the next block.

Theorem 9.7 *For each prime p, there is a discrete gradient vector field on* $\mathrm{ucell}(n,w)$ *with $\mathbb{F}_p$ coefficients such that the critical cells are in bijection with a basis for $H_*(\mathrm{ucell}(n,w);\mathbb{F}_p)$. The critical cells are those with the properties that every leader–follower pair has total weight greater than w, and every block that is not a follower is either $\circ^1$ or has the form $\circ^{2p^k}$ for some $k \geq 0$. These properties imply that consecutive blocks that are not followers appear in decreasing order of weight, weakly decreasing if $p=2$ and strictly decreasing if $p \neq 2$, except for singleton blocks $\circ^1$ which may occur consecutively for any p.*

Proof The proof is exactly analogous to that of Theorem 9.3, which addresses the case of $\mathbb{Q}$ coefficients. As in that proof, our discrete vector field is informally described by reading each symbol from left to right, breaking down any block of weight other than 1 or $2p^k$ into a leader–follower pair, and combining any leader–follower pair of total weight at most w.

We can describe the resulting critical cells more concretely as follows. To find all possibilities for leader–follower pairs in critical cells, for each $k \geq 0$ such that $2p^k \leq w$, we find the least multiple of $2p^k$ greater than w. If this multiple is not divisible by $2p^{k+1}$, it has the form $2p^k a$ for a not divisible by p, and the leader–follower pair $\circ^{2p^k} \mid \circ^{2p^k(a-1)}$ may appear in a critical cell. We know that $2p^k(a-1)$ is at most w, otherwise it would contradict the selection of $2p^k a$ as the least multiple of $2p^k$ greater than w. If the least multiple of $2p^k$ greater than w is divisible by $2p^{k+1}$, there is no leader–follower pair beginning with $2p^k$ that may appear in a critical cell. In addition, if w is even, the leader–follower pair $\circ^1 \mid \circ^w$ may appear in a critical cell.

For each critical cell, we can imagine dividing the symbol into strings, where the followers are the dividers. Each string consists of blocks $\circ^1$ and/or $\circ^{2p^k}$ for various $k \geq 0$, with constraints on the multiplicities and order because they may not form leader–follower pairs. The pair of blocks $\circ^{2p^k} \mid \circ^{2p^k}$ forms a leader–follower pair if $p \neq 2$, so it may not appear in one of these strings; if $p=2$, it does not form a leader–follower pair, so it may appear. The pairs $\circ^1 \mid \circ^{2p^k}$ for $k \geq 0$ and $\circ^{2p^k} \mid \circ^{2p^l}$ for $k < l$ do form leader–follower pairs regardless of p, so they may not appear in one of these strings. Thus, if $p=2$, each string is an arbitrary sequence of blocks $\circ^1$ and $\circ^{2p^k}$, any number of each, in weakly decreasing order. If $p \neq 2$, each string is an arbitrary sequence of any number of blocks $\circ^1$ and at most one of each block $\circ^{2p^k}$, in decreasing order. Note that, given a leader–follower pair, in the string preceding it, the blocks cannot have smaller weight than the leader, because of the condition that the entire string including the leader should be in decreasing order.

To complete the proof formally, we use the same ordering on cells of $\mathrm{ucell}(n,w)$ as in the proof of Theorem 9.3, except with the characteristic p definition of leader–follower pairs. The resulting discrete gradient vector field agrees with the informal description above. We define cycles $z(e)$ as in the proof of Theorem 9.3: for each leader–follower pair, we take the boundary of the block in $\mathrm{ucell}(n)$ resulting from merging the pair, while each block not in a leader–follower pair is already a cycle, and we take the concatenation product of all these cycles to get $z(e)$. Applying Lemma 3.1, we conclude that the cells $z(e)$ form a basis for $H_*(\mathrm{ucell}(n,w);\mathbb{F}_p)$. □

Corollary 9.8 *For each prime p, the homology $H_*(\mathrm{ucell}(*,w);\mathbb{F}_p)$ forms a bigraded algebra over $\mathbb{F}_p$ under concatenation product. It has the following generators:*

(1) *the singleton block $\circ^1$;*

(2) *the block $\circ^{2p^k}$ for each $k \geq 0$ with $2p^k \leq w$;*

(3) *if w is even but not equal to $2p^k$ for any $k \geq 0$, the cycle $\partial(\circ^{w+1})$; and*

(4) *the cycle $\partial(\circ^{n'})$, where n' is the least multiple of $2p^k$ greater than w, for each $k \geq 0$ such that $2p^k \leq \frac{1}{2}w$; in this case, we let $k(n')$ denote the power of p in the prime factorization of n'.*

It has the following relations:

(1) *the singleton block $\circ^1$ commutes with $\circ^{2p^k}$ whenever $1+2p^k \leq w$;*

(2) *the singleton block $\circ^1$ commutes with $\partial(\circ^{n'})$ whenever $1+n'-2p^{k(n')} \leq w$;*

(3) *if $p \neq 2$, then $\circ^{2p^k} \mid \circ^{2p^k}$ is null-homologous whenever $4p^k \leq w$;*

(4) *we have $\circ^{2p^l} \mid \circ^{2p^k} = -(\circ^{2p^k} \mid \circ^{2p^l})$ whenever $2p^l + 2p^k \leq w$ and $k \neq l$; and*

(5) *the block $\circ^{2p^l}$ commutes with $\partial(\circ^{n'})$ whenever $2p^l + n' - 2p^{k(n')} \leq w$.*

Proof The description of the basis shows that the specified generators do generate.

