

HOMOLOGICAL AND FINITENESS PROPERTIES OF PICTURE GROUPS

DANIEL S. FARLEY

ABSTRACT. Picture groups are a class of groups introduced by Guba and Sapir. Known examples include Thompson's groups F , T , and V .

In this paper, a large class of picture groups is proved to be of type F_∞ . A Morse-theoretic argument shows that, for a given picture group, the rational homology vanishes in almost all dimensions.

1. INTRODUCTION

There is now a substantial theory of diagram groups (see for example [11], [15], [16], [17], [18], and [19]). Guba and Sapir first introduced annular and braided diagram groups in [15] (see Section 16). Here, as in [10], these are called annular picture groups and picture groups, respectively. Some background on annular picture groups and picture groups is in Section 2.

One of the original reasons to study diagram groups was that Thompson's groups F , T , and V (see [9] for an introduction) occur as diagram groups over a particularly simple semigroup presentation, namely $\langle x \mid x = x^2 \rangle$. To be precise, T occurs as an annular diagram group, and V occurs as a braided diagram group (see [10] and [15]). It is often possible to prove that long-known theorems about Thompson's groups are also true more generally of diagram groups. For example, in [11], it was shown that any diagram group over a finite presentation of a finite semigroup is of type F_∞ , which generalized work of Brown and Geoghegan [8], who first showed that F is of type F_∞ . In [10] and [11], it was shown that all types of diagram groups act properly on CAT(0) cubical complexes (see [3], page 158, for the definition of CAT(0), and page 111 for the definition of a cubical complex). This had not been known, even in the case of Thompson's groups.

Theorem 1 of this paper states that a wide class of annular diagram groups and braided diagram groups (hereafter annular picture groups and picture groups) are of type F_∞ . The precise hypothesis of this theorem is too technical to state here, but the class of groups to which the theorem applies includes the groups T and V . Theorem 1 may thus be considered a generalization of Brown's work from [5] on Thompson's group. A proof is in Section 3.

The rest of the paper is devoted to proving that the rational homology of picture groups vanishes in almost all dimensions. Fix a semigroup presentation $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$, where $|\mathcal{R}| = n$. Theorem 11 states that the rational homology of

Received by the editors December 4, 2003.

2000 *Mathematics Subject Classification*. Primary 20J05, 20F65.

Key words and phrases. Picture groups, diagram groups, finiteness properties of groups, Morse theory.

©2004 American Mathematical Society
Reverts to public domain 28 years from publication

any picture group over \mathcal{P} vanishes in dimensions greater than n . Using a Morse-theoretic argument, one can then show that, under somewhat technical hypotheses, the rational homology vanishes in dimension n as well. This is Theorem 21. One corollary of this theorem is that Thompson's group V has the rational homology of a point, which was first proved by Brown [6].

I thank the referee for doing an extraordinarily helpful and efficient job. The argument in Section 3 of this paper is significantly clearer as a result, particularly because of the reference [2]. I thank Ken Brown for answering some of my questions regarding the chain complex in Section 4.

2. BACKGROUND

Let $\mathcal{P} = \{\Sigma \mid \mathcal{R}\}$ be a semigroup presentation. Thus, Σ is an alphabet and the set of relations \mathcal{R} is a set of equalities between positive, non-empty words. In this paper, the relations all have a preferred left side and no relation of the form $w = w$ occurs. If u is the left side of the relation $u = v$, then it will sometimes be useful to emphasize this by writing $u \rightarrow v$. Assume also that if $(u \rightarrow v) \in \mathcal{R}$, then $(v \rightarrow u) \notin \mathcal{R}$.

A *picture over \mathcal{P}* consists of a *frame* $\partial([0, 1]^2)$, a finite, possibly empty, collection of *transistors*, each homeomorphic to $[0, 1]^2$, and a finite non-empty collection of *wires*, each homeomorphic to $[0, 1]$. Each wire is labelled by a letter of the alphabet Σ . The frame and transistors come with well-defined top, bottom, left, and right sides.

