Homology of Tree Braid Groups

Daniel Farley

Abstract. The braid group on n-strands of a tree T, denoted B_nT, has free abelian integral homology groups in every dimension. The rank of $H_i(B_nT)$ is equal to the number of i-cells in UC^nT, the unlabelled configuration space of n points on T, which are critical with respect to a certain discrete Morse function.

1. Introduction

If Γ is a finite graph and n is a natural number, then the *labelled configuration space* $\mathcal{C}^n\Gamma$ is the n-fold Cartesian product of Γ, with the set $\Delta = \{(x_1, \ldots, x_n) \mid x_i = x_j \text{ for some } i \neq j\}$ removed. The *unlabelled configuration space* $U\mathcal{C}^n\Gamma$ is the quotient of $\mathcal{C}^n\Gamma$ by the action of the symmetric group S_n, where the action permutes the factors. The fundamental group of $U\mathcal{C}^n\Gamma$ (respectively, $\mathcal{C}^n\Gamma$) is the *braid group* (respectively, the *pure braid group*) *of* Γ *on* n *strands*, denoted $B_n\Gamma$ (respectively, $PB_n\Gamma$).

An earlier paper [**5**] used Robin Forman's discrete Morse theory [**7**] (see also [**3**]) in order to simplify the spaces $U\mathcal{C}^n\Gamma$ within their homotopy types. As a result, it was possible to compute presentations for B_nT, where n is an arbitrary positive integer and T is any finite tree. Here similar methods are used to compute the integral homology groups of any tree braid group, i.e., any braid group B_nT where T is a tree. In fact, the homology of $U\mathcal{C}^n\Gamma$ (equivalently, the homology of B_nT) can be computed from the Morse chain complex associated to the discrete flow from [**5**]. The group of i-dimensional chains in the Morse complex is the free $\mathbb{Z}$-module generated by critical i-cells. The main argument in this note shows that the differential maps in the Morse complex are all 0, so that $H_i(B_nT)$ is free abelian, of rank equal to the number of critical i-cells. The number of critical i-cells can be readily computed in individual cases, but here we give an explicit computation only for the case in which T is 3-regular.

2000 *Mathematics Subject Classification.* Primary 20F65, 20F36; Secondary 57M15, 55R80.

Key words and phrases. tree braid groups, homology groups, configuration spaces, discrete Morse theory.

It is worth mentioning that the Morse-theoretic methods of [**5**] and the current note can be applied to the groups $B_n\Gamma$ and the spaces $UC^n\Gamma$ for any graph Γ, although the results to date are less satisfactory than in the specific case of trees.

I thank Lucas Sabalka for participating in a number of conversations on the subject of this paper.

2. Discrete Morse Theory

The following exposition is taken from [**5**] and [**6**] (see also [**7**] and [**3**], which were the original sources for the Morse-theoretic arguments in [**6**] and [**5**]). Let X be a finite CW complex. Let K_i be the (finite) collection of the open i-cells of X. If $\sigma \in K_i$, $\tau \in K_{i+1}$, and $\sigma \subseteq \overline{\tau}$, then σ is a *regular face* of τ if there is a characteristic map $h : B^{i+1} \to \overline{\tau}$ such that $h : h^{-1}\sigma \to \sigma$ is a homeomorphism and $\overline{h^{-1}(\sigma)}$ is a closed i-ball.

A *discrete vector field* V on X is a sequence of partial functions $V_i : K_i \to K_{i+1}$ such that:

(1) each V_i is injective;
(2) if $V_i(\sigma) = \tau$, then σ is a regular face of τ;
(3) $im(V_i) \cap dom(V_{i+1}) = \emptyset$.

(A *partial function* from a set A to a set B is a function defined on a subset of A.)

Let V be a discrete vector field on X. A *V-path of dimension i* is a sequence of i-cells $\sigma_0, \sigma_1, \ldots, \sigma_r$ such that if $V(\sigma_j)$ is undefined, then $\sigma_{j+1} = \sigma_j$, and otherwise $\sigma_{j+1} \neq \sigma_j$ and σ_{j+1} is a face of $V(\sigma_i)$ (i.e., $\sigma_{j+1} \subseteq V(\sigma_i)$). The V-path is *closed* if $\sigma_r = \sigma_0$, and *non-stationary* if $\sigma_1 \neq \sigma_0$. A discrete vector field V is a *discrete gradient vector field* if V has no non-stationary closed paths.