Relation (1) is true because the boundary of the cell $\circ^{1+2p^k}$ mod p is $-\circ^1 \mid \circ^{2p^k} + \circ^{2p^k} \mid \circ^1$. All other faces have coefficients of 0, because $\binom{p^k}{k'} \equiv 0 \bmod p$ unless k' is 0 or p^k. Relation (2) comes from the relation $\partial^2(\circ^{1+n'}) = 0$ on $\mathrm{ucell}(n)$; expanding $\partial(\circ^{1+n'})$ gives $-\circ^1 \mid \circ^{n'} + \circ^{n'} \mid \circ^1$ plus a sum of cells in $\mathrm{ucell}(n,w)$, so applying ∂ again gives our desired homology relation. Relation (3) is true because $\circ^{2p^k} \mid \circ^{2p^k}$ is the only face of $\circ^{4p^k}$ with nonzero coefficient mod p when $p \neq 2$. (If $p = 2$, then $\circ^{4p^k}$ is a cycle.) To see that relation (4) is true, if $l < k$, the faces of $\circ^{2p^l+2p^k}$ have coefficients $\binom{p^l(1+p^{k-l})}{p^l k'} \equiv \binom{1+p^{k-l}}{k'}$ for various k'. The coefficients for $k' = 1$ and $k' = p^{k-l}$ are both 1 mod p and are the coefficients of $\circ^{2p^l} \mid \circ^{2p^k}$ and $\circ^{2p^k} \mid \circ^{2p^l}$. The other coefficients are all 0 mod p, because Pascal's triangle identity gives $\binom{1+p^{k-l}}{k'} = \binom{p^{k-l}}{k'-1} + \binom{p^{k-l}}{k'}$, which is 0 mod p unless k' is 0, 1, $p^{k-l}-1$ or p^{k-l}. Relation (5) comes from the relation $\partial^2(\circ^{2p^l+n'}) = 0$; using similar reasoning to relation (4), we find that the only faces of $\partial(\circ^{2p^l+n'})$ with nonzero coefficient have the form $\circ^{k'} \mid \circ^{2p^l+n'-k'}$, where k' is congruent to either 0 or $2p^l$ mod $2p^k$. Thus, $\partial(\circ^{2p^l+n'})$ is $\circ^{2p^l} \mid \circ^{n'} + \circ^{n'} \mid \circ^{2p^l}$ plus a sum of cells in $\mathrm{ucell}(n,w)$, and then we may apply ∂ again, keeping in mind that the sign convention for the Leibniz rule gives a negative sign to the term $\circ^{2p^l} \mid \partial(\circ^{n'})$.

The relations are enough to transform an arbitrary product of generators so that consecutive nonfollower blocks are in decreasing order, strictly decreasing if $p \neq 2$. The resulting cycle is in our basis. □

Corollary 9.9 *For fixed j, w and prime p, the Betti numbers $\beta_j(\mathrm{ucell}(n,w);\mathbb{F}_p)$ as a function of n grow with an upper bound of $O(n^q)$, where $q = \lfloor j/(w-1) \rfloor$. If w is odd, the Betti numbers become eventually constant in n; if w is even and $j \geq (w-1)(w-3)$, the Betti numbers grow as $\Theta(n^q)$.*

Proof As in Corollary 9.5, we can delete nonleader singleton blocks $\circ^1$ from each critical cell to form one of finitely many "skylines". To recover all critical cells with a given skyline, we insert singletons $\circ^1$ either on the far right, or immediately preceding a leader–follower pair $\circ^1 \mid \circ^w$, which can only exist if w is even. In this case, the maximum possible number of such pairs is $q = \lfloor j/(w-1) \rfloor$, so, summing over all skylines, we find that the total Betti number is eventually equal to the number of skylines if w is odd, and is bounded above by $O(n^q)$ if w is even.

In the case where $j \geq (w-1)(w-3)$, we can construct a skyline critical cell with q instances of $\circ^1 \mid \circ^w$ in exactly the same way as in Corollary 9.5. The existence of such a skyline implies that $\beta_j(\mathrm{ucell}(n,w);\mathbb{F}_p)$ grows with a lower bound of $\Omega(n^q)$, matching the upper bound. □

10 Open questions and further directions

(1) One generalization of the disks in a strip configuration spaces is the following. Let $E \to B$ be a locally trivial bundle, and consider the configuration space of n distinct points in E such that each fiber of the bundle may contain at most w of them. What do our methods say about the homology of this configuration space? We predict that, if some neighborhoods in B are one-dimensional, then the homology exhibits the same noncommutativity as that of disks in a strip, but that, if B is everywhere at least two-dimensional, then the homology of the configuration space is a finitely generated FI–module.