Any given wire connects a point on the bottom of one transistor to a point on the top of another, or a point on the top or bottom of the frame to a point on the top or bottom (respectively) of a transistor, or a point on the top of the frame to a point on the bottom. The wires are also required to be pairwise disjoint. The endpoints of a wire are called *contacts* and each contact in a picture inherits a label from the unique wire it touches. Let T_1 and T_2 be transistors. Write $T_1 > T_2$ if there is a wire connecting the bottom of T_1 to the top of T_2 , and let $>$ also denote the transitive closure of this relation. One further requires of the attaching maps of the wires that $>$ be a strict partial order on the set of transistors.

There is one further condition on the transistors. Fix a transistor T . A word u_T may be obtained by reading the labels of the top contacts of T from left to right. This is the *top label of T* . A word v_T , called the *bottom label of T* , is defined similarly. The final requirement is that $u_T = v_T$ (or $v_T = u_T$) be a member of \mathcal{R} . (It is not enough for $u_T = v_T$ or $v_T = u_T$ to be in the congruence generated by \mathcal{R} .) The relation $u_T \rightarrow v_T$ *labels T* if u_T is the top label and v_T is the bottom label.

Any labelled topological space satisfying the above conditions is a picture over \mathcal{P} . One can define the top and bottom labels for a picture Δ over \mathcal{P} in exactly the way the top and bottom labels of a transistor were defined. If the top label of Δ is the word u and the bottom label is v , then Δ is a (u, v) -*picture*. A $(u, *)$ -*picture* is a picture having u as its top label, and any word for the bottom label. Figure 1 portrays two $(acbd, abab)$ -pictures over the semigroup presentation $\mathcal{P} = \langle a, b, c, d \mid cb = bc, ab = cd, ab = ba \rangle$. The frames appear as dotted boxes. Ignore the dotted circles for now.

A picture without transistors is a *permutation picture*. A picture in which no two transistors are comparable with respect to the partial order on transistors is called *thin*.




FIGURE 1.

Two pictures Δ_1 and Δ_2 over \mathcal{P} are *isomorphic* if there is a homeomorphism between them which matches labels and preserves the orientations on the frame and transistors; i.e., both the top-bottom orientations and the left-right orientations. One writes that $\Delta_1 \equiv \Delta_2$.

Given a (u, v) -picture Δ_1 and a (v, w) -picture Δ_2 , one can define the *concatenation* $\Delta_1 \circ \Delta_2$, which is the (u, w) -picture obtained by identifying the bottom of the frame for Δ_1 with the top of the frame for Δ_2 by a homeomorphism which matches contacts, and then removing the line segment corresponding to the bottom of Δ_1 in the quotient, while keeping the wires passing through this line segment intact.

Two transistors $T_1 > T_2$ form a *dipole* if the top label of T_1 is identical (as a word in Σ^+ , the free semigroup on Σ) to the bottom label of T_2 , and the bottom contacts of T_1 are paired off by wires in order with the top contacts of T_2 . To *remove a dipole*, delete the transistors T_1 and T_2 and all wires connecting them, and then glue together in order the wires that formed top contacts of T_1 with those that formed bottom contacts of T_2 . The inverse operation is called *inserting a dipole*. Two pictures Δ_1 and Δ_2 are *equal modulo dipoles* if one can be obtained from the other by repeatedly inserting and removing dipoles. One writes that $\Delta_1 = \Delta_2$. Figure 1 shows a dipole (inside the circle in the left picture) and the effect of removing this dipole (the picture on the right).

A picture is *reduced* if it contains no dipoles. By Lemma 2.2 of [10], there is a unique reduced picture in any equivalence class of pictures modulo dipoles.

