msp

Stability phenomena in the homology of tree braid groups

ERIC RAMOS

For a tree G, we study the changing behaviors in the homology groups $H_i(B_nG)$ as n varies, where $B_nG := \pi_1(\mathrm{UConf}_n(G))$. We prove that the ranks of these homologies can be described by a single polynomial for all n, and construct this polynomial explicitly in terms of invariants of the tree G. To accomplish this we prove that the group $\bigoplus_n H_i(B_nG)$ can be endowed with the structure of a finitely generated graded module over an integral polynomial ring, and further prove that it naturally decomposes as a direct sum of graded shifts of squarefree monomial ideals. Following this, we spend time considering how our methods might be generalized to braid groups of arbitrary graphs, and make various conjectures in this direction.

1 Introduction

1.1 Introductory remarks and statements of the main theorems

In recent years there has been a push towards understanding the mechanisms connecting various well-known asymptotic stability results in topology and algebra. For instance, let $\mathcal{N}$ be a connected oriented manifold of dimension ≥ 2 which is the interior of a manifold with boundary, and write $\mathrm{Conf}_n(\mathcal{N})$ for the n–strand configuration space

$$\mathrm{Conf}_n(\mathcal{N}) := \{(x_1, \ldots, x_n) \in \mathcal{N}^n \mid x_i \neq x_j\}.$$

There is a natural action on $\mathrm{Conf}_n(\mathcal{N})$ by the symmetric group $\mathfrak{S}_n$, and we may therefore define the n–strand unordered configuration space

$$\mathrm{UConf}_n(\mathcal{N}) := \mathrm{Conf}_n(\mathcal{N})/\mathfrak{S}_n.$$

A classical theorem of McDuff [18, Theorem 1.2] implies that for each i and any $n \gg i$, the group $H_i(\mathrm{UConf}_n(\mathcal{N}))$ is independent of n. In contrast to the work of McDuff, it can be seen that the analogous statement is not true for the ordered configuration spaces $\mathrm{Conf}_n(\mathcal{N})$. What is true, however, is perhaps the next best thing. It follows

from work of Church, Ellenberg and Farb [4, Theorem 6.4.3] that for any i, there is a polynomial $P \in \mathbb{Q}[x]$ such that the Betti number $\dim_{\mathbb{Q}}\big(H_i(\mathrm{Conf}_n(U); \mathbb{Q})\big)$ agrees with $P(n)$ for $n \geq 0$. Results of this type fall under the heading of what one might call asymptotic algebra.

The modern philosophy of asymptotic algebra can be roughly stated as follows: a family of algebraic objects which display asymptotic stability phenomena can often times be encoded into a single object, which is finitely generated in an appropriate abelian category. In the case of McDuff, for each i and n the group $H_i(\mathrm{UConf}_n(\mathcal{N}))$ can be realized as the n–graded piece of some finitely generated graded module over $\mathbb{Z}[x]$. The result of Church, Ellenberg and Farb involves showing that the $\mathfrak{S}_n$–representations $H_i(\mathrm{Conf}_n(U); \mathbb{Q})$ are each constituents of some finitely generated representation of the category $\mathcal{FI}$ of finite sets and injections (see [4] for more on this notion). This philosophy is also heavily featured in Sam and Snowden's recent resolution [21] of Stembridge's conjecture. The goal of this paper is to apply similar techniques to the homologies of the unordered configuration spaces of trees. Note that this problem was considered by Lütgehetmann [16]. The results of that work are disjoint from the current work.

In this paper, a *graph* will always refer to a connected, compact CW complex of dimension 1. A *tree* is a graph which is contractible as a topological space. An *essential vertex* of G is a vertex of degree, or valency, at least 3, while an *essential edge* of G is a connected component of the space obtained by removing all essential vertices from G. Note that both the essential edges and vertices of a graph are unaffected by subdivision of edges, and can be thought of as the topologically essential pieces of the graph.

Our main result relates to the asymptotic behavior of the homologies of the braid group of a tree. To state this result, we first need to define a kind of connectivity invariant for trees.

Let G be a tree. Then we set

$$\Delta_G^i := \max_{\{\{v_{j_1}, \ldots, v_{j_i}\} \mid v_{j_k} \text{ essential}\}} \big\{\dim_{\mathbb{Q}}(H_0(G - \{v_{j_1}, \ldots, v_{j_i}\}; \mathbb{Q}))\big\}.$$

In words, Δ_G^i is the maximum number of connected components that G can be broken into by removing exactly i essential vertices. Therefore, Δ_G^1 is the maximum degree of a vertex in G, while, if G has N_G essential vertices, $\Delta_G^{N_G}$ is the number of essential edges of G. By convention, $\Delta_G^0 = 1$, while $\Delta_G^i = 0$ for $i > N_G$.

Theorem A *Let G be a tree and write $B_n G := \pi_1(\mathrm{UConf}_n(G))$ for the braid group. Then for each i there is a polynomial $P_i \in \mathbb{Q}[x]$ of degree $\Delta_G^i - 1$ such that, for all $n \geq 0$,*

$$P_i(n) = \dim_{\mathbb{Q}}(H_i(B_n G; \mathbb{Q})).$$

Remark 1.1 It follows from a theorem of Ghrist [13, Theorem 3.3], and, independently, a theorem of Świątkowski [22, Theorem 0.1], that given any graph G not homeomorphic to S^1, $H_i(B_n G) = 0$ for all i strictly greater than the number of essential vertices of G. This is realized in the case where G is a tree in Theorem A by the fact that $\Delta_G^i - 1 = -1$ in these cases.

It was brought to the author's attention that recent work of Maciążek and Sawicki [17] proves a version of Theorem A. Their work approaches the problem from a very different perspective, and is motivated by questions in physics. The author is grateful towards Professors Maciążek and Sawicki for informing him of their work.

The polynomial P_i of Theorem A is explicitly computed throughout the course of this work (see Theorem 3.17). This computation implies something somewhat surprising about these homology groups.

Corollary B *Let G be a tree, and let $i \geq 0$. Then the homology groups $H_i(B_n G)$ depend only on i, n and the degree sequence of G.*

It is interesting to note that the rank of $H_i(B_n G)$ agrees with a polynomial for all $n \geq 0$, as opposed to only agreeing for n sufficiently large. In the case of configuration spaces of manifolds, a result of this kind does have precedent. We have already discussed the result of Church, Ellenberg and Farb [4, Theorem 6.4.3], which states that if $\mathcal{N}$ is an oriented manifold which is the interior of a manifold with nonempty boundary, then for any i the dimension of $H_i(\mathrm{Conf}_n(\mathcal{N}); \mathbb{Q})$ agrees with a polynomial for all n. It is perhaps an interesting observation that trees can be thought of as graphs with nontrivial boundary. It is unclear whether the connection to the work of Church, Ellenberg and Farb goes any deeper than this, however.

1.2 An outline of the proof

To prove Theorem A, we will rely on classical techniques in commutative algebra, as well as more modern techniques in combinatorial topology. The first key ingredient is the discrete Morse theory of Forman [10]. Given a regular CW complex X (see

Definition 2.1), write $\mathcal{K}_i$ for the set of i–cells of X. A *discrete Morse function* is a map f from the cells of X to $\mathbb{R}$ satisfying the following two hypotheses for all cells $\sigma \in \mathcal{K}_i$:

(1) $\left| \{ \tau \in \mathcal{K}_{i+1} \mid \sigma \subseteq \overline{\tau} \text{ and } f(\sigma) \geq f(\tau) \} \right| \leq 1.$

(2) $\left| \{ \tau \in \mathcal{K}_{i-1} \mid \tau \subseteq \overline{\sigma} \text{ and } f(\sigma) \leq f(\tau) \} \right| \leq 1.$

We call $\sigma \in \mathcal{K}_i$ a *critical i–cell* with respect to f if the sets of conditions (1)–(2) are both empty. The main consequence of discrete Morse theory is that the critical cells of X determine its homotopy type. Formally, let $a \in \mathbb{R}$ and write $X(a)$ for the subcomplex of X comprised of the closures of cells σ with $f(\sigma) \leq a$. If there is no critical cell σ with $a < f(\sigma) < b$, then $X(a)$ is a deformation retract of $X(b)$. Otherwise, if there is a unique critical cell σ with $a < f(\sigma) < b$, then $X(b)$ is obtained from $X(a)$ by attaching the cell σ. Moreover, there is a complex

$$(1\text{-}1) \qquad \cdots \to \mathcal{M}_i \to \cdots \xrightarrow{\tilde{\partial}} \mathcal{M}_1 \to \mathcal{M}_0 \to 0,$$

where $\mathcal{M}_i$ is a free $\mathbb{Z}$–module on the critical i–cells of X, whose homology is the homology of the space X. We call the differential $\tilde{\partial}$ the *Morse differential* (see Definition 2.5).

Using work of Abrams [1], Farley and Sabalka were able to impose a discrete Morse structure on the spaces $\mathrm{UConf}_n(G)$, where G is any graph [8]. We will spend a good amount of time recounting the construction of Farley and Sabalka in Section 2.2. Once we have accomplished this, our strategy will be to develop a strong understanding of the critical cells. More specifically, we will work towards understanding the changing behaviors of the critical cells as n varies.

Let G be a graph with E essential edges, and write $S_G := \mathbb{Z}[x_1, \ldots, x_E]$ for the integral polynomial ring in E variables. We will prove the following in Section 3.2:

($*$) For each i, there exists a finitely generated graded S_G–module $\mathcal{M}_{i,\bullet}$ for which $\mathcal{M}_{i,n}$ is a free $\mathbb{Z}$–module with basis vectors indexed by the critical i–cells of $\mathrm{UConf}_n(G)$.