(2) The representation stability properties of the configuration space of n points on a given manifold come in some sense from the special case where the manifold is Euclidean space; the special case can be considered a local model. In particular, when the manifold has an end, the homology of the manifold configuration space, considered for all n at once, is a module over the twisted commutative algebra given by the homology of the Euclidean configuration space. The algebra acts by inserting cycles near infinity in the end of the manifold. Does our disks-in-a-strip configuration space act as a local model for other configuration spaces, for which the homology exhibits similar finite generation properties due to its being a module? For instance, what can we say about the homology of the configuration space of disks in the product of an interval with a noncompact 1–complex, such as the union of three rays with a common starting point?

(3) Our proof that $H_*(\mathrm{config}(n,w))$ is finitely generated as a twisted noncommutative algebra relies on fully computing $H_*(\mathrm{config}(n,w))$ and exhibiting the generators. Is there a more abstract algebraic framework that would prove finite generation without computing the homology?

(4) Having described $H_*(\mathrm{config}(n,w))$, but not at all equivariantly, we can ask about its S_n–action by permuting the disk labels. In its decomposition into irreducible representations of S_n, how do the multiplicities grow in n? For finitely generated twisted noncommutative algebras in general, by what

patterns can the multiplicities grow? In particular, can one use presentations by generators and relations as in Section 8 to recover this information?

Appendix Computer calculations for small n

In [Alpert et al. 2021], the Betti numbers of $\mathrm{config}(n, w)$ were computed for $n \leq 8$ using off-the-shelf software for computing persistent homology. This involved running the software on the complex $\mathrm{cell}(8)$, which has over 5 million cells. In general, $\mathrm{cell}(n)$ has $2^{n-1}n!$ cells.

We wrote a Python script that harnesses Theorems B and C to compute the persistence diagram of $H_*(\mathrm{cell}(n, *))$. Although the runtime still grows as $n!$, this is faster than the above method by an exponential factor; for $n = 12$, the script ran on a laptop in less than 90 minutes. Figures 13, 14, 15 and 16 show a graphical representation of the resulting persistence diagram. We would like to thank Matthew Kahle for making the first version of this series of figures.

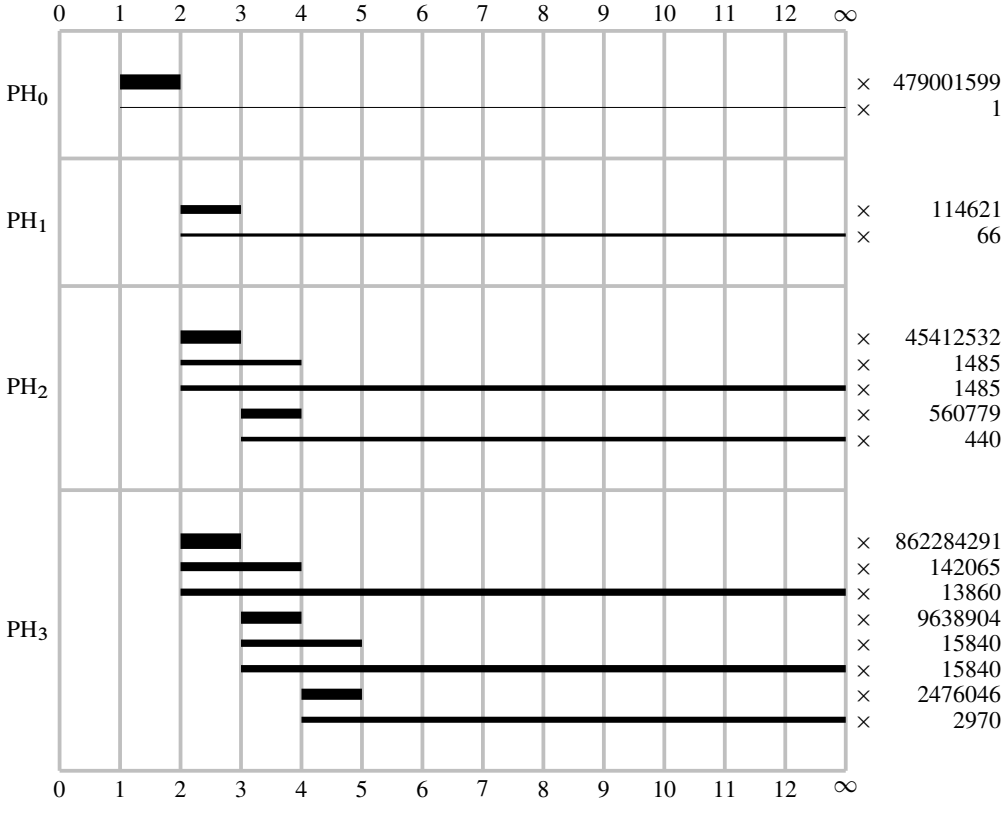

Figure 13: The persistent homology for the configuration space of 12 disks of unit diameter in a strip of width w. The thickness of a bar in the barcode is proportional to the logarithm of the multiplicity, and the exact multiplicity is in the rightmost column.