A discrete gradient vector field on X results in a classification of the cells on X into three categories: an open cell σ is *redundant* if σ is in the domain of V, *collapsible* if it is in the range of V, and *critical* if it is in neither the range nor the domain of V. The definition of a discrete vector field ensures that these categories are mutually exclusive, and it is clear that any open cell of X must be of one of these types.

Choose an orientation for each cell of X, and consider the standard cellular chain complex of X, denoted $C_*(X)$. Let $B_i(X)$ denote the standard basis of $C_i(X)$, consisting of positively oriented i-cells of X. Let $\pm B_i(X)$ denote $B_i(X) \cup -B_i(X) \cup \{0\}$. A discrete gradient vector field V corresponds to a collection of maps $V_i : B_i(X) \to \pm B_{i+1}(X)$ satisfying:

(1) if $V(b_1) = V(b_2) \neq 0$, then $b_1 = b_2$;
(2) $VV = 0$;
(3) for any $b \in B_i(X)$ $(i \geq 0)$, either $V(b) = 0$ or b occurs in $\partial V b$ with the coefficient -1;
(4) for $b_1, b_2 \in B_i(X)$, write $b_1 > b_2$ if b_2 occurs in $\partial V b_1$ or $V\partial b_1$ with a non-zero coefficient and $b_1 \neq b_2$.

(a) for each $b_0 \in B_i(X)$, there are only finitely many $b_1 \in B_i(X)$ satisfying $b_0 > b_1$;
(b) the irreflexive, transitive closure of $>$, also denoted $>$, is a strict partial order with no infinite descending chains.

The maps V_i may be extended linearly to maps $V_i : C_i(X) \to C_{i+1}(X)$. The terms *collapsible cell*, *redundant cell*, and *critical cell* have the same definitions as before, if we use the maps $V_i : C_i(X) \to C_{i+1}(X)$ in the earlier definition.

Let $f = 1 + \partial V + V\partial$. Thus $f : C_*(X) \to C_*(X)$ is a chain map, called the *discrete flow* corresponding to V. For any finite chain $c \in C_*(X)$, there is some $m \in \mathbb{N}$ such that $f^m(c) = f^{m+1}(c)$, so there is a well-defined chain map $f^\infty : C_*(X) \to C_*(X)$ ([**6**]; Proposition 16). The *Morse complex* $(M_*(X), \widetilde{\partial}_*)$ consists of the groups $M_i(X)$, which are defined to be the free abelian groups on the basis of critical i-cells, together with the differentials $\pi\partial f^\infty$, where ∂ is the standard boundary map in cellular homology and $\pi : C_i(X) \to C_i(X)$ is the projection onto the free abelian subgroup spanned by the critical i-cells.

THEOREM 2.1. *([**6**]; Theorem 18) The complexes $\left(M_*(X), \widetilde{\partial}_*\right)$ and $(C_*(X), \partial_*)$ have isomorphic homology groups in all dimensions.* □

It will be useful to have a somewhat simplified version of the boundary map in the Morse complex. Let $F = 1 + \partial V$.

LEMMA 2.2.

(1) F^∞ is well-defined. That is, for any finite chain c in $C_*(X)$, there is some $m \in \mathbb{N}$ such that $F^m(c) = F^{m+1}(c)$.
(2) $\pi F^\infty \partial = \pi \partial f^\infty$.

PROOF.

(1) If c is a collapsible or critical cell, then $V(c) = 0$, so $F(c) = c$. If c is redundant, then the argument of [**6**], Lemma 14, shows that either $F(c) = 0$ or that any redundant cell occurring with a non-zero coefficient in $F(c)$ has smaller height than c, where height is as defined in [**6**]. Since the height is a positive integer, there must exist some m such that $F^m(c) = F^{m+1}(c)$; (1) follows.
(2) Pick some finite chain c. We can choose m so large that $f^\infty(c) = f^m(c)$ and $F^\infty(\partial c) = F^m(\partial c)$. Thus it is sufficient to show that $F^m\partial = \partial f^m$. But this follows immediately from the identity $F\partial = \partial f$, which follows from the definitions of F and f.