Specializing to the case where G is a tree, work of Farley [7] implies that the Morse differential is always trivial. Using the complex (1-1) we obtain the following:

Theorem C *Let G be a tree with E essential edges, and let S_G denote the integral polynomial ring in E variables. Then, for each i and n, the action of S_G on the critical*

cells of $\mathrm{UConf}_n(G)$ described in Section 3.2 imposes the structure of a finitely generated graded S_G–module on the abelian group $\bigoplus_{n\geq 0} H_i(\mathrm{UConf}_n(G))$. In particular, there exists a finitely generated graded S_G–module $\mathcal{H}_i$ such that

$$(\mathcal{H}_i)_n \cong H_i(\mathrm{UConf}_n(G)).$$

A result of Ghrist and Abrams (see Theorem 2.13) implies the spaces $\mathrm{UConf}_n(G)$ are aspherical. It follows immediately from this that $H_i(\mathrm{UConf}_n(G)) \cong H_i(B_nG)$. With this in mind, the existence of a polynomial $P_i \in \mathbb{Q}[x]$ which agrees with the Betti numbers of $H_i(B_nG)$ follows from the existence of the Hilbert polynomial in the study of graded modules over polynomial rings. Of course, the above theorem does not tell us anything about the degree of the Hilbert polynomial, nor does it bound its obstruction. To accomplish this, we must first prove a structure theorem about the modules $\mathcal{H}_i$.

Remark 1.2 We will, throughout the paper, make casual use of the theory of graded modules over polynomial rings. More specifically, we will often rely upon the existence of the Hilbert polynomial in this setting. For the reader unfamiliar with this material, it is recommended that they read through the first chapter of Eisenbud [5].

To state this theorem, we first recall the definition of a *squarefree monomial ideal*. We say an ideal $I \subseteq \mathbb{Q}[x_1, \ldots, x_d]$ is a squarefree monomial ideal if it contains a generating set of monomials, none of which are divisible by a square in $\mathbb{Q}[x_1, \ldots, x_d]$. These ideals are the subject of Stanley–Reisner theory, and have many desirable properties. For instance, much is known about their Hilbert polynomial (see Miller and Sturmfels [19] for a reference on the subject).

Theorem D Let G be a tree, and let $\mathcal{H}_i$ denote the S_G–module of Theorem C. Then $\mathcal{H}_i \otimes_{\mathbb{Z}} \mathbb{Q}$ is isomorphic to a direct sum of graded twists of squarefree monomial ideals, each having dimension at most Δ_G^i.

We will find that this theorem implies the conclusions of Theorem A. In fact, we will be able to compute the polynomial P_i associated to $\mathcal{H}_i$ explicitly in terms of invariants of the tree G (see Theorem 3.17).

1.3 An overview of the paper

In the next section, we will spend time developing necessary background. This includes short summaries of discrete Morse theory (Section 2.1), and configuration spaces of

graphs (Sections 2.2 and 2.3). Following this, we use the machinery developed in these preliminary sections to prove the statement $(*)$ (Section 3.2). Finally, we specialize to the case of trees, and use enumerative combinatorial methods to prove Theorem A via an explicit computation of the Hilbert polynomial (Sections 3.3 and 3.4).

To finish the paper, we briefly consider the case of a general graph G. Note that while most of the explicit results in this paper are limited to the case where G is a tree, the statement $(*)$ will hold for any graph G. A result like Theorem C will therefore hold for general graphs so long as we know that the Morse differential commutes with the action of S_G on $\mathcal{M}_{i,\bullet}$. It is the belief of the author that the action of S_G, or perhaps a slight alteration thereof, will indeed commute with this differential. Unfortunately, it is known that this differential can become tremendously complicated as G increases in complexity (see Ko and Park [15]). In any case, the methods in this work therefore provide at least a strategy for proving stability results for more general graphs. In the final sections, we discuss some implications of this.

Acknowledgments

2 Preliminaries

2.1 Discrete Morse theory

We now take the time to briefly summarize the key points in Forman's discrete Morse theory [10]. We will largely be following the exposition of Forman [10; 11], Farley and Sabalka [8] and Ko and Park [15].

In the introduction, we spent some time discussing the notion of discrete Morse function. One thing that should stand out about this definition is that the literal values of the

function are immaterial. Namely, the classification of critical cells is unchanged by composition with any strictly monotone function $\mathbb{R} \to \mathbb{R}$. In many cases it is often easier to construct the relationships between the cells, rather than the discrete Morse function itself. This hints towards the construction of what are known as discrete vector fields. We will use this approach during the exposition of this and all future sections.

To begin, we first must place certain light restrictions on the spaces we will be working with.

Definition 2.1 Let X be a CW complex. A *cell* of X will always refer to an open cell in X. Given a cell σ of dimension i, we will often write $\sigma^{(i)}$ to indicate that σ has dimension i. We will write $\mathcal{K}$ to denote the set of cells of X, and $\mathcal{K}_i$ to denote the set of i–cells of X.

A cell $\tau^{(i)} \subseteq \overline{\sigma^{(i+1)}}$ is said to be a *regular face* of a cell $\sigma^{(i+1)}$ if, given a characteristic map $\Phi_\sigma \colon D^{i+1} \to X$ for $\sigma^{(i+1)}$, $\Phi_\sigma^{-1}(\overline{\tau})$ is a closed ball and the map $\Phi_\sigma|_{\Phi_\sigma^{-1}(\overline{\tau})}$ is a homeomorphism. We say that the complex X is *regular* if $\tau^{(i)}$ is a regular face of $\sigma^{(i+1)}$ for any pair of cells $\tau^{(i)} \subseteq \overline{\sigma^{(i+1)}}$. Equivalently, X is regular if and only if the attaching map of each of its cells is a homeomorphism.

We will assume throughout most of our exposition that X is a regular CW complex. In his original paper on discrete Morse theory, Forman [10] proves some of his results without the requirement that X be regular. The spaces $\mathrm{UConf}_n(G)$, for any graph G, are actually cubical complexes, which are certainly regular CW complexes. Therefore, the condition that X be regular is not restrictive for what we need.

Definition 2.2 Let X be a regular CW complex. A *discrete vector field* V on X is a collection of partially defined functions $V_i \colon \mathcal{K}_i \to \mathcal{K}_{i+1}$ satisfying the following three conditions for each i:

(1) V_i is injective.

(2) The image of V_i is disjoint from the domain of V_{i+1}.

(3) $\sigma^{(i)}$ is a face of $V_i(\sigma^{(i)})$ for any $\sigma^{(i)}$ in the domain of V_i.

Given a regular CW complex X equipped with a discrete vector field V, a *cellular path* between two cells $\alpha^{(i)}$ and $\beta^{(i)}$ is a finite sequence of i–cells

$$\alpha^{(i)} = \alpha_0^{(i)}, \ \alpha_1^{(i)}, \dots, \ \alpha_{l-1}^{(i)}, \ \alpha_l^{(i)} = \beta^{(i)}$$

such that $\alpha^{(i)}_{j+1}$ is a face of $V_i(\alpha^{(i)}_j)$. We say that the path is *closed* if $\alpha^{(i)} = \beta^{(i)}$ and we say it is *trivial* if $\alpha^{(i)}_j = \alpha^{(i)}_k$ for all j and k.

A discrete vector field V is said to be a *discrete gradient vector field* if it admits no nontrivial closed cellular paths.

If V is a discrete gradient vector field on a regular CW complex X, then we call a cell σ of X *redundant* if σ is in the domain of V_i for some i, *collapsible* if it is in the image of V_i for some i, and *critical* otherwise.

Proposition 2.3 [8, Proposition 2.2; 10, Theorem 3.4] *Let X be a regular CW complex equipped with a discrete gradient vector field V. Consider the filtration*

$$\varnothing = X''_0 \subseteq X'_0 \subseteq X''_1 \subseteq X'_1 \subseteq \cdots \subseteq X''_n \subseteq X'_n \subseteq \cdots,$$

where X'_i is the i–skeleton of X with the redundant i–cells removed, and X''_i is the i–skeleton of X with both the redundant and critical i–cells removed. Then:

(1) *For any i, X'_i is obtained from X''_i by attaching m_i i–cells to X''_n along their boundaries, where m_i is the number of critical i–cells of the discrete gradient vector field V.*

(2) *For any i, X''_{i+1} deformation retracts onto X'_i.*

The above proposition leads to one notable corollary, which we record now.

Corollary 2.4 [8, Proposition 2.3; 10, Corollary 3.5] *Let X be a regular CW complex equipped with a discrete gradient vector field V. Then X is homotopy equivalent to a CW complex with precisely m_i i–cells for each i, where m_i is the number of critical i–cells of V.*

Just as is the case with traditional Morse theory, the decomposition of the space X given by Corollary 2.4 can be used to compute the homology groups of X. For simplicity, we will state the construction for cubical complexes, although the general case is similar.

Definition 2.5 Let X be a cubical complex, and write $C_i(X)$ for the free abelian group of i–cells of X. If $c \cong I_1 \times \cdots \times I_i$ is an i–cell of X, with $I_j \cong [0, 1]$, then we define

$$c^j_\tau = I_1 \times \cdots \times I_{j-1} \times \{0\} \times I_{j+1} \times \cdots \times I_i, \quad c^j_\iota = I_1 \times \cdots \times I_{j-1} \times \{1\} \times I_{j+1} \times \cdots \times I_i.$$

This allows us to define a boundary morphism

$$\partial \colon C_i(X) \to C_{i-1}(X)$$

given by

$$\partial(c) := \sum_j c_\tau^j - c_i^j,$$

turning $C_\bullet(X)$ into a chain complex. It is a well-known fact that the homology of this chain complex is the usual homology of the space X.

Further assume that X is equipped with a discrete gradient vector field V. Then we have a map $R \colon C_i(X) \to C_i(X)$ defined by

$$R(c) = \begin{cases} 0 & \text{if } c \text{ is collapsible,} \\ c & \text{if } c \text{ is critical,} \\ \pm \partial(V_i(c)) + c & \text{otherwise,} \end{cases}$$

where the sign of $\partial(V_i(c))$ in the above definition is chosen so that c has a negative coefficient. The property that V has no nontrivial closed paths implies that $R^m(c) = R^{m+1}(c)$ for all $m \gg 0$ and all i–cells c [10]. We set $R^\infty(c)$ to be this stable value.