For a fixed word $w \in \Sigma^+$, the set of all (w, w) -pictures over \mathcal{P} , modulo dipoles, forms a group $\mathcal{D}_b(\mathcal{P}, w)$ under the operation of concatenation. These groups are the same as the braided diagram groups defined in [15]; they were called picture groups in [10]. One can also define *annular picture groups* (as in [10]). Given a picture Δ , identify the left side of the frame $\{0\} \times [0, 1]$ with the right side $\{1\} \times [0, 1]$ in an orientation-preserving way. Now remove the open line segment that corresponds to $\{0\} \times (0, 1)$ in the quotient. The result looks like a picture, except that the frame has been replaced by two disjoint circles. Each circle receives an orientation coming from the left-right orientation of the top and bottom of the frame, respectively. If the new space $\bar{\Delta}$ may be embedded in the plane in an orientation-preserving fashion, then the original picture Δ is called an *annular picture*. The set of annular (w, w) -pictures forms a group denoted $\mathcal{D}_a(\mathcal{P}, w)$. These groups were also studied in [15], where they were called annular diagram groups.

The groups $\mathcal{D}_b(\mathcal{P}, w)$ and $\mathcal{D}_a(\mathcal{P}, w)$ act on CAT(0) cubical complexes ([10, Theorem 4.5]). The CAT(0) cubical complexes for $\mathcal{D}_b(\mathcal{P}, w)$ and $\mathcal{D}_a(\mathcal{P}, w)$ are denoted




FIGURE 2.

$\tilde{K}_b(\mathcal{P}, w)$ and $\tilde{K}_a(\mathcal{P}, w)$, respectively. Here the complexes $\tilde{K}_b(\mathcal{P}, w)$ will be described; the complexes $\tilde{K}_a(\mathcal{P}, w)$ are similar, but involve only annular pictures in their definition. A *vertex* (or *formal vertex*) is an equivalence class of reduced pictures, where the pictures Δ_1 and Δ_2 are considered equivalent if $\Delta_1 \circ \Phi \equiv \Delta_2$, for some picture Φ containing no transistors. A *vertex of $\mathcal{D}_b(\mathcal{P}, w)$* is a vertex consisting of $(w, *)$ -pictures. (The vertices of $\tilde{K}_a(\mathcal{P}, w)$ are defined in the same way, but all of the pictures in question are annular.)

It is convenient to depict a vertex as a picture in which the bottom wires have been clipped, so they do not touch the bottom of the frame. Figure 2a) portrays a vertex in $\tilde{K}_b(\mathcal{P}, x)$, where $\mathcal{P} = \langle x \mid x = x^2 \rangle$. Let $C_b(\Delta, \Psi) = \{[\Delta \circ \Psi'] \mid [\Psi'] \leq [\Psi]\}$, where Δ is a $(w, *)$ -picture, Ψ is *thin*, i.e., no two of the transistors of Ψ are comparable in the partial order on transistors, and $[\Psi'] \leq [\Psi]$ if there is some picture Φ such that $\Psi \circ \Phi$ and Ψ' are isomorphic without reducing dipoles [10]. The set $C_b(\Delta, \Psi)$ is an *abstract cube* (or *cube*) of $\tilde{K}_b(\mathcal{P}, w)$. A cube can be represented by a picture in which the bottom wires have been clipped, and some subset of the set of minimal transistors of the picture are painted white (see Lemma 3.4 in [10]). (Here “minimal” means minimal in the sense of the partial order $<$ on transistors.) Figure 2b) represents a cube in $\tilde{K}_b(\mathcal{P}, x)$, where $\mathcal{P} = \langle x \mid x = x^2 \rangle$.

Choose some arbitrary white transistor to be numbered 1, another to be numbered 2, and so on, until every white transistor is assigned a natural number. See Figure 3. If there are n white transistors, then this numbering describes a labelling




FIGURE 3.




FIGURE 4.

of the standard n -cube $[0, 1]^n$. The label of $(i_1, \dots, i_n) \in [0, 1]^n$ is the vertex obtained by removing the k th white transistor if $i_k = 0$, shading the k th white transistor if $i_k = 1$, and removing any dipoles that occur. Figure 4 shows how the numbering in Figure 3 describes how to label $[0, 1]^2$. The cubes of $\tilde{K}_b(\mathcal{P}, w)$ are glued together along faces sharing the same labellings. Details are in [10]. The standard n -cube $[0, 1]^n = I^n$ inherits a natural orientation from \mathbb{R}^n , so that the identification of a cube C with $[0, 1]^n$ gives C an orientation as well.