□

3. Proof of the Main Theorem

We will apply the discrete version of Morse theory from the previous section to a certain cubical complex $UD^n\Gamma$. Consider the product

$$\prod_{\ell=1}^{n} \Gamma.$$

This space has a natural cubical structure, in which the cubes are products of the edges and vertices of Γ. Let Δ' denote the union of those open cells whose closures intersect the *fat diagonal* Δ:

$$\Delta = \left\{ (x_1, x_2, \ldots, x_n) \in \prod_{\ell=1}^{n} \Gamma \mid x_i = x_j \text{ for some } i \neq j \right\}.$$

Let $\mathcal{D}^n\Gamma$ denote the space $\prod^n \Gamma - \Delta'$. The symmetric group S_n acts on $\mathcal{D}^n\Gamma$ by permuting the factors. The action is free by the definition of $\mathcal{D}^n\Gamma$, and the quotient is denoted $U\mathcal{D}^n\Gamma$. The spaces $\mathcal{D}^n\Gamma$ and $U\mathcal{D}^n\Gamma$ are "discretized" versions of the configuration spaces $\mathcal{C}^n\Gamma$ and $U\mathcal{C}^n\Gamma$, respectively. Note that the precise definition of $\mathcal{D}^n\Gamma$ and $U\mathcal{D}^n\Gamma$ depends upon a subdivision of the original graph Γ. Aaron Abrams [**1**] has shown that the discretized configuration spaces are homotopy equivalent with their originals under appropriate hypotheses:

THEOREM 3.1. [**1**] *For any $n > 1$ and any graph Γ with at least n vertices, the labelled (respectively, unlabelled) configuration space of n points on Γ strong deformation retracts onto $\mathcal{D}^n\Gamma$ ($U\mathcal{D}^n\Gamma$) if*

(1) each path between distinct vertices of degree not equal to 2 passes through at least $n-1$ edges; and
(2) each path from a vertex to itself that cannot be shrunk to a point in Γ passes through at least $n+1$ edges. □

Any graph can be subdivided to satisfy the hypotheses of this theorem, and we assume from now on that any graph under consideration has already been subdivided so. We confine our attention to the unlabelled, discretized spaces $U\mathcal{D}^nT$, where T is a (finite) tree. Since the closed cubes of $\mathcal{D}^nT$ are products (of closed cells) $c_1 \times \ldots \times c_n$ such that $c_i \cap c_j = \emptyset$, the closed cubes of $U\mathcal{D}^nT$ can be usefully identified with collections $\{c_1, \ldots, c_n\}$ of closed cells such that $c_i \cap c_j = \emptyset$.

A discrete gradient vector V on $U\mathcal{D}^nT$ will be defined using a certain ordering of the vertices of T. Embed the tree in the plane, and choose a basepoint $*$ of degree 1. Assign numbers to the other vertices by the following method: trace the tree, beginning with the vertex $*$, taking the leftmost turn at any vertex of degree greater than or equal to 3, and turning around upon reaching a vertex of degree 1. The vertices are numbered $1, 2, \ldots$ according to the order in which they are first encountered, with the vertex adjacent to $*$ being numbered 1. If v_1 and v_2 are vertices of T, we sometimes write that $v_1 > v_2$ if v_1 has a higher number than v_2. The edges e of T are oriented so

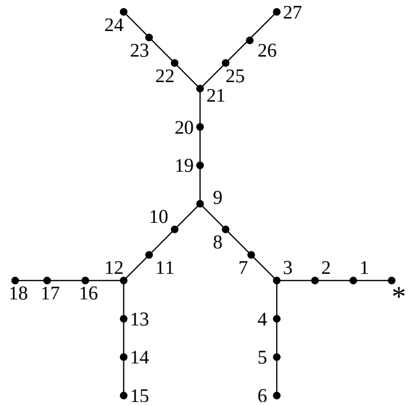

Figure 1

that the initial vertex $\iota(e)$ is farther from $*$ than the terminal vertex $\tau(e)$ (and thus $\iota(e) > \tau(e)$).