For each i, let $\mathcal{M}_i$ denote the free abelian group with basis indexed by the critical i–cells of V. Then the *Morse complex* associated to V is defined to be

$$\mathcal{M}_\bullet \colon \cdots \to \mathcal{M}_n \to \cdots \to \mathcal{M}_1 \xrightarrow{\tilde{\partial}} \mathcal{M}_0 \to 0,$$

where the boundary map $\tilde{\partial}$ is given by

$$\tilde{\partial}(c) := R^\infty(\partial(c))$$

The map $\tilde{\partial}$ is known as the *Morse differential*.

Theorem 2.6 [10, Theorem 8.2; 11, Theorem 7.3] *For all i there are isomorphisms*

$$H_i(X) \cong H_i(\mathcal{M}_\bullet).$$

2.2 The configuration spaces of graphs

In this section we review necessary facts about configuration spaces of graphs. In the next section, we will explain how the techniques of discrete Morse theory apply to these spaces.

Definition 2.7 A *graph* is any compact, connected CW complex of dimension one. A *tree* is a topologically contractible graph. Given a graph G with vertex v, we will write $\mu(v)$ to denote the degree of the vertex v. Note that any loop on v contributes 2 to its degree count. We say that v is *essential* if $\mu(v) \geq 3$. If $\mu(v) = 1$, we say that v is a *boundary* vertex, and the unique edge connected to v is called a *boundary edge*.

The *configuration space on n points* of G is the topological space

$$\operatorname{Conf}_n(G) := \{(x_1, \ldots, x_n) \in G^n \mid x_i \neq x_j \text{ if } i \neq j\}.$$

We note that there is a natural action without fixed points on $\operatorname{Conf}_n(G)$ by the symmetric group $\mathfrak{S}_n$. The *unordered configuration space on n points* of G is the quotient space

$$\operatorname{UConf}_n(G) := \operatorname{Conf}_n(G)/\mathfrak{S}_n.$$

For the majority of this paper, we will work with the spaces $\operatorname{UConf}_n(G)$. Note that many of the structural theorems discussed in this section will apply to both spaces.

In order to apply discrete Morse theory to questions about these configuration spaces, we will first need to place a CW complex structure on them. To accomplish this, we use a theorem of Abrams [1].

Definition 2.8 The *discretized configuration space* $D_n(G)$ *in n points* over G is the CW subcomplex of G^n spanned by cells of the form

$$\sigma_1 \times \cdots \times \sigma_n,$$

where $\sigma_i \subseteq G$ is a cell (ie an edge or vertex) of G and $\overline{\sigma}_i \cap \overline{\sigma}_j = \varnothing$ whenever $i \neq j$. We write $UD_n(G)$ to denote the quotient of $D_n(G)$ by the action of the symmetric group.

Theorem 2.9 [1, Theorem 2.1; 14, Theorem 2.4; 20] *Let G be a graph, and assume that G satisfies the following two properties:*

(1) *Each path connecting distinct vertices of degree $\neq 2$ has length at least $n - 1$.*

(2) *Each homotopically essential path connecting a vertex to itself has length at least $n + 1$.*

Then the inclusions $D_n(G) \hookrightarrow \operatorname{Conf}_n(G)$ and $UD_n(G) \hookrightarrow \operatorname{UConf}_n(G)$ are homotopy equivalences.

Remark 2.10 In the first cited source, Abrams states the theorem assuming that each path connecting distinct vertices of degree $\neq 2$ has length at least $n + 1$. It is noted after the proof that the version of the theorem stated above is true, and a brief argument is given for how it is proven. In the second source, Kim, Ko and Park give a formal argument for this improvement. In the third source, Prue and Scrimshaw provide a discrete Morse theory argument, which is independent of the first two sources. For our purposes, the exact number of vertices needed is unimportant, as we can always just subdivide the edges of G more if needed.

Note that subdividing edges of a graph G does not impact the configuration spaces $\mathrm{Conf}_n(G)$ and $\mathrm{UConf}_n(G)$. For much of what follows, we will often just assume, without explicit mention, G is subdivided enough that the homotopy equivalence of Theorem 2.9 holds.

Theorem 2.9 implies that $\mathrm{Conf}_n(G)$ and $\mathrm{UConf}_n(G)$ are homotopy equivalent to cubical complexes of dimension n. In fact, we will be able to do better than this.

Theorem 2.11 [13, Theorem 3.3; 22, Theorem 0.1] *Let G be a graph which is not homeomorphic to S^1. Then $\mathrm{Conf}_n(G)$ and $\mathrm{UConf}_n(G)$ are homotopy equivalent to CW complexes of dimension N_G, where N_G is the number of essential vertices of G.*

Remark 2.12 If G is homeomorphic to S^1, then $\mathrm{UConf}_n(G)$ is easily seen to be homotopy equivalent to a circle for all n. Throughout this work, and the literature in general, it is a recurring theme that certain theorems only apply to graphs which are neither S^1 nor the interval. It is interesting that these two graphs are precisely those which are homeomorphic to compact manifolds.

Note that Ghrist originally proved Theorem 2.11 using more classical topological means. Świątkowski accomplished something similar in [22]. We will later see that it naturally falls out of the discrete Morse structure that Farley and Sabalka placed on $UD_n(G)$. This was first noted by Farley and Sabalka [8].

One remarkable thing to note about the Ghrist–Świątkowski theorem is that the dimension of the configuration spaces of graphs is independent of n. This behavior is in stark contrast to the behavior of configuration spaces of smooth manifolds of dimension ≥ 2.

In his paper [13], Ghrist also proves that configuration spaces of graphs are, in fact, aspherical. This result was later reproven by Abrams [1, Theorem 3.10], where he shows that both $D_n(G)$ and $UD_n(G)$ are universally covered by CAT(0) complexes.

Theorem 2.13 [13, Theorem 3.1; 1, Corollary 3.11] *Let G be a graph. Then $\mathrm{Conf}_n(G)$, and hence $\mathrm{UConf}_n(G)$, are aspherical. That is to say, $\pi_k(\mathrm{Conf}_n(G)) = 0$ for $k > 1$.*

Note that this theorem is analogous to that which says the configuration spaces of the plane are aspherical. In that case, the fundamental groups of the ordered and unordered spaces are the Artin pure braid groups and the Artin braid groups, respectively. We borrow this terminology for our context as well.

Definition 2.14 Let G be a graph. The *braid group on n strands* of G is defined to be

$$B_n G := \pi_1(\mathrm{UConf}_n(G)).$$

We similarly define the *pure braid group on n strands* as

$$P_n G := \pi_1(\mathrm{Conf}_n(G)).$$

The study of the braid groups of graphs is still a very active area of research. See [14; 15; 9] for more on these groups.

As an immediate corollary to Theorem 2.13, we obtain the following:

Corollary 2.15 *Let G be a graph. Then there are isomorphisms*

$$H_i(\mathrm{Conf}_n(G)) \cong H_i(P_n G), \quad H_i(\mathrm{UConf}_n(G)) \cong H_i(B_n G).$$

The goal for this paper is to establish a methodology for understanding stability phenomena in the groups $H_i(B_n G)$, in the spirit of modern trends of asymptotic algebra. Note that we will spend very little time considering pure braid groups. In fact, there doesn't seem to be much literature about the homology of these groups, as they are vastly more complicated than the braid groups [15; 3].

The following theorem of Farley exactly computes the homology groups when G is a tree:

Theorem 2.16 [7] *Let G be a tree. Then the group $H_i(B_n G)$ is free for all $n \geq 0$.*

Note that the theorem of [7] is slightly more general than this. We will discuss the more general version in later sections.

Expanding upon the work of Farley, the following theorem of Ko and Park suggests that the tree case is indicative of a more general phenomenon:

Theorem 2.17 [15, Theorem 3.6] *Let G be a graph. Then G is planar if and only if $H_1(B_nG)$ is torsionfree for any, and therefore all, $n \geq 2$. Moreover, in the case where G is not planar, all torsion is 2–torsion.*

Kim, Ko and Park [14, Theorem 5.5] proved that G is planar if and only if $H_1(B_2G, \mathbb{Z})$ is torsionfree. It is also conjectured in that work that Theorem 2.17 is true. It should be noted that [15, Theorem 3.6] is far stronger than what we have written above. In fact, their result explicitly computes the groups $H_1(B_nG, \mathbb{Z})$ in terms of combinatorial invariants of the graph G. It follows from their computation that the amount of 2–torsion of the group is unvarying in n. Moreover, if G is biconnected — that is, G requires the removal of at least 2 vertices to disconnect it — then $H_1(B_nG, \mathbb{Z}) \cong H_1(B_2G, \mathbb{Z})$ for all $n \geq 2$ [15, Lemma 3.12]. In the final section of this paper we will conjecture an extension of this fact to the higher homologies (see Conjecture 4.3).

To finish this section, we record a result of Gal on the Euler characteristic of these spaces. Note that Gal proves a more general theorem for computing the Euler characteristic of the configuration space of any simplicial complex.

Theorem 2.18 [12, Theorem 2] *Let G be a graph and set*

$$\mathfrak{e}(t) := \sum_{n \geq 0} \frac{\chi(\mathrm{Conf}_n(G))}{n!} t^n.$$

Then

$$\mathfrak{e}(t) = \frac{\prod_v \left(1 + (1 - \mu(v))t\right)}{(1-t)^E},$$

where E is the number of edges of G.

Remark 2.19 $\mathrm{Conf}_n(G)$ is an $n!$–fold cover of $\mathrm{UConf}_n(G)$. It follows from this that the above formula can be easily used to compute the Euler characteristic of $\mathrm{UConf}_n(G)$ as well.