A cube is said to be in *standard form* if any white transistor T is labelled by a relation $u \rightarrow v$ in \mathcal{R} . A thin picture is in standard form if each of its transistors satisfies the same condition. Any cube may be put into standard form. Figure 5 illustrates in a simple case how to convert a cube into standard form. The cube represented on the left is an edge in $\tilde{K}_b(\mathcal{P}, w)$, where $\mathcal{P} = \langle x \mid x^2 \rightarrow x \rangle$ and $w = x$. The idea is to shade any offending transistor, attach a new transistor in order to form a dipole, and then paint the new (bottom) transistor white.

A *b-word* is an equivalence class of words under the equivalence relation \sim , where $w_1 \sim w_2$ if w_1 and w_2 are equal in Σ^+ after a permutation of their letters. Let $\langle w \rangle_b$ denote the *b-word* containing the word w . A *cyclic word* is an equivalence class




FIGURE 5.

of words under the equivalence relation \approx , where $w_1 \approx w_2$ if w_1 and w_2 are cyclic shifts of each other. Let $\langle w \rangle_a$ denote the cyclic word containing the word w . Define the *bottom label of a vertex v in $\mathcal{D}_b(\mathcal{P}, w)$* to be the b -word formed by the collection of bottom labels of pictures $\Delta \in v$. The *bottom label of a vertex in $\mathcal{D}_a(\mathcal{P}, w)$* can be similarly defined as a cyclic word.

A *rewrite system* is an oriented graph; the vertices of a rewrite system are *objects* and the directed edges are *moves*. If there is a move from one vertex v_1 to another, say v_2 , then write $v_1 \rightarrow v_2$. If there is no move from a given object v , that object is *reduced*. A rewrite system is *terminating* if every sequence $a_1 \rightarrow \dots \rightarrow a_n \rightarrow \dots$ is finite. The reflexive, transitive closure of the relation \rightarrow is a partial order on the vertices of a terminating rewrite system.

To a given semigroup presentation \mathcal{P} , one can associate three different rewrite systems. Let $\Gamma(\mathcal{P})$ be the rewrite system having the words of Σ^+ as its vertex set, and having an oriented edge originating at u and terminating at v if there is a relation $u_2 \rightarrow v_2$ such that $u = u_1 u_2 u_3$ and $v = u_1 v_2 u_3$, where one or both of the words u_1 and u_3 may be empty, and the “=” means equality in the free semigroup Σ^+ . Let $\Gamma_a(\mathcal{P})$ be the rewrite system obtained from $\Gamma(\mathcal{P})$ by identifying any two vertices that represent the same cyclic word. Let $\Gamma_b(\mathcal{P})$ be the rewrite system obtained by identifying any two vertices that represent the same b -word. If some (or all) of these rewrite systems are terminating, then there is a partial order on the set of words, cyclic words, or b -words in the alphabet, as the case may be.

3. SOME GROUPS OF TYPE F_∞

The goal in this section is to prove the following theorem:

Theorem 1. *If \mathcal{P} is a finite semigroup presentation and $\Gamma_b(\mathcal{P})$ (respectively, $\Gamma_a(\mathcal{P})$) is a terminating rewrite system with only finitely many reduced objects, then $\mathcal{D}_b(\mathcal{P}, w)$ (respectively, $\mathcal{D}_a(\mathcal{P}, w)$) is of type F_∞ .*

The arguments for $\mathcal{D}_b(\mathcal{P}, w)$ and $\mathcal{D}_a(\mathcal{P}, w)$ are the same in almost every respect, so it is enough to make the argument for $\mathcal{D}_b(\mathcal{P}, w)$ and note what needs to be changed to cover the other case.

It will be helpful to use a proposition from [5].