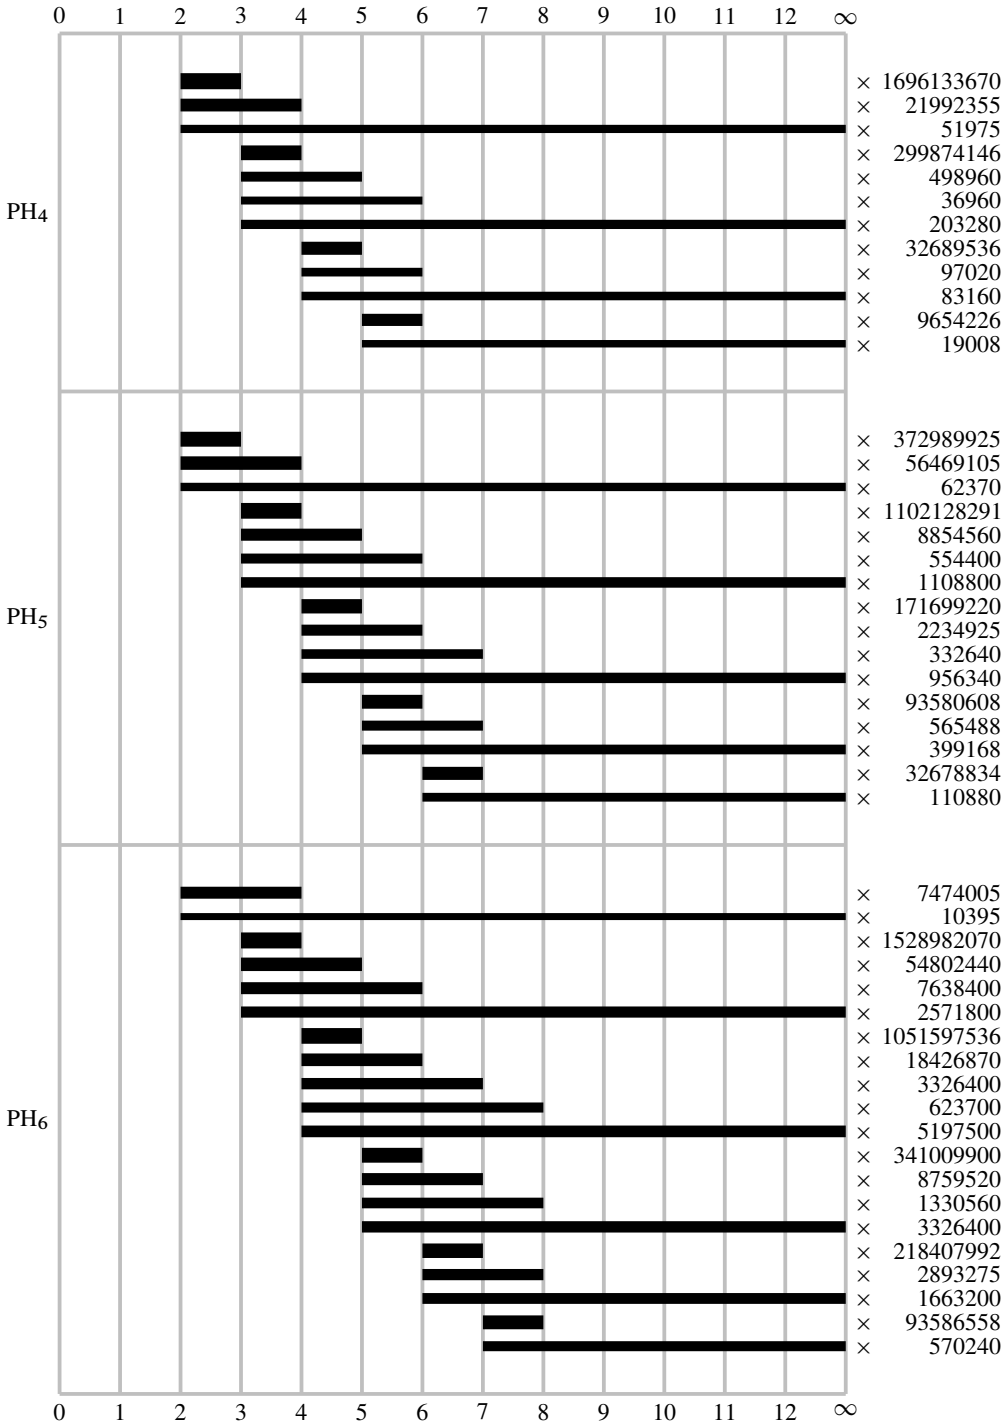

Figure 14: The persistent homology for the configuration space of 12 disks of unit diameter in a strip of width w. The thickness of a bar in the barcode is proportional to the logarithm of the multiplicity, and the exact multiplicity is in the rightmost column.