Definition 2.20 Let G be a graph which has at least one essential vertex, and let $\mathcal{E}$ denote the set of essential vertices of G. Then an *essential edge* of G is a connected component of $G - \mathcal{E}$.

Corollary 2.21 *Let G be a graph with at least one essential vertex and let E denote the number of essential edges of G. Then the function*

$$n \mapsto \chi(\mathrm{UConf}_n(G))$$

is a polynomial of degree $E - 1$.

Proof We first note that smoothing the degree 2 vertices of G does not impact the spaces $\mathrm{UConf}_n(G)$. We may therefore assume without loss of generality that G does not have any such vertices. In this case, the Euler characteristic $\chi(\mathrm{UConf}_n(G))$ is the coefficient of t^n in the power series expansion of

$$\frac{\prod_{v \text{ essential}}\left(1 + (1 - \mu(v))t\right)}{(1-t)^E}.$$

A straightforward enumerative combinatorics argument implies that the n^{th} coefficient of the power series expansion of this expression is given by a polynomial of degree exactly $E - 1$ for n sufficiently large. It remains to show that this agreement begins at $n = 0$.

Assume that G has N_G essential vertices. It is a standard fact from enumerative combinatorics that the coefficients of the power series expansion of a rational function of the form $f(t)/(1-t)^l$ will agree with a polynomial for all n so long as $\deg(f) < l$. This is the case in our specific instance, so long as $N_G > 1$. □

Theorem 2.18 was largely the inspiration for this work. It suggests that asymptotically the Betti numbers of $H_i(B_nG)$ are polynomial. In this work we will prove this suggestion in the case where G is a tree. We will also provide a setup in the case where G is a general graph, which hopefully will be able to illustrate this behavior in that case as well.

2.3 The discrete Morse theory of $\mathrm{UConf}_n(G)$

In this section we will outline the discrete Morse structure on $\mathrm{UConf}_n(G)$ developed in the work of Farley and Sabalka [8]. We begin by fixing n, and assuming the graph G is sufficiently subdivided for Theorem 2.9 to hold for $\mathrm{UConf}_n(G)$.

Fix a spanning tree T for G, as well as an embedding of T into the plane. We will label the vertices of T by applying a depth-first search. Concretely, begin by choosing a boundary vertex v_0 to be the root of T, and label it with the number 0. Continue down the boundary edge adjacent to v_0, labeling any vertices encountered along the way with increasing labels. If at any point an essential vertex is encountered, then travel down the leftmost — relative to the current direction of travel — edge whose vertices have not been labeled. If a boundary vertex is encountered, then one returns to the most recently passed essential vertex. An example of a correctly labeled tree is given in Figure 1.
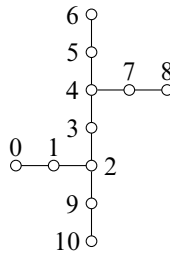
Figure 1: A tree which is properly labeled. Note that this tree is also sufficiently subdivided to apply Theorem 2.9 with $n = 3$.

For the remainder of this section, G and T will be as in the previous paragraphs. As before we will write $\mathcal{K}$ to denote the set of cells of $\mathrm{UConf}_n(G)$, while $\mathcal{K}_i$ will denote the set of i–cells of $\mathrm{UConf}_n(G)$.

Definition 2.22 Given an edge e of G, we write $\iota(e)$ to denote its largest — with respect to the given labeling — vertex, and $\tau(e)$ to denote its smallest vertex. If v is a vertex of G which is not v_0, then we write $e(v)$ to denote the unique edge of T for which $\iota(e(v)) = v$. We note that $e(v)$ will be the first edge on the unique path within T from v to v_0.

Let $c = \{v_{j_1}, \ldots, v_{j_{n-i}}, e_{r_1}, \ldots, e_{r_i}\}$ be an i–cell of $UD_n(G)$. If

$$\tau(e(v_{j_l})) \cap \{v_{j_1}, \ldots, v_{j_{n-i}}, \iota(e_{r_1}), \tau(e_{r_1}), \ldots, \iota(e_{r_i}), \tau(e_{r_i})\} \neq \varnothing$$

for some $1 \leq l \leq n-i$, then we say that v_{j_l} is *blocked* in c. Note that by convention v_0 is always blocked.

We define a partial function $V_i \colon \mathcal{K}_i \to \mathcal{K}_{i+1}$ inductively in the following way. Given an i–cell $c = \{v_{j_1}, \ldots, v_{j_{n-i}}, e_{r_1}, \ldots, e_{r_i}\}$, assume without loss of generality that v_{j_1} is an unblocked vertex of c which is smallest with respect to the labeling of T. Assuming that c is not in the image of V_{i-1}, we set

$$V_i(\{v_{j_1}, \ldots, v_{j_{n-i}}, e_{r_1}, \ldots, e_{r_i}\}) := \{v_{j_2}, \ldots, v_{j_{n-i}}, e_{r_1}, \ldots, e_{r_i}, e(v_{j_1})\}.$$

If c does not have any unblocked vertices, or if c is in the image of V_{i-1}, then $V_i(c)$ is undefined.

Theorem 2.23 [8, Sections 3.1–3.2] *The collection of partial functions $V_i \colon \mathcal{K}_i \to \mathcal{K}_{i+1}$ forms a discrete gradient vector field on $UD_n(G)$.*
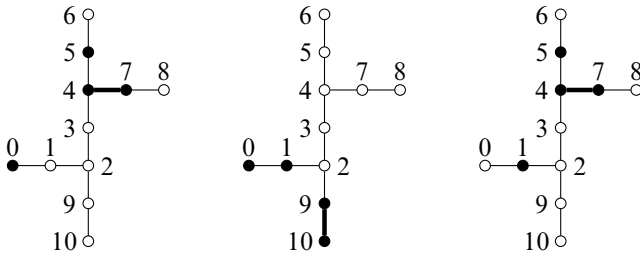
Figure 2: The critical 1–cell $\{v_0, v_5, e_{4,7}\}$, the collapsible 1–cell $\{v_0, v_1, e_{9,10}\}$ and the redundant 1–cell $\{v_1, v_5, e_{4,7}\}$

While the spaces $\mathrm{UConf}_n(G)$ can be high-dimensional, we can still visualize their cells. Indeed, it is often useful to think of the cells of $\mathrm{UConf}_n(G)$ as being subsets of edges and vertices of the graph G itself. In Figure 2 we see examples of a critical, a collapsible, and a redundant 1–cell of V, where G is taken to be the tree in Figure 1 and $n = 3$. In these examples, the cell $\{v_0, v_1, e_{9,10}\}$ is collapsible, as

$$\{v_0, v_1, e_{9,10}\} = V_0(\{v_0, v_1, v_{10}\}),$$

and v_{10} is the smallest unblocked vertex. The cell $\{v_1, v_5, e_{4,7}\}$ is redundant, as v_1 is not blocked, and it is not in the image of V_0. Indeed, the only cell that could possibly map to it would be $\{v_1, v_5, v_7\}$. However, in this case v_1 is the smallest unblocked vertex. Finally, the cell $\{v_0, v_5, e_{4,7}\}$ is critical, as its vertices are all blocked, and it is not collapsible.

From these examples, one can imagine there existing a way to classify all cells of the three types. To write down such a classification, we first need a bit of nomenclature.

Definition 2.24 Let $c = \{v_{j_1}, \ldots, v_{j_{n-i}}, e_{r_1}, \ldots, e_{r_i}\}$ be an i–cell of $UD_n(G)$. We say that an edge $e \in c$ is *order-respecting in* c if e is in T and, for all $v \in c$ such that $\tau(e(v)) = \tau(e)$, the label of v in T is larger than the label of $\iota(e)$.

Remark 2.25 Edges $e \in G - T$ — which we call *deleted edges* — are, by definition, never order-respecting.

Theorem 2.26 (the classification theorem [8, Theorem 3.6]) *Let c be a cell of $UD_n(G)$ and let T be a choice of spanning tree of G equipped with a planar embedding. Then:*

(1) *c is critical if and only if it contains no order-respecting edges and all of its vertices are blocked.*

(2) *c is redundant if and only if*

    (a) *it contains no order-respecting edges and at least one of its vertices is un-blocked, or*

    (b) *it contains an order-respecting edge (and thus a minimal order-respecting edge e) and there is some unblocked vertex v such that the label of v in T is less than that of $\iota(e)$.*

(3) *c is collapsible if and only if it contains an order-respecting edge (and thus a minimal order-respecting edge e) and any vertex v such that the label of v in T is less than that of $\iota(e)$ is blocked.*

It is not hard to check that the previous examples satisfy the conditions of the classification theorem.

One useful fact about the discrete gradient vector field V is that its critical cells are, in some sense, the most restrictive of the three possible types. Indeed, we may expand upon the classification of critical cells in the following way:

Lemma 2.27 *Let $c = \{v_{j_1}, \dots, v_{j_{n-i}}, e_{r_1}, \dots, e_{r_i}\}$ be a critical i–cell of V.*

(1) *For each k in $\{1, \dots, i\}$, either e_{r_k} is a deleted edge or $\tau(e_{r_k})$ is an essential vertex of T.*

(2) *Let c' be a critical i–cell with edges $e_{r_1}, \dots, e_{r_i}$ such that the number of its vertices in each component of $T - \{\overline{e}_{r_1}, \dots, \overline{e}_{r_i}\}$ agrees with the number of vertices of c in these same components. Then $c = c'$.*

Proof For the first statement, let e be an edge of c which is not deleted. If $\tau(e) = v_0$, then it is clear that e is order-respecting. We note that $\tau(e)$ is not a boundary vertex in any other case by how the labeling on T was defined. It therefore remains to show that $\tau(e)$ does not have degree 2. Indeed, if this were the case, then for any vertex v of G, $\tau(e(v)) \neq \tau(e)$. In particular, e must be order-preserving. This proves the first claim.