Proposition 2 ([5], Corollary 3.3, part a.). *Let X be a contractible G -complex such that the stabilizer of every cell is of type F_∞ . Let $\{X_j\}$ be a filtration such that each X_j is finite mod G . Suppose the connectivity of the pair (X_{j+1}, X_j) tends to infinity as j tends to infinity. Then G is of type F_∞ . \square*

Assume now that \mathcal{P} is a finite semigroup presentation and the rewrite system $\Gamma_b(\mathcal{P})$ is terminating and has only finitely many reduced objects. Choose a sequence $(\langle w_n \rangle_b)$ of b -words (i.e., vertices of $\Gamma_b(\mathcal{P})$) to have the property that, for any word $\bar{w} \in \Gamma_b(\mathcal{P})^0$, whenever $\bar{w} \leq w_j$, one has $\bar{w} \in \{w_1, \dots, w_j\}$. This can be done by induction under the given hypotheses. For each j , let $V_j = \{\Delta \in \tilde{K}_b(\mathcal{P}, w)^0 \mid \text{the bottom label of } \Delta \text{ is in } \{w_1, \dots, w_j\}\}$, and let K_j be the largest subcomplex of $\tilde{K}_b(\mathcal{P}, w)$ with the property that $K_j^0 = V_j$. Each K_j is $\mathcal{D}_b(\mathcal{P}, w)$ -finite since each V_j is $\mathcal{D}_b(\mathcal{P}, w)$ -finite and $\tilde{K}_b(\mathcal{P}, w)$ itself is locally finite (see the proof of Theorem 4.9 in [10] for a discussion of the orbit of a vertex under the action of $\mathcal{D}_b(\mathcal{P}, w)$, and

for the proof that $\tilde{K}_b(\mathcal{P}, w)$ is locally finite). Moreover, the stabilizer in $\mathcal{D}_b(\mathcal{P}, w)$ of any cell in $\tilde{K}_b(\mathcal{P}, w)$ is finite (see Theorem 4.9 in [10]), so the first condition of Proposition 2 is satisfied.

Consider the effect of passing from K_k to K_{k+1} . The bottom label of each of the vertices in K_{k+1} is the b -word $\langle w_{k+1} \rangle_b$. Let $[\Delta]$ be an arbitrary vertex in $K_{k+1}^0 - K_k^0$. If ϵ is small, then $K_{k+1} - N_\epsilon([\Delta])$ will strong deformation retract onto the complex K' by radial projection, where K' is the largest subcomplex of $\tilde{K}_b(\mathcal{P}, w)$ having $K_{k+1}^0 - \{[\Delta]\}$ as its zero skeleton.

The strong deformation retractions from each vertex in $K_{k+1}^0 - K_k^0$ are compatible with one another, because no cube of K_{k+1} can contain two vertices having the bottom label $\langle w_{k+1} \rangle_b$ (cf. [11], Lemma 4.7). The boundary of each neighborhood $N_\epsilon([\Delta])$ is homeomorphic to $|\text{lk}_\downarrow(\langle w_{k+1} \rangle_b)|$, where $\text{lk}_\downarrow(\langle v \rangle_b)$ is defined below. Each such ϵ -neighborhood is itself a cone on $|\text{lk}_\downarrow(\langle w_{k+1} \rangle_b)|$, and hence contractible. It follows that K_{k+1} is homotopy equivalent to the complex obtained by attaching to K_k countably many cones on $|\text{lk}_\downarrow(\langle w_{k+1} \rangle_b)|$ along their bases. From this observation, it is easy to argue, using the Mayer-Vietoris sequence, van Kampen's theorem, and the Hurewicz theorem (see, e.g., [22]), that

Proposition 3. *If $|\text{lk}_\downarrow(\langle w_{k+1} \rangle_b)|$ is n -connected, then the pair (K_{k+1}, K_k) is n -connected. \square*

Now one needs to define $\text{lk}_\downarrow(\langle v \rangle_b)$. First, some additional definitions are necessary. The formal vertex $[\Psi]$ uses a contact on the top of its frame if the wire leading from this contact is attached to the top of some transistor of $[\Psi]$. If two formal vertices $[\Psi_1]$ and $[\Psi_2]$ have the same top label v , then $[\Psi_1]$ and $[\Psi_2]$ overlap if there is some i such that $[\Psi_j]$, for $j = 1, 2$, uses the i th contact from the left on the top of its frame.