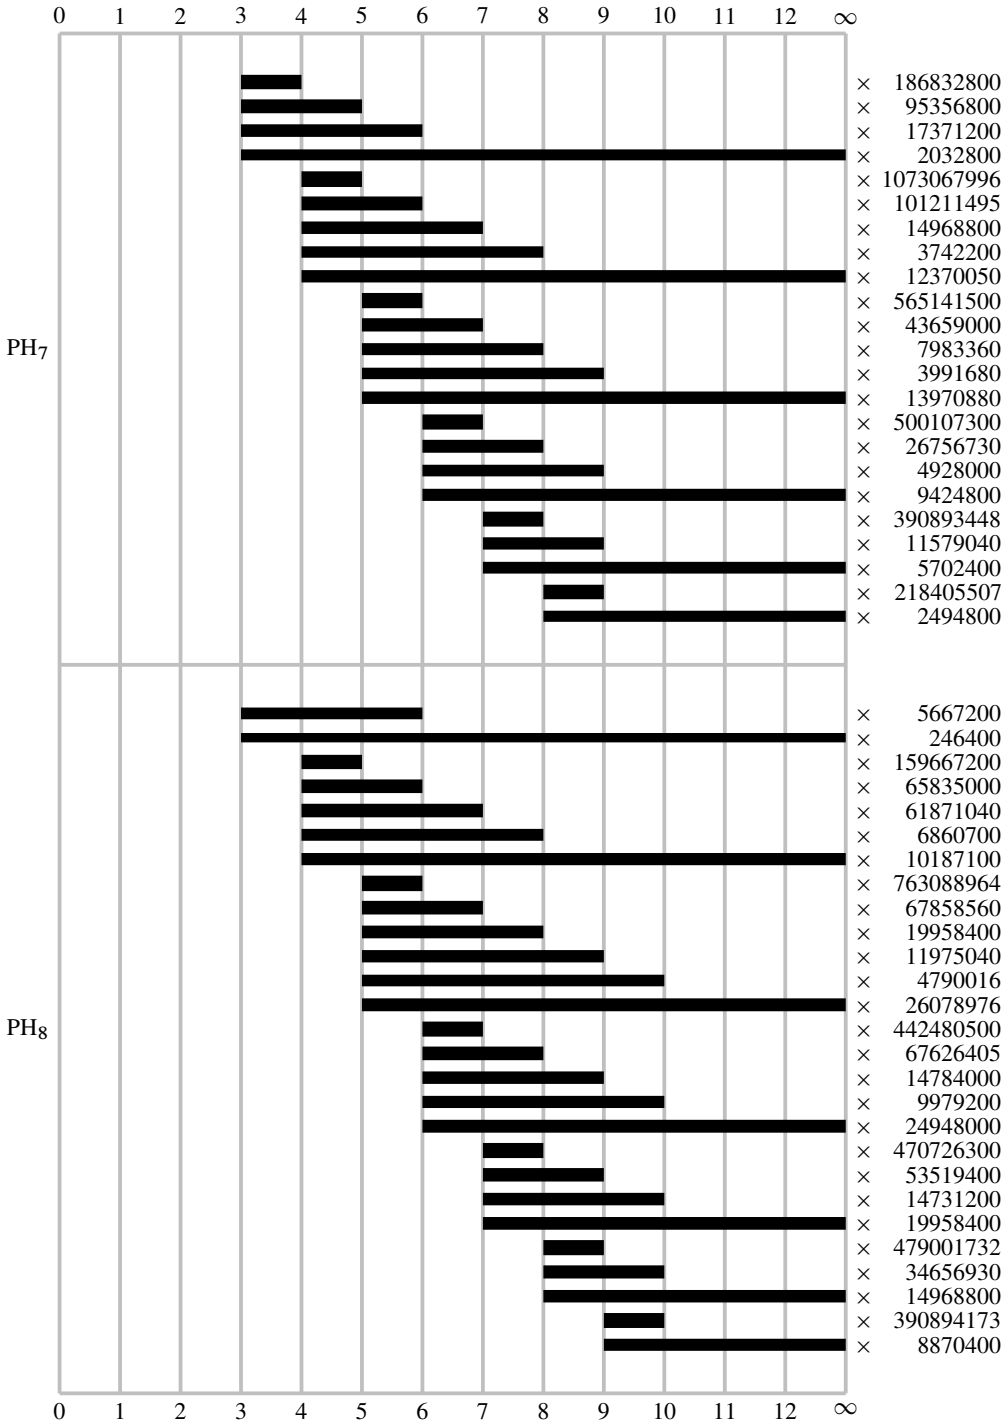

Figure 15: The persistent homology for the configuration space of 12 disks of unit diameter in a strip of width w. The thickness of a bar in the barcode is proportional to the logarithm of the multiplicity, and the exact multiplicity is in the rightmost column.