We will need the following technical lemma about the order.

LEMMA 3.2. *If e is an edge of T, and v_1, v_2 are vertices of T such that $v_1 \in (\tau(e), \iota(e)) = \{v \in T^0 \mid \tau(e) < v < \iota(e)\}$ and $v_2 \notin (\tau(e), \iota(e))$, then v_1 and v_2 are in different path components of $T - e$. In particular, if e_1 is an edge satisfying $e_1 \cap e = \emptyset$, then the endpoints of e_1 are either both in $(\tau(e), \iota(e))$, or neither is.* □

If v is any vertex of T other than $*$, then let $e(v)$ be the initial edge of the unique geodesic issuing from v and connecting to $*$. If $v \in \{c_1, \ldots, c_n\} = \{c_1, \ldots, c_{n-1}, v\}$, then v is *unblocked* if $\{c_1, \ldots, c_{n-1}, e(v)\}$ is a cell of $U\mathcal{D}^nT$, i.e., if $e(v) \cap c_j = \emptyset$ for $j = 1, \ldots, n-1$. Otherwise, v is *blocked*.

The discrete gradient vector field V is defined inductively. If $c = \{c_1, \ldots, c_n\}$ is a 0-cell, i.e., if $c_1, c_2, \ldots, c_n$ are all vertices of T, then $V_0(c)$ is obtained by replacing the smallest unblocked vertex v of $\{c_1, \ldots, c_n\}$ with $e(v)$, where "smallest" is meant in the sense of the numbering of vertices. If there are no unblocked vertices, then $V(c)$ is undefined. Now assume that V_{i-1} has been defined. If $c = \{c_1, \ldots, c_n\}$ is an (i-dimensional) cell in the range of V_{i-1}, then $V_i(c)$ is undefined (this is forced by the definition of a discrete vector field). If c is not in the range of V_{i-1}, then $V_i(c)$ is defined exactly as in the 0-dimensional case; $V_i(c)$ is obtained by replacing the smallest unblocked vertex v of c with $e(v)$, and $V_i(c)$ is undefined if there are no unblocked vertices.

It is proved in [**5**] that the family of partial functions V are a discrete gradient vector field. The resulting classification of cells into critical, collapsible, and redundant types is easiest to describe with another definition. An edge e in $c = \{c_1, \ldots, c_n\}$ is *order-respecting* with respect to c if, for every vertex v in c, $e(v) \cap e = \tau(e)$ implies that $v > \iota(e)$. Note that the property of being order-respecting depends very much on the cell c, and a given edge might be order-respecting in one cell c but fail to be order-respecting in another cell c'.

EXAMPLE 3.3. This example is intended to clarify the notion of "order-respecting" edge. Consider the discretized configuration space $U\mathcal{D}^2Y$, where the graph Y is simply the capital letter "Y", consisting of the wedge of three intervals $[0, 1]$, such that 0 is the wedge point in each. The figure depicts two 1-cells of the configuration space $U\mathcal{D}^2Y$. Note that the numbering of the vertices is correct for the embedding of Y into the plane and the choice of basepoint.

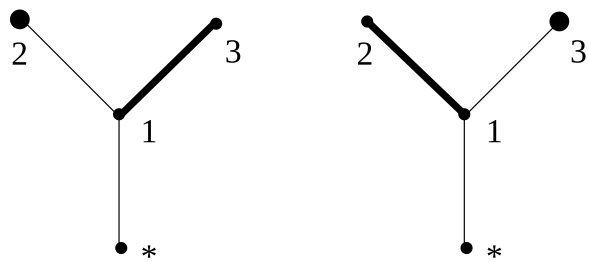

Figure 2

The cell on the left, $\{v_2, [v_1, v_3]\}$, has an edge, $[v_1, v_3]$, which fails to be order-respecting, since $e(v_2) \cap [v_1, v_3] = v_1 = \tau([v_1, v_3])$, but $v_2 \leq v_3 = \iota([v_1, v_3])$. The edge $[v_1, v_2]$ in the cell on the right is order-respecting.