For the second claim, one uses the fact that the vertices of c must all be blocked. □

Note that the observations made in the above lemma were originally pointed out by Farley and Sabalka when they defined the vector field V [8, Section 3.3]. We collected these observations in Lemma 2.27 so that they could be easily referred back to during future sections. We will find that this lemma is critical in defining the polynomial ring

structure on $\bigoplus_n H_i(B_nG)$. In the next section, we will find that the deleted edges of G can be chosen so that they remain unchanged as n increases and we repeatedly subdivide the graph. The above lemma also suggests that those edges of critical cells contained in the tree T are unchanging in n. While this is not literally the case, as all edges of T are being subdivided as n increases, Lemma 2.27 tells us that the important information encoded by such an edge is the essential vertex of its tail, as well as the direction it is leaving this essential vertex. For this reason we will often be a bit loose with our language and claim that two critical cells with different choices of n have the "same" edges.

Definition 2.28 Let c be a critical cell of the discrete gradient vector field and let $e \in T$ be an edge in c. The classification theorem implies that there must exist some vertex $v \in c$ such that the label of v in T is smaller than that of $\iota(e)$, while $\tau(e(v)) = \tau(e)$. We refer to such a vertex as a *witness* for e in c. If v is the only witness for e in c, then we say v is *necessary*.

3   The proof of Theorem A

3.1   The setup

We use this section to fix the notation which will be used throughout the proofs of the main theorems. Note that all of the constructions presented in the previous sections were already well established in the literature prior to this work. To the knowledge of the author, the main construction of the following sections — namely the action of S_G on the critical cells of $\mathrm{UConf}_n(G)$ — does not appear elsewhere in the literature.

As was the case in all previous sections, let G denote a graph, which is neither a circle nor a line segment. We will reserve E to denote the number of essential edges of G, while N_G will denote the number of essential vertices of G.

Next, we must construct our spanning tree T, as in Section 2.3. To do this, we first subdivide every edge of G, which connects two essential vertices, once. Note that this includes any loops of G. Once this is done, we choose a spanning tree T of G which satisfies the following:

(3-1)   Every edge of G not in T is adjacent to an essential vertex of G.

Note that it is not entirely obvious that such a T exists for an arbitrary graph G. This fact was proven by Farley and Sabalka during the proof of [8, Theorem 4.4].

Remark 3.1 It is noted by Farley and Sabalka that if one chooses T to satisfy (3-1), then Lemma 2.27 and Corollary 2.4 imply Theorem 2.11.

Having chosen our spanning tree T to satisfy (3-1), we observe that for $n \geq 2$, we can sufficiently subdivide G for Theorem 2.9 by subdividing T. This follows from the fact that we began by subdividing all edges of G which connected essential vertices. We will often not differentiate between the spanning trees chosen for each n for this very reason.

3.2 The modules $\mathcal{M}_{i,\bullet}$

Recall from Section 2.1 the Morse complex

$$\cdots \to \mathcal{M}_i \to \cdots \to \mathcal{M}_1 \to \mathcal{M}_0 \to 0,$$

where $\mathcal{M}_i$ is the free $\mathbb{Z}$–module with basis indexed by critical cells. In our case, we write $\mathcal{M}_{i,n}$ to denote the free $\mathbb{Z}$–module with basis indexed by the critical cells of the Farley–Sabalka discrete gradient vector field of $UD_n(G)$.

Let $\mathfrak{e}_1, \ldots, \mathfrak{e}_E$ denote the essential edges of the graph G. Then we set

$$S_G := \mathbb{Z}[x_{\mathfrak{e}_1}, \ldots, x_{\mathfrak{e}_E}].$$

Our goal for the remainder of this section will be to argue, for each i, that $\mathcal{M}_{i,\bullet}$ has the structure of a finitely generated graded S_G–module.

Definition 3.2 Fix i and n, and let $\mathfrak{e}$ denote an essential edge of G. Given a critical i–cell

$$c = \{v_{j_1}, \ldots, v_{j_{n-i}}, e_{r_1}, \ldots, e_{r_i}\}$$

of $UD_n(G)$, we define $x_{\mathfrak{e}} \cdot c$ to be the unique critical i–cell of $UD_n(G)$ obtained from c by adding a vertex to the connected component of $T - \{\bar{e}_{r_1}, \ldots, \bar{e}_{r_i}\}$ containing $\mathfrak{e}$. More precisely, $x_{\mathfrak{e}} \cdot c$ is obtained from c by adding the smallest vertex on the connected component containing $\mathfrak{e}$. This is well defined by Lemma 2.27, as well as the choice of the tree T, as T will contain a unique representative of each essential edge of G by construction. This turns the collection $\{\mathcal{M}_{i,n}\}_n$ into a graded S_G–module, which we denote by $\mathcal{M}_{i,\bullet}$.

One can visualize the action of S_G described in the above definition in the following way. The labeling on T induces a natural flow on T. Namely, all edges flow towards
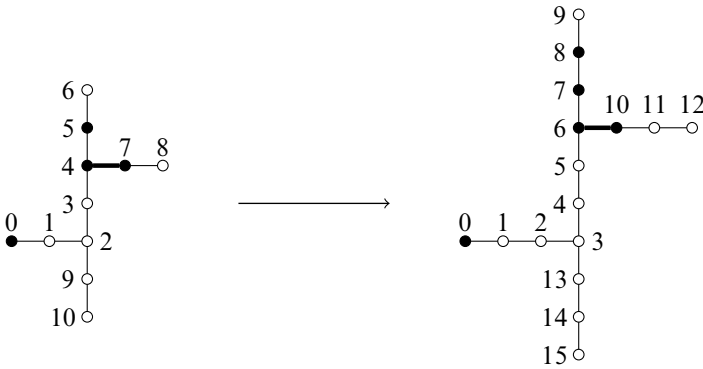
Figure 3: An illustration of multiplication by $x_{\mathfrak{e}}$, where $\mathfrak{e}$ is the essential edge containing the boundary vertex v_6 of the tree G of Figure 1.

the root of T. Let $c = \{v_{j_1}, \ldots, v_{j_{n-i}}, e_{r_1}, \ldots, e_{r_i}\}$ be a critical cell. We imagine the vertices v_j as being particles drifting in the direction of the flow, while the edges e_j, along with their endpoints, are stationary blockades. For any essential edge $\mathfrak{e}$, the action of $x_{\mathfrak{e}}$ on c involves placing a new particle somewhere on $\mathfrak{e}$, and allowing it to flow until it too is blocked.

We provide an illustration of the action of S_G in Figure 3.

Lemma 3.3 *For each i, the module $\mathcal{M}_{i,\bullet}$ is finitely generated over S_G.*

Proof We claim that every critical cell in $\mathcal{M}_{i,n}$, with $n > 2i$, can be obtained from a critical cell in $\mathcal{M}_{i,n-1}$. Indeed, if c is a critical cell in $\mathcal{M}_{i,n}$, then at least one vertex is not a necessary witness of some edge in e. Remove the vertex, among those which are not necessary witness vertices, which occupies the maximal position with respect to the labeling on T. Removing this vertex leaves us with a critical cell c' in $\mathcal{M}_{i,n-1}$, as this removal cannot create unblocked vertices nor order-respecting edges. It follows that c is the image of c' under the action of the appropriate essential edge.

Because $\mathcal{M}_{i,n}$ has a finite basis for each n, it follows that $\mathcal{M}_{i,\bullet}$ is finitely generated. $\square$

Remark 3.4 It follows from the above lemma that the number of critical i–cells grows, as a function of n, like a polynomial of degree $\leq E - 1$. Indeed, if $V_{\bullet}$ is a graded module over a polynomial ring in d variables, then the function $n \mapsto \operatorname{rank}_{\mathbb{Z}} V_n$ will always grow like a polynomial of degree $\leq d$. In fact, if $i < N_G$, we claim that this polynomial must have degree strictly less than $E - 1$. First, we note that the action

of S_G on a critical cell c does not affect the edges of c. Because of this it follows that $\mathcal{M}_{i,\bullet}$ can be expressed as a direct sum of graded S_G–modules, where each summand corresponds to a choice of edges in our critical i–cell. For any fixed critical i–cell c, one observes that two variables $x_{\mathfrak{e}}$ and $x_{\mathfrak{e}'}$ will act identically whenever $\mathfrak{e}$ and $\mathfrak{e}'$ are on the same connected component of G with the edges of c removed. It follows that the summand of $\mathcal{M}_{i,\bullet}$ containing this cell has Hilbert polynomial of degree strictly less than E whenever $i < N_G$. The Hilbert polynomial of the whole of $\mathcal{M}_{i,\bullet}$ is a sum of such polynomials, and therefore also has degree strictly less than E. This will be used in the final section of the paper.

Lemma 3.3, and Theorem 2.6 now imply some asymptotic data about the homology groups $H_i(B_n G)$.

Theorem 3.5 *Let G be a graph with E essential edges. Then for all $i \geq 0$, there exists a polynomial $P_i \in \mathbb{Q}[t]$ of degree $\leq E - 1$ such that, for all $n \gg i$,*

$$\dim_{\mathbb{Q}}(H_i(B_n G; \mathbb{Q})) \leq P_i(n).$$

If $i < N_G$, then this polynomial has degree $< E - 1$.

3.3 The case of trees

In this section, we begin to explore the specific case where G is a tree. The work of Farley [7] will allow us to conclude quite a bit more than we are able to in the general case. To begin with, we state a refined version of Theorem 2.16.

Theorem 3.6 [7] *Let G be a tree. Then for each n the Morse complex $\mathcal{M}_{i,n}$ has trivial differential. In particular,*

$$\mathcal{H}_i \cong \mathcal{M}_{i,\bullet}$$

as S_G–modules, where $\mathcal{H}_i$ is the module of Theorem C.