Definition 4 (cf. [1]). The descending link of the b -word $\langle v \rangle_b$, denoted $\text{lk}_\downarrow(\langle v \rangle_b)$, is the simplicial complex satisfying:

- i) $\text{lk}_\downarrow(\langle v \rangle_b)^0 = \{[\Psi] \mid [\Psi] \text{ is a vertex with one transistor, } \Psi \text{ is in standard form, and the top label of } [\Psi] \text{ is } v\}$.
- ii) a collection of vertices in $\text{lk}_\downarrow(\langle v \rangle_b)^0$ forms a simplex if and only if no two of them overlap.

The isomorphism type of $\text{lk}_\downarrow(\langle v \rangle_b)$ does not depend on the choice in i) of the word v from $\langle v \rangle_b$.

The descending link of a cyclic word $\langle v \rangle_a$, denoted $\text{lk}_\downarrow(\langle v \rangle_a)$, is defined in the same way, except that all pictures occurring in the definition are assumed to be annular. The descending link of a word v , denoted $\text{lk}_\downarrow(v)$, was defined in [11], but it can be defined in an equivalent way with the above definition, if every picture in this definition is understood to be "ordinary", i.e., planar.

The above characterization of the descending link follows easily from the isomorphism between $\text{lk}_{ab}([\Pi])$ and (V_v, S_v) ([10], Proposition 4.7); here the bottom label of $[\Pi]$ is $\langle v \rangle_b$. It also follows that $\text{lk}_\downarrow(\langle v \rangle_b)$ is a flag complex, from the same reasoning that is used to show that $\text{lk}_{ab}([\Pi])$ is a flag complex (see the proof of Theorem 4.9 in [10]). (Recall that a simplicial complex L is a *flag complex* if every finite subset of L^0 that is pairwise joined by edges spans a simplex ([3], page 210).)

The rest of the argument requires some definitions, most of which can be found in [21]. Let K be a simplicial complex. The subcomplex L is *full* in K if L is the

largest subcomplex of K having L^0 as its 0-skeleton. The *star* of a vertex v in K , denoted $\text{st}(v, K)$, or $\text{st}(v)$ if the ambient complex is understood, is the subcomplex of K consisting of all the simplices containing v , together with the faces of these simplices. It is not difficult to check that the star L of a vertex v in a flag complex K is a full subcomplex of K . Thus, the star of a vertex v in a flag complex is simply the full subcomplex generated by v and all adjacent vertices. The star of any vertex in a simplicial complex is contractible.

The *link* of a vertex v in a simplicial complex K , denoted $\text{lk}(v, k)$, consists of those simplices spanned by vertices v_1, v_2, \dots, v_k such that $v \notin \{v_1, v_2, \dots, v_k\}$ and v, v_1, v_2, \dots, v_k span a simplex in K . Thus the link of a vertex v in a flag complex K is also a full subcomplex, and $\text{lk}(v, K)^0 = \text{st}(v, K)^0 - \{v\}$.

Recall that if \mathcal{C} is a cover of a topological space X , then the *nerve* of this cover is the simplicial complex having elements of \mathcal{C} as its set of vertices, and a collection of vertices spans a simplex if the corresponding elements of \mathcal{C} have non-empty intersection. It will be useful to appeal to the Nerve Theorem (see, e.g., A. Björner: “Topological Methods” in [2], Theorem 10.6, page 1850):

Theorem 5 (Nerve Theorem). *Let Δ be a simplicial complex and $(\Delta_i)_{i \in I}$ a family of subcomplexes such that $\Delta = \bigcup_{i \in I} \Delta_i$. If every non-empty finite intersection $\Delta_{i_1} \cap \dots \cap \Delta_{i_t}$ is $(k - t + 1)$ -connected, then Δ is k -connected if and only if the nerve $\mathcal{N}(\Delta_i)$ is k -connected. \square*

The following lemma will be used in the remainder of the proof of Theorem 1.