We have the following theorem:

THEOREM 3.4. *([**5**], Theorem 3.6)*

(1) A cell is critical if and only if it contains no order-respecting edges and all of its vertices are blocked.
(2) A cell is redundant if and only if
   (a) it contains no order-respecting edges and at least one of its vertices is unblocked OR
   (b) it contains an order-respecting edge (and thus an order-respecting edge e such that $\iota(e)$ is minimal in the set of initial vertices of order-respecting edges e' of c) and there is some unblocked vertex v such that $v < \iota(e)$.
(3) A cell is collapsible if and only if it contains an order-respecting edge e (and thus an order-respecting edge e such that $\iota(e)$ is minimal in the set of initial vertices of order-respecting edges e' of c) and, for any $v < \iota(e)$, v is blocked. □

Here is a simple way to determine whether a given cell $c = \{c_1, \ldots, c_n\}$ is critical, collapsible, or redundant. Arrange the cells $c_1, \ldots, c_n$ from least to greatest, where the edges are to be ordered according to the numbering of their initial vertices. Now scan the cells beginning with the least and proceeding in order to the greatest. If you find an unblocked vertex first, then the cell c is redundant. If you find an order-respecting edge first, the cell is collapsible. If you scan the entire list of cells $c_1, \ldots, c_n$, finding neither an unblocked vertex nor an order-respecting edge, then the cell is critical. This way of classifying a given cell is a consequence of the above theorem.

We now turn to a description of the boundary operator in the cellular chain complex $C_*(U\mathcal{D}^nT)$ of $U\mathcal{D}^nT$. Each cell of $U\mathcal{D}^nT$ is a cube, so we can use the standard boundary operator for a cubical complex [**4**]. Let $[0,1]^k$ be a standard k-dimensional cube. Let A_i (respectively, B_i) be the face corresponding to freezing the ith coordinate of $[0,1]^k$ at 0 (respectively, 1). With these conventions,

$$\partial\left([0,1]^k\right) = \sum_{i=1}^{k} (-1)^i (A_i - B_i).$$

Identify a given k-dimensional cube c of $U\mathcal{D}^nT$ with $[0,1]^k$ as follows: if $e_1, e_2, e_3, \ldots, e_k$ are the edges of c listed in increasing order, i.e., so that $\iota(e_1) < \iota(e_2) < \iota(e_3) < \ldots < \iota(e_k)$, then regard e_1 as the first factor in $[0,1]^k$, e_2 as the second, etc., where the initial vertex of an edge e_i is identified with 0 (and the terminal vertex of e_i is identified with 1). This identification induces an orientation of each cell in $C_*(U\mathcal{D}^nT)$. The resulting boundary operator $\partial : C_*(U\mathcal{D}^nT) \to C_*(U\mathcal{D}^nT)$ has the property that opposite faces of a cube c have opposite signs in ∂c.

Let c be an i-dimensional redundant cell. According to the classification of cells, this means that c contains some unblocked vertex, and thus an unblocked vertex v that is minimal among unblocked vertices in the ordering (i.e., numbering) of vertices. Define $r(c)$, the *elementary reduction* of c, to be the result of replacing this minimal unblocked vertex v with the vertex v_1 lying adjacent to v on the geodesic segment $[v, *]$ in T. Note that $r(c)$ is the face of $V(c)$ opposite to c. If c is not redundant, let $r(c) = c$. Since applying r repeatedly must eventually result in a collapsible or critical cell, the function r^∞ is well-defined.

The *order defect* of a redundant cell c is the difference $n(v) - n(v_1)$, where $n(v)$ denotes the number of v, and the vertices v and v_1 are as above. The order defect of a redundant cell is necessarily a positive integer.

The following lemma (cf. [**5**], Lemma 5.1) simplifies calculations greatly.

Lemma 3.5. *(Redundant cubes lemma for trees)*

(1) If the i-cell c contains an order-respecting edge e such that every vertex $v \in (\tau(e), \iota(e)) \cap c$ is blocked, then $\pi F^\infty(c) = 0$.
(2) If c is any redundant cell, then $\pi F^\infty(c) = \pi F^\infty r(c)$. Thus, in particular, $\pi F^\infty(c) = \pi F^\infty r^\infty(c)$.

Proof.