To state our main stability theorem, we begin with the following definition:

Definition 3.7 Let G be a tree. Then we define the quantity

$$\Delta_G^i := \max_{\{\{v_{j_1},\dots,v_{j_i}\}\,|\,v_{j_k}\text{ essential}\}} \left\{\dim_{\mathbb{Q}}(H_0(G - \{v_{j_1},\dots,v_{j_i}\}; \mathbb{Q}))\right\}.$$

For example, Δ_G^1 is the maximum degree of G, while $\Delta_G^{N_G}$ is E, the number of essential edges. By convention, $\Delta_G^0 = 1$ while $\Delta_G^i = 0$ for $i > N_G$.

Theorem 3.8 *Let G be a tree, and fix $i \geq 0$. Then the Betti number $b_i(n) := \dim_{\mathbb{Q}}\big(H_i(B_n(G); \mathbb{Q})\big)$ is equal to a polynomial of degree $\Delta_G^i - 1$ for $n \gg i$.*

Proof Theorem 3.6 and the work of the previous sections imply that $b_i(n)$ is eventually a polynomial; it remains only to show that this polynomial has the claimed degree.

We begin by partitioning the basis vectors of $\mathcal{M}_{i,n}$ according to the collection of edges which appear. The action of S_G does not impact the edges, and so each of these partitions will correspond to a summand of $\mathcal{M}_{i,n}$. Moreover, considering any one of these summands, we observe that the variables induce at most Δ_G^i distinct operators on $\mathcal{M}_{i,n}$. It follows that the degree of the Hilbert polynomial of $\mathcal{M}_{i,\bullet}$ is at most $\Delta_G^i - 1$. To finish the proof, we will show that $\mathcal{M}_{i,\bullet}$ has a summand whose Hilbert polynomial has degree $\Delta_G^i - 1$.

Fix the i essential vertices of G whose absence realizes the maximum in Definition 3.7. For each essential vertex in this fixed list, let e be an edge adjacent to it in the second leftmost direction, relative to the direction of the root. Note that any critical cell containing these edges has a unique essential edge which can house a witness vertex. There is therefore a critical i–cell in $\mathcal{M}_{i,2i}$ constructed by choosing each of these edges, along with the unique witness vertex for each. Denote this cell by c. The submodule generated by c is a summand of $\mathcal{M}_{i,\bullet}$. We claim that this submodule is our desired summand.

Any i–cell of $\mathcal{M}_{i,n}$ which shares the same edges of c — ie those basis vectors in the n^{th} graded piece of $S_G \cdot c$ — is entirely determined by distributing $n - 2i$ vertices to one of Δ_G^i components. Thus,

$$\mathrm{rank}_{\mathbb{Z}}((S_G \cdot c)_n) = \binom{n - 2i + \Delta_G^i - 1}{\Delta_G^i - 1},$$

as desired. $\square$

The only piece of Theorem A which remains to be proven is in showing that $b_i(n)$ agrees with a polynomial for all $n \geq 0$. To accomplish this, we must first prove Theorem D. As a first simplification, we note that our answer is unaffected by changing basis to $\mathbb{Q}$, and may therefore assume we are working over a rational polynomial ring. This grants us access to many classical approaches to solving the problem. We will also make heavy use of the direct sum decomposition of $\mathcal{M}_{i,\bullet}$ described in the proof of Theorem 3.8.

Definition 3.9 Let G be a tree. The notation

$$\bar{v} := (v_{j_1}, \dots, v_{j_i})$$

will always refer to an i–tuple of essential vertices of G, appearing in the order induced by the labeling on G. We write $\mu(\bar{v})$ to denote the number of connected components of $G - \{v_{j_1}, \dots, v_{j_i}\}$. Given $\bar{v}$, let $\bar{l} := (l_1, \dots, l_i)$ denote an i–tuple of positive integers such that

$$1 \le l_k \le \mu(v_{j_k}) - 2 \quad \text{for all } k.$$

We will also set

$$|\bar{l}| := \sum_k l_k.$$

Given a pair $(\bar{v}, \bar{l})$, we associate to it the summand of $\mathcal{M}_{i, \bullet} \otimes_{\mathbb{Z}} \mathbb{Q}$ generated in degree $2i$ by critical cells c satisfying the following two properties:

(1) If $e \in c$ is an edge, then $\tau(e) = v_{j_k}$ for some k.

(2) if $e \in c$ has $\tau(e) = v_{j_k}$, then there are precisely l_k essential edges which can house witness vertices of e.

Note that these two pieces of data uniquely determine the edges which are allowed to appear in cells which form the bases for the graded pieces of the summand. Denote this summand by $N^{(\bar{v}, \bar{l})}$.

We observe that two variables $x_{\mathfrak{e}_1}$ and $x_{\mathfrak{e}_2}$ of S_G act identically on $N^{(\bar{v}, \bar{l})}$ if and only if $\mathfrak{e}_1$ and $\mathfrak{e}_2$ are contained in the same connected component of $G - \{v_{j_1}, \dots, v_{j_i}\}$. Let $S^{(\bar{v}, \bar{l})}$ denote the quotient of $S_G \otimes_{\mathbb{Z}} \mathbb{Q}$ by $x_{\mathfrak{e}_2} - x_{\mathfrak{e}_2}$. In particular, $S^{(\bar{v}, \bar{l})}$ is isomorphic to a rational polynomial ring in $\mu(\bar{v})$ variables, and $N^{(\bar{v}, \bar{l})}$ is a finitely generated graded module over $S^{(\bar{v}, \bar{l})}$.

Remark 3.10 If $i = 0$, then $\bar{v}$ and $\bar{l}$ are empty tuples with $\mu(\bar{v}) = 1$. If $i > N_G$, then, by convention, $N^{(\bar{v}, \bar{l})} = 0$.

Our next goal will be to show that the Hilbert function

$$n \mapsto \dim_{\mathbb{Q}}(N^{(\bar{v}, \bar{l})})$$

agrees with a polynomial for all n. Note that we already know that this Hilbert function is a polynomial for n sufficiently large, and must only show that this agreement is the case for all n. This will imply the same about the Hilbert function for $\mathcal{M}_{i, \bullet}$. We

accomplish this through the computations in the next section, although we spend time now to set some ground work in this direction. Following this, we will explicitly compute this Hilbert polynomial. We begin with the following:

Proposition 3.11 *The module $N^{(\bar{v},\bar{l})}(i)$ is isomorphic to a squarefree monomial ideal of $S^{(\bar{v},\bar{l})}$, where for a graded module V over a polynomial ring, $V(i)$ denotes the graded module with $V(i)_n = V_{n+i}$.*

Proof We recall that a monomial ideal is called squarefree if it is generated by monomials, none of which are divisible by a square.

By definition, $N^{(\bar{v},\bar{l})}(i)_i = N_{2i}^{(\bar{v},\bar{l})}$ is spanned by critical cells containing the edges $e_1, \dots, e_i$ determined by the pair $(\bar{v}, \bar{l})$ (see Definition 3.9), along with a single witness vertex for each. To define a map $\phi \colon N^{(\bar{v},\bar{l})}(i) \to S^{(\bar{v},\bar{l})}$, it suffices to specify where each of these cells is mapped. We set, for any such cell c,

$$(3\text{-}2) \qquad\qquad \phi(c) = \prod_{\mathfrak{e}} x_{\mathfrak{e}},$$

where the product is over essential edges containing a witness vertex of c. Note that it is impossible for a single essential edge to house a witness vertex for multiple edges, and so the above monomial is indeed squarefree. Extending this map through the action of $S^{(\bar{v},\bar{l})}$ defines our desired isomorphism. $\qquad\qquad\square$

Squarefree monomial ideals are some of the most well-understood objects in commutative algebra. They are also the subject of the so-called Stanley–Reisner theory (see [19] for a comprehensive reference on the subject).

3.4 Computing the Hilbert polynomial

We note that for $i = 1$, the polynomial describing the Betti number $b_1(n)$ has been explicitly computed in terms of invariants of the tree G. Indeed, this follows from the structure theorem of Ko and Park [15, Theorem 3.6], and also from the work of Farley and Sabalka [8]. The Hilbert polynomial has also been computed for all i in the case where G is a tree with maximal degree 3 by Farley [7]. Aside from these cases, no explicit description of the polynomial is known. That being said, the work in this paper implies that computing this polynomial is no more difficult than computing the Hilbert polynomials of some squarefree monomial ideals. In other words, for any tree G, the work of this paper reduces the task of computing $H_i(B_n G)$ for all n to a finite computation. We will proceed with this computation in this section.

Remark 3.12 Proposition 3.11 reveals that the monomial ideals which appear as graded shifts of summands of $\mathcal{H}_i$ are quite simple. Because of this, we will find that the computation of the Hilbert polynomial of $\mathcal{H}_i$ falls more in the realm of enumerative combinatorics than Stanley–Reisner theory. It is likely that studying the cases of more general graphs will require more robust commutative algebra, and we therefore proceed with the case that G is a tree using that language. Our hope is that once the case of general graphs is completed, this will aid in creating a uniform means of approaching asymptotic behaviors in the configuration spaces of graphs.