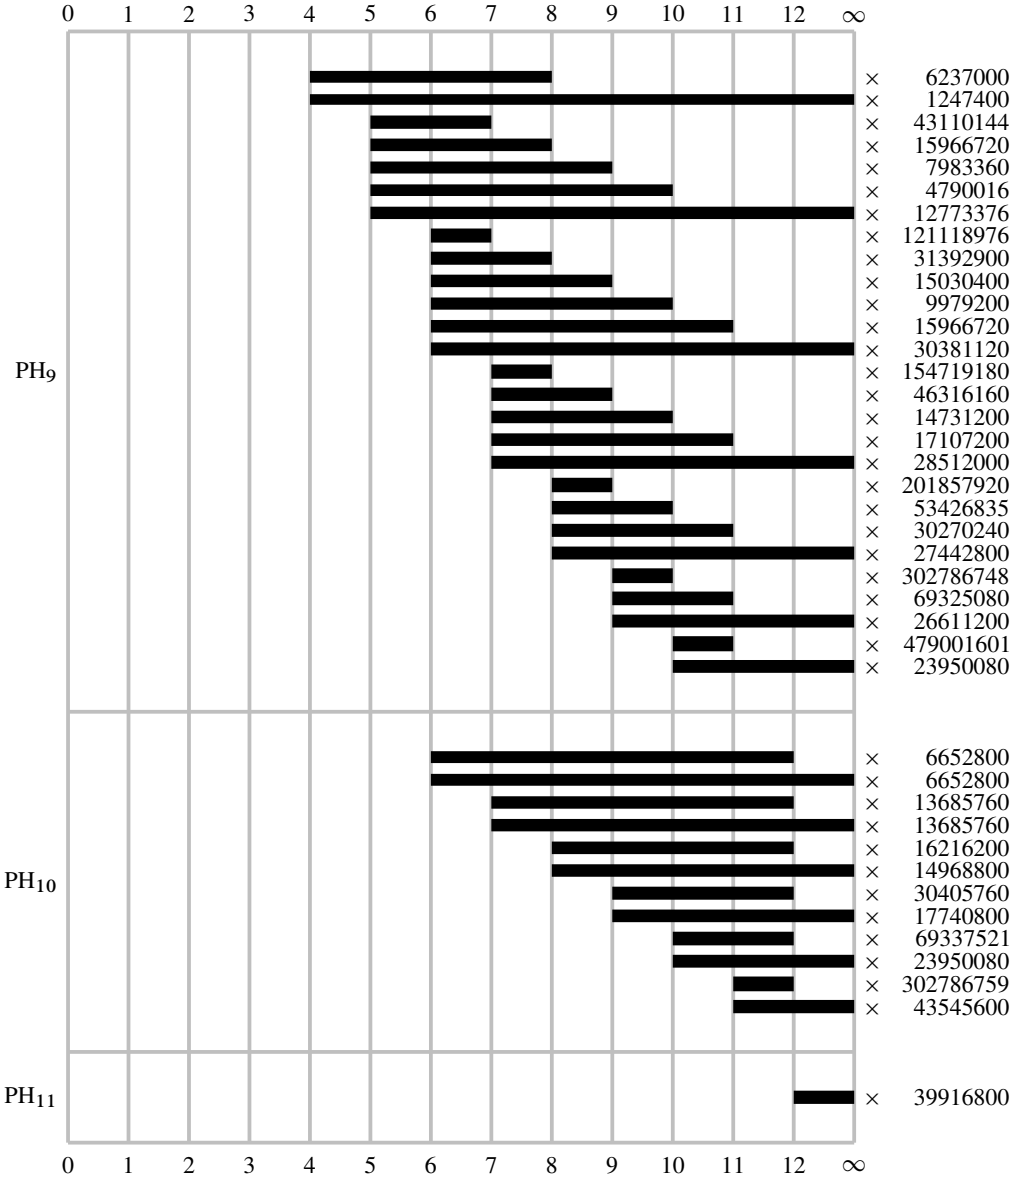

Figure 16: The persistent homology for the configuration space of 12 disks of unit diameter in a strip of width w. The thickness of a bar in the barcode is proportional to the logarithm of the multiplicity, and the exact multiplicity is in the rightmost column.

References

[Alpert 2020] **H Alpert**, *Discrete configuration spaces of squares and hexagons*, J. Appl. Comput. Topol. 4 (2020) 263–280 MR Zbl

[Alpert et al. 2021] **H Alpert**, **M Kahle**, **R MacPherson**, *Configuration spaces of disks in an infinite strip*, J. Appl. Comput. Topol. 5 (2021) 357–390 MR Zbl

[Alpert et al. 2023] **H Alpert**, **U Bauer**, **M Kahle**, **R MacPherson**, **K Spendlove**, *Homology of configuration spaces of hard squares in a rectangle*, Algebr. Geom. Topol. 23 (2023) 2593–2626 MR Zbl