(1) If c contains an order-respecting edge, then it must be collapsible or redundant. If c is collapsible, then

$$F(c) = (1 + \partial V)(c) = c + \partial V V(c') = c.$$

Thus, there is nothing left to prove when c is collapsible.

We can therefore assume that c is redundant. According to the classification of redundant cells, c must contain a minimal unblocked vertex v such that $v < \iota(e')$, for any order-respecting edge e' in c. In particular, the vertex v must satisfy $v < \tau(e)$, since certainly $v < \iota(e)$ and there are also no unblocked vertices in $(\tau(e), \iota(e))$. The cell $V(c)$ is obtained by replacing v with the edge $[v, v_1]$, where v_1 is the vertex adjacent to v on the geodesic from v to $*$. The edge $[v, v_1]$ must be order-respecting (see [**5**], Lemma 3.4) and every vertex in $V(c) \cap (v_1, v)$ must be blocked, for otherwise v would fail to be the minimal unblocked vertex in c. We must have $v_1 < v < \tau(e) < \iota(e)$. The faces of $V(c)$ (i.e., the cells occurring with a non-zero coefficient in $\partial V(c)$) are each obtained by replacing a single edge $\tilde{e}$ in $V(c)$ with either $\tau(\tilde{e})$ or $\iota(\tilde{e})$. The result will be to unblock certain vertices which were adjacent to the edge $\tilde{e}$, and introduce a new vertex (either $\tau(\tilde{e})$ or $\iota(\tilde{e})$, as the case may be) which will necessarily be unblocked. It is impossible to introduce unblocked vertices into both (v_1, v) and $(\tau(e), \iota(e))$ in this fashion, however, and at least one of $[v, v_1]$ and e must remain order-respecting, so each cell occurring in $F(c)$ with a non-zero coefficient will again satisfy the hypothesis on c. We conclude by induction that each cell in $F^{\infty}(c)$ satisfies the hypothesis on c. This implies that none of these cells is critical, so $\pi F^{\infty}(c) = 0$.

(2) The proof is by induction on the order defect of c. Suppose first that the order defect of c is 1. Note that $V(c)$ satisfies the hypothesis of the first part of the lemma, as any collapsible cell must. Since the order defect of c is 1, there is some edge e in $V(c)$ such that $\tau(e)$ and $\iota(e)$ are consecutively numbered (and $\tau(e) < \iota(e)$). Such an edge is necessarily order-respecting in any cell containing it. Every cell occurring in the chain $\partial V(c)$ contains the edge e except for c (which occurs with the coefficient -1) and $r(c)$ (occurring with coefficient 1). Thus,

$$F(c) = (1 + \partial V)(c) = r(c) + c',$$

where c' is a chain consisting entirely of cells satisfying the hypothesis of part 1) of the lemma. It follows that

$$F^{\infty}(c) = F^{\infty}(r(c)).$$

Now suppose that the order defect of $c = \{c_1, \ldots, c_{n-1}, v\}$ is k. Let v be the minimal unblocked vertex of c. Thus, $V(c) = \{c_1, c_2, \ldots, c_{n-1}, e(v)\}$. The chain $F(c) = (1 + \partial V)(c)$ consists of the term $r(c)$, as well as a chain consisting of cells $\tilde{c}$ such that $\tilde{c}$

is obtained by replacing an edge $\tilde{e} \neq e(v)$ of $\{c_1, \ldots, c_{n-1}, e(v)\}$ with either its initial or its terminal vertex. By Lemma 3.2, either both endpoints of $\tilde{e}$ are in (v_1, v) or neither endpoint is. If neither endpoint of $\tilde{e}$ is in (v_1, v), then the faces corresponding to $\iota(\tilde{e})$ and $\tau(\tilde{e})$ satisfy the hypotheses of part (1) of the lemma. Let c' be the cellular chain consisting of all such faces. Part (1) of the lemma implies that $\pi F^\infty(c') = 0$.