In this section we will use the decomposition of Proposition 3.11 to compute the Hilbert polynomial of $\mathcal{M}_{i,\bullet}$. Proposition 3.11 implies that $N^{(\overline{v},\overline{l})}(i)$ is isomorphic to a squarefree monomial ideal of $S^{(\overline{v},\overline{l})}$. To ease computations, we reorder the variables of $S^{(\overline{v},\overline{l})}$ to satisfy the following. Begin by labeling the essential edges adjacent to v_{j_1} which house witness vertices by x_1 through x_{l_1}, chosen in the order induced by our ordering of the vertices of G. Next, label the essential edges housing a witness vertex adjacent to v_{j_2} with the variables x_{l_1+1} to $x_{l_1+l_2}$, using the same rule. Continue in this fashion until $|\overline{l}|$ essential edges have been labeled. We then use the remaining variables to label the final $\mu(\overline{v}) - |\overline{l}|$ essential edges, chosen in any order. Our first observation is the following:

Lemma 3.13 There is an isomorphism of $S^{(\overline{v},\overline{l})}$–modules

$$N^{(\overline{v},\overline{l})}(i) \cong (x_1,\ldots,x_{l_1}) \otimes_{\mathbb{Q}} (x_{l_1+1},\ldots,x_{l_1+l_2}) \otimes_{\mathbb{Q}} \cdots \otimes_{\mathbb{Q}} (x_{\sum_{j=1}^{i-1} l_j+1},\ldots,x_{|\overline{l}|})$$
$$\otimes_{\mathbb{Q}} \mathbb{Q}[x_{|\overline{l}|+1},\ldots,x_{\mu(\overline{v})}],$$

where the tuple $(x_1,\ldots,x_{l_1})$ is considered as a module over the ring $\mathbb{Q}[x_1,\ldots,x_{l_1}]$, $(x_{l_1+1},\ldots,x_{l_1+l_2})$ a module over the ring $\mathbb{Q}[x_{l_1+1},\ldots,x_{l_1+l_2}]$, and so on.

Proof This follows immediately from the isomorphism (3-2). $\qquad\square$

Definition 3.14 The *Hilbert series* associated to the module $N^{(\overline{v},\overline{l})}$ is the formal power series

$$H_{(\overline{v},\overline{l})}(t) := \sum_{n\geq 0} (\dim_{\mathbb{Q}} N_n^{(\overline{v},\overline{l})}) t^n.$$

Our next goal will be to express $H_{(\overline{v},\overline{l})}(t)$ as a rational function in t. From here, standard enumerative combinatorics will allow us to compute the Hilbert polynomial.

Proposition 3.15   *The Hilbert series $H_{(\bar{v},\bar{l})}(t)$ can be expressed as the rational function*

$$H_{(\bar{v},\bar{l})}(t) = \frac{t^i \prod_j (1 - (1-t)^{l_j})}{(1-t)^{\mu(\bar{v})}}.$$

Proof   We will prove this using the fact that the Hilbert series is multiplicative over tensor products. It is easy to see that the Hilbert series for the augmentation ideal of a rational polynomial ring in d variables is equal to the rational function

$$\frac{1}{(1-t)^d} - 1 = \frac{1 - (1-t)^d}{(1-t)^d}.$$

Lemma 3.13 therefore implies

$$H_{(\bar{v},\bar{l})}(t) = t^i \left(\prod_j \frac{1 - (1-t)^{l_j}}{(1-t)^{l_j}} \right) \left(\frac{1}{(1-t)^{\mu(\bar{v})-|\bar{l}|}} \right) = \frac{t^i \prod_j (1 - (1-t)^{l_j})}{(1-t)^{\mu(\bar{v})}},$$

where the factor t^i arises from the graded shift of $N^{(\bar{v},\bar{l})}$.   □

This will allow us to prove our result on the obstruction to the Hilbert polynomial of $\mathcal{M}_{i,\bullet}$.

Corollary 3.16   *Let G be a tree and fix $i \geq 0$. Then the dimension of $\mathcal{M}_{i,n} \otimes_{\mathbb{Z}} \mathbb{Q}$ agrees with a polynomial for all $n \geq 0$.*

Proof   It suffices to prove the claim for the summands $N^{(\bar{v},\bar{l})}$. Proposition 3.15 tells us that the Hilbert series of $N^{(\bar{v},\bar{l})}$ takes the form

$$\frac{t^i \prod_j (1 - (1-t)^{l_j})}{(1-t)^{\mu(\bar{v})}}.$$

The proof of Corollary 2.21 implies that it suffices to show the total degree of the numerator $t^i \prod_j (1 - (1-t)^{l_j})$ is strictly smaller than $\mu(\bar{v})$. Namely, we must argue

$$i + |\bar{l}| < \mu(\bar{v}).$$

This follows from the fact that each essential vertex of $\bar{v}$, say v_{j_k}, is adjacent to at least one essential edge which can never house a witness vertex, namely, the essential edge which is maximal among all essential edges adjacent to v_{j_k} with respect to the ordering of the vertices of G. Moreover, the essential edge containing the root can never house a witness vertex for any essential vertex.

Put another way,

$$i + |\bar{l}| \le i + \sum_k (\mu(v_{j_k}) - 2) = i + \mu(\bar{v}) - i - 1 = \mu(\bar{v}) - 1.$$

This concludes the proof. □

We have now proven enough to explicitly compute the Hilbert polynomial of $\mathcal{M}_{i,\bullet}$, and consequently $H_i(B_n G)$. The following theorem follows from the results of this section, as well as the combinatorics of generating functions.

Theorem 3.17 *Let G be a tree, fix $i, r \ge 0$ and, for any pair $(\bar{v}, \bar{l})$, write*

$$a_{i,r}^{(\bar{v},\bar{l})} := \sum_{\substack{d_1 + \cdots + d_i = r \\ d_j \ge 1}} (-1)^{r+i} \binom{l_1}{d_1} \cdots \binom{l_i}{d_i}.$$

Also let $P_i^{\bar{v}}(n)$ be the polynomial

$$P_i^{\bar{v}}(n) = \binom{\mu(\bar{v}) - 1 + n - i}{\mu(\bar{v}) - 1}.$$

Then for all $n \ge 0$, the dimension of $H_i(B_n, \mathbb{Q})$ is equal to the polynomial

$$(3\text{-}3) \qquad \sum_{(\bar{v},\bar{l})} \sum_{r=i}^{|\bar{l}|} a_{i,r}^{(\bar{v},\bar{l})} P_i^{\bar{v}}(n - r).$$

Proof It suffices to understand the polynomials which describe the coefficients of the rational functions

$$\frac{t^i \prod_j (1 - (1 - t)^{l_j})}{(1 - t)^{\mu(\bar{v})}}.$$

This is a simple exercise in the combinatorics of generating functions. □

Remark 3.18 From the above formulation, the polynomial $P_i^{\bar{v}}(n)$ as well as the constants $a_{i,r}^{(\bar{v},\bar{l})}$ only depend on the degree sequence of the graph G. It follows that the homologies of the braid group of a tree only depend on its degree sequence, as we recorded in Corollary B

Specializing Theorem 3.17 to the case wherein $i = 1$ yields a much simpler form.

Corollary 3.19 Let G be a tree. Then for all $n \geq 0$, the dimension of $H_1(B_n G, \mathbb{Q})$ agrees with the polynomial

$$\sum_{v \text{ essential}} \sum_{l=1}^{\mu(v)-2} \binom{\mu(v)-2+n}{\mu(v)-1} - \binom{\mu(v)-2+n-l}{\mu(v)-1-l}.$$

Proof In this case we note that for an essential vertex v and a chosen integer $0 \leq l \leq \mu(v)-2$, we have

$$a_{1,r}^{v,l} = (-1)^{r+1}\binom{l}{r}.$$

Therefore, the polynomial (3-3) can be written as

$$\sum_{v \text{ essential}} \sum_{l=1}^{\mu(v)-2} \sum_{r=1}^{l}(-1)^{r+1}\binom{l}{r}\binom{\mu(v)-2+n-r}{\mu(v)-1}.$$

Using the principle of inclusion–exclusion, it can be seen that

$$\sum_{r=1}^{l}(-1)^{r+1}\binom{l}{r}\binom{\mu(v)-2+n-r}{\mu(v)-1} = \binom{\mu(v)-2+n}{\mu(v)-1} - \binom{\mu(v)-2+n-l}{\mu(v)-1-l},$$

as desired. $\square$

A result of this kind appears in the work of Farley and Sabalka, where they compute the number of generators of the group $B_n G$ [8]. A formula for the case $i = 1$ is also found in the results of Ko and Park [15].

Another simple case is that in which every essential vertex of G has degree 3. In this case, Theorem 3.17 can be used to recover the following result of Farley [7]:

Corollary 3.20 Let G be a tree whose every essential vertex has degree 3. Then for all $n \geq 0$, the dimension of $H_i(B_n G; \mathbb{Q})$ agrees with the polynomial

$$\binom{N_G}{i}\binom{n}{2i}.$$

Proof The major thing to note in this case is that for any $\bar{v}$, the only associated vector $\bar{l}$ is $(1, \ldots, 1)$. Therefore,

$$a_{i,r}^{(\bar{v},\bar{l})} = \begin{cases} 1 & \text{if } r = i, \\ 0 & \text{otherwise.} \end{cases}$$

The formula (3-3) now becomes

$$\binom{N_G}{i} P_i^{\bar{v}}(n-i) = \binom{N_G}{i}\binom{\mu(\bar{v})-1+n-2i}{\mu(\bar{v})-1}$$
$$= \binom{N_G}{i}\binom{2i+1-1+n-2i}{2i} = \binom{N_G}{i}\binom{n}{2i}.$$

This completes the proof. □

4 Concluding remarks

In this section we consider what this work implies for the homology of graph braid groups when G is a general graph rather than a tree.

Let G be a graph with N_G essential vertices and E essential edges. Also assume that G is homeomorphic to neither S^1 nor a line segment. Then the work of Section 3.2 implies that $\mathcal{M}_{i,\bullet}$ carries the structure of a finitely generated S_G–module. If G is not a tree, however, it is no longer the case that the Morse differential is trivial. We therefore pose the following question:

Question 1 Does the action of S_G on $\mathcal{M}_{i,\bullet}$ commute with the Morse differential? In other words, is the Morse differential a morphism of graded S_G–modules?

We note that altering our choice of spanning tree T has a nontrivial influence on the Morse differential. It is therefore unclear whether the Morse differential will commute with the action of S_G for all choices of T. We therefore refine Question 1 as follows:

Question 2 Does there exist a choice of spanning tree T of G such that the action of S_G on $\mathcal{M}_{i,\bullet}$ commutes with the Morse differential?