[Arnold 1969] **V I Arnold**, *The cohomology ring of the colored braid group*, Mat. Zametki 5 (1969) 227–231 MR Zbl In Russian; translated in Math. Notes 5 (1969) 138–140

[Baryshnikov et al. 2014] **Y Baryshnikov**, **P Bubenik**, **M Kahle**, *Min-type Morse theory for configuration spaces of hard spheres*, Int. Math. Res. Not. 2014 (2014) 2577–2592 MR Zbl

[Bauer 2021] **U Bauer**, *Ripser: efficient computation of Vietoris–Rips persistence barcodes*, J. Appl. Comput. Topol. 5 (2021) 391–423 MR Zbl

[Björner 2008] **A Björner**, *Random walks, arrangements, cell complexes, greedoids, and self-organizing libraries*, from "Building bridges" (M Grötschel, G O H Katona, editors), Bolyai Soc. Math. Stud. 19, Springer (2008) 165–203 MR Zbl

[Björner and Welker 1995] **A Björner**, **V Welker**, *The homology of "k–equal" manifolds and related partition lattices*, Adv. Math. 110 (1995) 277–313 MR Zbl

[Blagojević and Ziegler 2014] **P V M Blagojević**, **G M Ziegler**, *Convex equipartitions via equivariant obstruction theory*, Israel J. Math. 200 (2014) 49–77 MR Zbl

[Carlsson et al. 2012] **G Carlsson**, **J Gorham**, **M Kahle**, **J Mason**, *Computational topology for configuration spaces of hard disks*, Phys. Rev. E 85 (2012) art. id. 011303 Zbl

[Cohen 1976] **F R Cohen**, *The homology of C_{n+1}-spaces, $n \geq 0$*, from "The homology of iterated loop spaces" (F R Cohen, T J Lada, J P May, editors), Lecture Notes in Math. 533, Springer (1976) 207–351 MR Zbl

[Diaconis 2009] **P Diaconis**, *The Markov chain Monte Carlo revolution*, Bull. Amer. Math. Soc. 46 (2009) 179–205 MR Zbl

[Edelsbrunner and Harer 2008] **H Edelsbrunner**, **J Harer**, *Persistent homology: a survey*, from "Surveys on discrete and computational geometry" (J E Goodman, J Pach, R Pollack, editors), Contemp. Math. 453, Amer. Math. Soc., Providence, RI (2008) 257–282 MR Zbl

[Farber 2008] **M Farber**, *Invitation to topological robotics*, Eur. Math. Soc., Zürich (2008) MR Zbl

[Forman 2002] **R Forman**, *A user's guide to discrete Morse theory*, Sém. Lothar. Combin. 48 (2002) art. id. B48c MR Zbl

[Kahle 2012] **M Kahle**, *Sparse locally-jammed disk packings*, Ann. Comb. 16 (2012) 773–780 MR Zbl

[Mac Lane 1998] **S Mac Lane**, *Categories for the working mathematician*, 2nd edition, Graduate Texts in Math. 5, Springer (1998) MR Zbl

[Miller and Wilson 2019] **J Miller**, **J C H Wilson**, *Higher-order representation stability and ordered configuration spaces of manifolds*, Geom. Topol. 23 (2019) 2519–2591 MR Zbl

[Ramos 2017] **E Ramos**, *Generalized representation stability and FI_d–modules*, Proc. Amer. Math. Soc. 145 (2017) 4647–4660 MR Zbl

[Ramos 2019] **E Ramos**, *Configuration spaces of graphs with certain permitted collisions*, Discrete Comput. Geom. 62 (2019) 912–944 MR Zbl

[Riehl 2017] **E Riehl**, *Category theory in context*, Dover, Mineola, NY (2017) Zbl

[Sam and Snowden 2012] **S V Sam**, **A Snowden**, *Introduction to twisted commutative algebras*, preprint (2012) arXiv 1209.5122

[Sam and Snowden 2017] **S V Sam**, **A Snowden**, *Gröbner methods for representations of combinatorial categories*, J. Amer. Math. Soc. 30 (2017) 159–203 MR Zbl

[Sinha 2013] **D P Sinha**, *The (non-equivariant) homology of the little disks operad*, from "OPERADS 2009" (J-L Loday, B Vallette, editors), Sémin. Congr. 26, Soc. Math. France, Paris (2013) 253–279 MR Zbl

[Ziegler 1995] **G M Ziegler**, *Lectures on polytopes*, Graduate Texts in Math. 152, Springer (1995) MR Zbl

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

Department of Mathematics and Statistics, Auburn University
Auburn, AL, United States