Lemma 6. *Let K be a finite flag complex.*

i) *K is $(n - 1)$ -connected provided, for any collection of vertices $S \subseteq K^{(0)}$, the intersection $\bigcap_{v \in S} \text{lk}(v)$ is $(n - |S|)$ -connected.*

ii) *If S is any collection of vertices of K and $\bigcap_{v \in S} \text{lk}(v)$ is n -connected, then so is $\bigcap_{v \in S} \text{st}(v)$.*

Proof. For each set S of vertices, put

$$K_S := \bigcap_{v \in S} \text{st}(v) \quad \text{and} \quad L_S := \bigcap_{v \in S} \text{lk}(v).$$

Let L_S^c denote the *complementary complex* of L_S in K_S that contains precisely those simplices of K_S that are disjoint from L_S .

The next step is to show that the complex L_S^c is a simplex unless it is empty. Since K_S is flag (being an intersection of stars in a flag complex), it suffices to show that any two distinct vertices u and w in L_S^c are joined by an edge. Since u is not in L_S , there is a vertex $v \in S$ not connected to u by an edge. On the other hand, $u \in K_S$ whence $u \in \text{st}(v)$. Thus, $v = u$ and so u belongs to S . Now, $w \in K_S$ implies $w \in \text{st}(u)$, so w and u span an edge. Thus L_S^c is a simplex.

By the same reasoning, every vertex in L_S is connected to every vertex in L_S^c . Since K_S is flag, $K_S = L_S * L_S^c$ whence it is either contractible (if L_S^c is non-empty) or at least n -connected. This proves ii). Statement i) of the lemma follows from the Nerve Theorem applied to the cover of K by the stars of all of the vertices in K . Indeed, all of the subcomplexes K_S have the required connectivity. Moreover, K_S is non-empty unless $|S| > n + 1$. Hence the n -skeleton of the nerve N is isomorphic to the n -skeleton of a huge simplex. Thus, the n -skeleton of N is $(n - 1)$ -connected, which implies that K is $(n - 1)$ -connected. \square

The proof of Theorem 1 is, up to this point, exactly the same for the groups $\mathcal{D}_b(\mathcal{P}, w)$ and $\mathcal{D}_a(\mathcal{P}, w)$. The only difference of any substance between the two cases occurs here:

Proposition 7. *Suppose \mathcal{P} is a finite semigroup presentation and $\Gamma_b(\mathcal{P})$ (respectively, $\Gamma_a(\mathcal{P})$) is a terminating rewrite system with only finitely many reduced objects. Let N_b be an integer such that $N_b > \ell(v)$ for any reduced b -word in $\Gamma_b(\mathcal{P})^0$. Let N be an integer such that $N > \ell(v)$ for any reduced word in $\Gamma(\mathcal{P})^0$. Let R be the length of the longest word occurring as the left side of some relation in \mathcal{R} . Let $n \in \mathbb{Z}$.*

- i) *If v is a b -word and $\ell(v) \geq 2(n + 1)R + N_b$, then $\text{lk}_\downarrow(\langle v \rangle_b)$ is n -connected.*
- ii) *If v is a cyclic word and $\ell(v) \geq (n + 2)N + (n + 3)(R - 1)$, then $\text{lk}_\downarrow(\langle v \rangle_a)$ is n -connected.*

Proof. i) The proof is by induction on n . In the first non-trivial case, $n = -1$, one has $\ell(v) \geq N_b$. By the assumption on N_b , this implies that $\langle v \rangle_b$ is not reduced, so there is some picture in standard form having v as its top label and having exactly one transistor, which implies that $\text{lk}_\downarrow(\langle v \rangle_b)$ is non-empty.

For the inductive step, fix $v \in \langle v \rangle_b$. Each $[\Psi] \in \text{lk}_\downarrow(\langle v \rangle_b)^0$ is a formal vertex in standard form with a single transistor, and having the top label v .