If the answer to these questions is affirmative, then we can immediately conclude many things about the asymptotic behavior of these homology groups. Indeed, in this case, the Hilbert basis theorem would imply that there is a finitely generated S_G–module $\mathcal{H}_i$ such that $(\mathcal{H}_i)_n \cong H_i(B_nG)$.

Corollary 4.1 *Assume that Questions 1 and 2 have an affirmative answer. For each i and n, write*

$$H_i(B_nG) \cong \mathbb{Z}^{b_i(n)} \oplus \bigoplus_{p \text{ prime}} \bigoplus_{j \geq 1} (\mathbb{Z}/p^j\mathbb{Z})^{m_{p,i,j}(n)}.$$

Then:

(1) *There exists a polynomial $P_i(t) \in \mathbb{Q}[t]$ of degree $\leq E-1$ such that $P_i(n) = b_i(n)$ for all $n \gg 0$.*

(2) *There exists a positive integer e_i, independent of n, such that the exponent of $H_i(B_n G)$ is at most e_i.*

(3) *For each prime p and positive integer j, there is a polynomial $P_{i,p^j}(t) \in \mathbb{Q}[t]$ such that $P_{i,p^j}(n) = m_{p,i,j}(n)$ for $n \gg 0$. In particular, there is a finite set of primes $p_{i,1}, \ldots, p_{i,l}$ and polynomials $m_{i,p_j}(n)$ such that the order of the torsion subgroup of $H_i(B_n G)$ is precisely $p_1^{m_{p_{i,1}}(n)} \cdots p_l^{m_{p_{i,l}}(n)}$.*

Proof This follows from standard facts in the study of graded modules over integral polynomial rings. $\square$

The work of Ko and Park [15, Theorem 3.6] shows that $H_1(B_n G)$ has a torsion subgroup annihilated by 2, and that the multiplicity of this torsion is eventually constant in n. Their work also implies that the rank of $H_1(B_n G)$ will be constant whenever G is sufficiently connected. In contrast to this, affirmative answers to Questions 1 and 2 would imply that the top homology of $B_n G$ always grows as fast as possible.

Theorem 4.2 *Assume that Questions 1 and 2 have an affirmative answer. Then the Hilbert function*
$$n \mapsto \dim_{\mathbb{Q}}(H_{N_G}(B_n G; \mathbb{Q}))$$
is equal to a polynomial of degree $E-1$ for $n \gg 0$.

Proof Corollary 2.21 implies that the Euler characteristic of $\mathrm{UConf}_n(G)$ is a polynomial of degree $E-1$. On the other hand, Remark 3.4 implies the lower homology groups have Hilbert functions which grow strictly slower than this. It follows that the rank of $H_{N_G}(B_n G)$ must eventually grow like a polynomial of degree $E-1$. $\square$

Note that Theorem 2.6 implies that $H_{N_G}(B_n G)$ is torsionfree. We have therefore seen that $H_1(B_n G)$ can only have 2–torsion, while the top homology $H_{N_G}(B_n G)$ does not have any torsion. It is an interesting question to ask what other types of torsion can appear in the intermediate homologies.

While an affirmative answer to Questions 1 and 2 would be a step in the right direction, the work of Ko and Park [15] would suggest that the action of S_G on $\mathcal{M}_{i,\bullet}$ described in Section 3.2 is perhaps nonoptimal. We have already discussed that the result [15,

Theorem 3.6] implies that the braid groups of biconnected graphs have eventually constant first homologies. However, the action of S_G defined in Section 3.2 only detects the connectivity of a chosen spanning tree of G. In particular, is it possible that one can improve on this action by accounting for connectivity granted by the deleted edges?

Question 3 Let G be a graph which is neither a circle nor a line segment. Does there exist an action of S_G on $\mathcal{M}_{i,\bullet}$ which interacts in a nontrivial way with the deleted edges of G? More specifically, can one define this action so that if c is a critical i–cell, then there is a $\mathbb{Z}$–linear relation between $x_{\mathfrak{e}_1} \cdot c$ and $x_{\mathfrak{e}_2} \cdot c$ whenever the two essential edges $\mathfrak{e}_1$ and $\mathfrak{e}_2$ are in the same connected component of G with the edges of c removed?

The existence of such an action would actually suggest something quite strong about what kind of invariants of G are encoded by the homology of its braid groups. If G is a graph, recall that we define Δ_G^i to be the maximum number of connected components G can be broken into by removing exactly i vertices.

Conjecture 4.3 *Let G be a graph which is neither a line segment nor a circle. Then for all i and all $n \gg i$, the Betti number*

$$b_i(n) := \dim_{\mathbb{Q}}(H_i(B_n G; \mathbb{Q}))$$

agrees with a polynomial of degree $\leq \Delta_G^i - 1$.

Note that Theorem 4.2 is a specific case of this conjecture. Also note, if G is biconnected, then the above implies that $b_1(n)$ is eventually constant. This fact is proven in the work of Ko and Park [15, Theorem 3.6]. The above conjecture therefore proposes a generalization of one aspect of that work.

Finally, we think about the ramifications of this work on the spaces $\mathrm{UConf}_n(G)$. The action of S_G on $\mathcal{M}_{i,\bullet}$ is defined in a purely formal way. However, in the study of more classical configuration spaces, stability phenomena often arise from maps on the underlying spaces $\mathrm{UConf}_n(G) \to \mathrm{UConf}_{n+1}(G)$ (see for instance [6; 18]). Is that also the case here?

Question 4 If G is a tree, is the action of S_G on $\mathcal{M}_{i,\bullet}$ induced by a map of topological spaces $\mathrm{UConf}_n(G) \to \mathrm{UConf}_{n+1}(G)$? If G is a general graph, and Questions 1 and 2 have affirmative answers, is the action of S_G on the homologies of $B_n G$ induced in a similar way?

Of course, Question 4 is highly relevant for Questions 1 and 2. If one understands the action of S_G topologically, then it might make the existence of such an action on the homologies of general graph braid groups more apparent.

Remark 4.4   Very recent work of An, Drummond-Cole and Knudsen [2] drastically extends the scope of the current work. In doing so, they provide affirmative answers to most of the questions of this section. It is unclear, however, whether their work directly implies Conjecture 4.3. This remains an interesting open question.

References

[1]   **A Abrams**, *Configuration spaces and braid groups of graphs*, PhD thesis, University of California, Berkeley (2000)   Available at http://home.wlu.edu/~abramsa/publications/thesis.ps

[2]   **B H An**, **G C Drummond-Cole**, **B Knudsen**, *Subdivisional spaces and graph braid groups*, preprint (2017)   arXiv

[3]   **K Barnett**, **M Farber**, *Topology of configuration space of two particles on a graph, I*, Algebr. Geom. Topol. 9 (2009) 593–624   MR

[4]   **T Church**, **J S Ellenberg**, **B Farb**, *FI–modules and stability for representations of symmetric groups*, Duke Math. J. 164 (2015) 1833–1910   MR

[5]   **D Eisenbud**, *The geometry of syzygies: a second course in commutative algebra and algebraic geometry*, Graduate Texts in Mathematics 229, Springer (2005)   MR

[6]   **J S Ellenberg**, **J D Wiltshire-Gordon**, *Algebraic structures on cohomology of configuration spaces of manifolds with flows*, preprint (2015)   arXiv

[7]   **D Farley**, *Homology of tree braid groups*, from "Topological and asymptotic aspects of group theory" (R Grigorchuk, M Mihalik, M Sapir, Z Šunić, editors), Contemp. Math. 394, Amer. Math. Soc., Providence, RI (2006) 101–112   MR

[8]   **D Farley**, **L Sabalka**, *Discrete Morse theory and graph braid groups*, Algebr. Geom. Topol. 5 (2005) 1075–1109   MR

[9]   **D Farley**, **L Sabalka**, *Presentations of graph braid groups*, Forum Math. 24 (2012) 827–859   MR

[10]   **R Forman**, *Morse theory for cell complexes*, Adv. Math. 134 (1998) 90–145   MR

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

[12]   **Ś R Gal**, *Euler characteristic of the configuration space of a complex*, Colloq. Math. 89 (2001) 61–67   MR

[13] **R Ghrist**, *Configuration spaces and braid groups on graphs in robotics*, from "Knots, braids, and mapping class groups: papers dedicated to Joan S Birman" (J Gilman, W W Menasco, X-S Lin, editors), AMS/IP Stud. Adv. Math. 24, Amer. Math. Soc., Providence, RI (2001) 29–40 MR

[14] **J H Kim**, **K H Ko**, **H W Park**, *Graph braid groups and right-angled Artin groups*, Trans. Amer. Math. Soc. 364 (2012) 309–360 MR

[15] **K H Ko**, **H W Park**, *Characteristics of graph braid groups*, Discrete Comput. Geom. 48 (2012) 915–963 MR

[16] **D Lütgehetmann**, *Configuration spaces of graphs*, Master's thesis, Freie Universität Berlin (2014) Available at http://luetge.userpage.fu-berlin.de/pdfs/masters-thesis-luetgehetmann.pdf

[17] **T Maciążek**, **A Sawicki**, *Homology groups for particles on one-connected graphs*, J. Math. Phys. 58 (2017) art. id. 062103 MR

[18] **D McDuff**, *Configuration spaces of positive and negative particles*, Topology 14 (1975) 91–107 MR

[19] **E Miller**, **B Sturmfels**, *Combinatorial commutative algebra*, Graduate Texts in Mathematics 227, Springer (2005) MR

[20] **P Prue**, **T Scrimshaw**, *Abrams's stable equivalence for graph braid groups*, preprint (2009) arXiv

[21] **S V Sam**, **A Snowden**, *Proof of Stembridge's conjecture on stability of Kronecker coefficients*, J. Algebraic Combin. 43 (2016) 1–10 MR

[22] **J Świątkowski**, *Estimates for homological dimension of configuration spaces of graphs*, Colloq. Math. 89 (2001) 69–79 MR

Department of Mathematics, University of Michigan
Ann Arbor, MI, United States

egramos@umich.edu