Now consider the other faces of $V(c)$, which come from replacing an edge $\tilde{e} \in V(c) - \{e(v)\}$ such that $v_1 < \tau(\tilde{e}) < \iota(\tilde{e}) < v$, with either $\tau(\tilde{e})$ or $\iota(\tilde{e})$. Each of these cells will be redundant, since some of the vertices near $\tilde{e}$ will be unblocked. (At the very least, the vertex $\tau(\tilde{e})$ or $\iota(\tilde{e})$ will have to be unblocked in such a face.) Any new unblocked vertex in any one of these faces will lie between v_1 and v in the order on vertices.

Fix a specific edge $\tilde{e}$ such that $v_1 < \tau(\tilde{e}) < \iota(\tilde{e}) < v$, and consider the corresponding pair of opposite faces A_j and B_j. Each of these cells is redundant. We repeatedly apply the induction hypothesis, using the fact that the order defect of A_j, $r(A_j)$, etc., is always less than k:

$$\pi F^\infty(A_j) = \pi F^\infty r(A_j) = \ldots = \pi F^\infty r^\infty(A_j).$$

Similar reasoning applies to the face B_j:

$$\pi F^\infty(B_j) = \pi F^\infty r^\infty(B_j).$$

But $r^\infty(A_j) = r^\infty(B_j)$, and the faces A_j and B_j appear with opposite signs in $F(c)$, so, if c'' is the chain consisting of the remaining faces of $V(c)$, i.e., the ones which correspond to edges $\tilde{e}$ such that $v_1 < \tau(\tilde{e}) < \iota(\tilde{e}) < v$, then

$$\pi F^\infty\left(c''\right) = 0,$$

since the terms cancel in pairs.

We conclude that

$$\begin{aligned}\pi F^\infty(c) &= \pi F^\infty r(c) + \pi F^\infty\left(c'\right) + \pi F^\infty\left(c''\right)\\ &= \pi F^\infty r(c).\end{aligned}$$

This completes the induction.

□

Example 3.6. A simple example may help to clarify the proof of the second part of the above lemma. Refer to the numbered tree in Figure 1, which is correctly subdivided for $n = 4$. We consider the redundant cell $c = \{*, v_{13}, e_{16}, v_{19}\}$, where e_{16} is the edge having v_{16} as its initial vertex. The smallest unblocked vertex is v_{19}, and the edge $e(v_{19})$ connects v_{19} to v_9, so the order defect of c is 10. If we apply F, we get a chain consisting of $r(c)$, together with the chains c' and c'' (as described in the proof above). The chain c' is 0 in this case. The chain c'' consists of the cells $c_1 = \{*, v_{12}, v_{13}, e_{19}\}$ and $c_2 = \{*, v_{13}, v_{16}, e_{19}\}$, appearing with opposite signs. The order defect of each is 1. We can repeatedly apply the induction hypothesis to these cells, since all of the resulting redundant cells have order defect less than 10 (in fact, the largest-occurring order defect is 4). It is rather clear that $r^\infty(c_1) = r^\infty(c_2)$, so we conclude that the terms from c'' cancel in

$$\pi F^\infty(c) = \pi F^\infty r(c) + \pi F^\infty \left(c''\right),$$

and the conclusion of the lemma follows for this case.

THEOREM 3.7. *For any critical i-cell c in $\mathcal{C}_i\left(U\mathcal{D}^n T\right)$, where T is a tree,*

$$\pi F^\infty \partial(c) = 0.$$

In particular, $H_i\left(B_n T; \mathbb{Z}\right)$ is free abelian of rank equal to the number of critical i-cells.

PROOF. Let c be a critical i-cell. We consider pairs A_j, B_j of opposite faces in ∂c. By the description of the cubical boundary operator, these faces appear with opposite signs in ∂c. Without loss of generality, assume that $A_j = \{c_1, c_2, \ldots, c_{i-1}, \iota(e)\}$ and $B_j = \{c_1, c_2, \ldots, c_{i-1}, \tau(e)\}$. Since all of the vertices in the original cell c were blocked, all of the vertices of A_j and B_j are also blocked, except for a collection in each clustered around the vertex $\tau(e)$. After repeatedly applying r to A_j and B_j, the vertices near $\tau(e)$ move until they are all blocked, and we arrive at $r^\infty\left(A_j\right)$ and $r^\infty\left(B_j\right)$. These are the same cell. Thus,

$$\begin{aligned} \pi F^\infty\left((-1)^j\left(A_j - B_j\right)\right) &= (-1)^j\left(\pi F^\infty\left(A_j\right) - \pi F^\infty\left(B_j\right)\right) \\ &= (-1)^j\left(\pi F^\infty r^\infty\left(A_j\right) - \pi F^\infty r^\infty\left(B_j\right)\right) \\ &= 0. \end{aligned}$$

The first statement of the theorem now follows by applying the same reasoning to each pair of opposing faces.