Department of Mathematics, University of California, Santa Barbara
Santa Barbara, CA, United States

hcalpert@auburn.edu, manin@math.ucsb.edu

http://webhome.auburn.edu/~hca0013/, http://web.math.ucsb.edu/~manin/

Proposed: Benson Farb
Seconded: David Fisher, Mladen Bestvina

Received: 13 July 2021
Revised: 9 June 2022

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITOR

András I Stipsicz — Alfréd Rényi Institute of Mathematics, stipsicz@renyi.hu

BOARD OF EDITORS

Mohammed Abouzaid	Stanford University, abouzaid@stanford.edu
Dan Abramovich	Brown University, dan_abramovich@brown.edu
Ian Agol	University of California, Berkeley, ianagol@math.berkeley.edu
Arend Bayer	University of Edinburgh, arend.bayer@ed.ac.uk
Mark Behrens	University of Notre Dame, mbehren1@nd.edu
Mladen Bestvina	University of Utah, bestvina@math.utah.edu
Martin R Bridson	University of Oxford, bridson@maths.ox.ac.uk
Jim Bryan	University of British Columbia, jbryan@math.ubc.ca
Dmitri Burago	Pennsylvania State University, burago@math.psu.edu
Tobias H Colding	Massachusetts Institute of Technology, colding@math.mit.edu
Simon Donaldson	Imperial College, London, s.donaldson@ic.ac.uk
Yasha Eliashberg	Stanford University, eliash-gt@math.stanford.edu
Benson Farb	University of Chicago, farb@math.uchicago.edu
David M Fisher	Rice University, davidfisher@rice.edu
Mike Freedman	Microsoft Research, michaelf@microsoft.com
David Gabai	Princeton University, gabai@princeton.edu
Stavros Garoufalidis	Southern U. of Sci. and Tech., China, stavros@mpim-bonn.mpg.de
Cameron Gordon	University of Texas, gordon@math.utexas.edu
Jesper Grodal	University of Copenhagen, jg@math.ku.dk
Misha Gromov	IHÉS and NYU, Courant Institute, gromov@ihes.fr
Mark Gross	University of Cambridge, mgross@dpmms.cam.ac.uk
Rob Kirby	University of California, Berkeley, kirby@math.berkeley.edu
Bruce Kleiner	NYU, Courant Institute, bkleiner@cims.nyu.edu
Sándor Kovács	University of Washington, skovacs@uw.edu
Urs Lang	ETH Zürich, urs.lang@math.ethz.ch
Marc Levine	Universität Duisburg-Essen, marc.levine@uni-due.de
Ciprian Manolescu	University of California, Los Angeles, cm@math.ucla.edu
Haynes Miller	Massachusetts Institute of Technology, hrm@math.mit.edu
Tomasz Mrowka	Massachusetts Institute of Technology, mrowka@math.mit.edu
Aaron Naber	Northwestern University, anaber@math.northwestern.edu
Peter Ozsváth	Princeton University, petero@math.princeton.edu
Leonid Polterovich	Tel Aviv University, polterov@post.tau.ac.il
Colin Rourke	University of Warwick, gt@maths.warwick.ac.uk
Roman Sauer	Karlsruhe Institute of Technology, roman.sauer@kit.edu
Stefan Schwede	Universität Bonn, schwede@math.uni-bonn.de
Natasa Sesum	Rutgers University, natasas@math.rutgers.edu
Gang Tian	Massachusetts Institute of Technology, tian@math.mit.edu
Ulrike Tillmann	Oxford University, tillmann@maths.ox.ac.uk
Nathalie Wahl	University of Copenhagen, wahl@math.ku.dk
Anna Wienhard	Universität Heidelberg, wienhard@mathi.uni-heidelberg.de

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2024 is US $805/year for the electronic version, and $1135/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

© 2024 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 28 Issue 2 (pages 497–1003) 2024

On the top-weight rational cohomology of $\mathcal{A}_g$ 497

MADELINE BRANDT, JULIETTE BRUCE, MELODY CHAN, MARGARIDA MELO, GWYNETH MORELAND and COREY WOLFE

Algebraic uniqueness of Kähler–Ricci flow limits and optimal degenerations of Fano varieties 539

JIYUAN HAN and CHI LI

Valuations on the character variety: Newton polytopes and residual Poisson bracket 593

JULIEN MARCHÉ and CHRISTOPHER-LLOYD SIMON

The local (co)homology theorems for equivariant bordism 627

MARCO LA VECCHIA

Configuration spaces of disks in a strip, twisted algebras, persistence, and other stories 641

HANNAH ALPERT and FEDOR MANIN

Closed geodesics with prescribed intersection numbers 701

YANN CHAUBET

On endomorphisms of the de Rham cohomology functor 759

SHIZHANG LI and SHUBHODIP MONDAL

The nonabelian Brill–Noether divisor on $\overline{\mathcal{M}}_{13}$ and the Kodaira dimension of $\overline{\mathcal{R}}_{13}$ 803

GAVRIL FARKAS, DAVID JENSEN and SAM PAYNE

Orbit equivalences of $\mathbb{R}$–covered Anosov flows and hyperbolic-like actions on the line 867

THOMAS BARTHELMÉ and KATHRYN MANN

Microlocal theory of Legendrian links and cluster algebras 901

ROGER CASALS and DAPING WENG

Correction to the article Bimodules in bordered Heegaard Floer homology 1001

ROBERT LIPSHITZ, PETER OZSVÁTH and DYLAN P THURSTON