The statement about homology follows from the fact that the space $U\mathcal{D}^n T$ admits a metric of non-positive curvature [**1**], and thus is aspherical [**2**]. As a result, we have

$$H_i\left(B_n T; \mathbb{Z}\right) \cong H_i\left(U\mathcal{D}^n T; \mathbb{Z}\right),$$

and the rest of the theorem follows. □

4. The Case of 3-regular Trees

The main theorem reduces the problem of computing the homology of a tree braid group to the problem of counting the number of critical cells with respect to the discrete gradient vector field V. A general solution to the latter problem appears to be rather complicated, but it is possible to give an easy solution in a special case. A tree T is *k-regular* if every vertex of T has degree 1, 2, or k.

Proposition 4.1. *(3-regular case) If T is a 3-regular tree having V vertices of degree 3, then $H_i\left(B_nT;\mathbb{Z}\right)$ is free abelian of rank* $\left(\begin{array}{c} V \\ i \end{array}\right)\left(\begin{array}{c} n \\ 2i \end{array}\right)$.

Proof. We count critical i-cells c. Begin by choosing the locations of the edges e in c. Since these edges must all fail to be order-respecting by the classification of critical cells, it must be that $\tau(e)$ is a vertex of degree 3, for each edge e of c. Fix a specific edge e. Suppose that $\tau(e) = v$, where v has degree 3. There are exactly two possibilities for e: either e lies to the left of v or to the right. (That is, when we walk from the basepoint $*$ to v, e lies either to our left or right. This is well-defined since the tree T comes with a specific embedding into the plane.) We can rule out the possibility that e lies to the left of v. In fact, if e were the edge to the left of v, then $\tau(e)$ and $\iota(e)$ would necessarily be consecutively numbered, and this would force e to be an order-respecting edge. It follows that we can identify an edge e in c by $\tau(e)$. Since $\tau(e)$ must be a vertex of degree 3, and there are i edges in c (since c has dimension i), it follows that there are V choose i possible locations for the edges of c.

Now we choose the locations of the vertices in c. The locations of i of these vertices are already determined by the locations of the edges, and by the fact that each edge must fail to be order-respecting, i.e., a vertex must sit at the place adjacent to $\tau(e)$ and to the left of it, for each edge e in c, for otherwise e would be order-respecting. We've now determined the locations of $2i$ cells in c. Since all vertices in c must be blocked, The cell c will now be completely determined by specifying how many of the $n-2i$ remaining cells (all of which are vertices) are in each connected component of $T - \{\tau(e) \mid e \textit{ is an edge of } c\}$. There are $2i+1$ such components, so we must decide how to distribute $n-2i$ indistinguishable vertices among $2i+1$ distinguishable connected components. There are exactly n choose $2i$ ways to do this. □

References

[1] Aaron Abrams. *Configuration spaces of braid groups of graphs.* PhD thesis, UC Berkeley, 2000.

[2] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[3] Kenneth S. Brown. The geometry of rewriting systems: a proof of the Anick-Groves-Squier theorem. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 137–163. Springer, New York, 1992.

[4] Kenneth S. Brown and Ross Geoghegan. An infinite-dimensional torsion-free FP_∞ group. *Invent. Math.*, 77(2):367–381, 1984.

[5] Daniel Farley and Lucas Sabalka. Discrete Morse theory and graph braid groups. arXiv:math.GR/0410539.

[6] Daniel S. Farley. Homological and finiteness properties of picture groups. *Trans. Amer. Math. Soc.*, 357(9):3567–3584 (electronic), 2005.

[7] Robin Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.

Department of Mathematics, University of Illinois at Urbana-Champaign, Champaign, IL 61820

E-mail address: farley@math.uiuc.edu