Let $\ell(v) \geq 2(n + 1)R + N_b$. To show that $\text{lk}_\downarrow(\langle v \rangle_b)$ is n -connected, it is enough, by Lemma 6 i), to demonstrate that $|\text{lk}([\Psi_1])| \cap \dots \cap |\text{lk}([\Psi_k])|$ is $(n - k + 1)$ -connected, for any choice $[\Psi_1], \dots, [\Psi_k]$ of vertices. At most kR contacts are used by the formal vertices $[\Psi_1], \dots, [\Psi_k]$. The complex $|\text{lk}([\Psi_1])| \cap \dots \cap |\text{lk}([\Psi_k])|$ is the full subcomplex of $\text{lk}_\downarrow(\langle v \rangle_b)$ spanned by the vertices of $\text{lk}_\downarrow(\langle v \rangle_b)$ which do not overlap with any of the $[\Psi_i]$, for $i = 1, \dots, k$. Thus $|\text{lk}([\Psi_1])| \cap \dots \cap |\text{lk}([\Psi_k])|$ is homeomorphic to $|\text{lk}_\downarrow(\langle v_1 \rangle_b)|$, where v_1 satisfies

$$\begin{aligned} \ell(v_1) &\geq 2(n + 1)R - kR + N \\ &\geq 2(n - k + 2)R + N. \end{aligned}$$

This implies that $|\text{lk}_\downarrow(\langle v_1 \rangle_b)|$ is $(n - k + 1)$ -connected, completing the proof.

ii) The first non-trivial case to consider occurs when $n = -1$. In this case $\ell(v) \geq N + 2(R - 1) \geq N$. Thus, for any word $v \in \langle v \rangle_a$, v is not reduced. It follows that $\langle v \rangle_a$ itself is not reduced, so $\text{lk}_\downarrow(\langle v \rangle_a)$ is non-empty.

Now let

$$\ell(v) \geq (n + 2)N + (n + 3)(R - 1) \quad (n \geq 0)$$

and fix a specific word $v \in \langle v \rangle_a$. Fix a vertex $[\Phi] \in \text{lk}_\downarrow(\langle v \rangle_a)$ such that the top label of the transistor of $[\Phi]$ contains no subword that is the left side of some relation in \mathcal{P} . Consider the cover of $\text{lk}_\downarrow(\langle v \rangle_a)$ by the stars of the vertices in the collection $\{[\Psi_i] \mid i = 1, \dots, p\}$ of vertices that overlap with $[\Phi]$.

A first step is to show that the intersection $|\text{st}([\Psi_{i_1}])| \cap \dots \cap |\text{st}([\Psi_{i_m}])|$ is $(n - 1)$ -connected, for any choice $[\Psi_{i_1}], \dots, [\Psi_{i_m}]$ of vertices as above ($2 \leq m \leq p$). To do this, it is enough, by Lemma 6 ii), to show that $|\text{lk}([\Psi_{i_1}])| \cap \dots \cap |\text{lk}([\Psi_{i_m}])|$ is $(n - 1)$ -connected. Since every vertex $[\Psi_{i_j}], \dots, [\Psi_{i_m}]$ overlaps with $[\Phi]$, no more than $3R - 2$ contacts are used by the $[\Psi_{i_j}], j = 1, \dots, m$, and these contacts occur consecutively (in the cyclic sense). The subcomplex $|\text{lk}([\Psi_{i_1}])| \cap \dots \cap |\text{lk}([\Psi_{i_m}])|$ is the same as the full subcomplex of $\text{lk}_\downarrow(\langle v \rangle_a)$ generated by the vertices which do not overlap with any of the $[\Psi_{i_j}], j = 1, \dots, m$. Thus, by the choice of $[\Phi]$, $|\text{lk}([\Psi_{i_1}])| \cap \dots \cap |\text{lk}([\Psi_{i_m}])|$

is homeomorphic to $|\text{lk}_\downarrow(v_1)|$, where v_1 is a word satisfying

$$\begin{aligned} \ell(v_1) &\geq (n + 2)N + (n + 3)(R - 1) - (3R - 2) \\ &= (n + 2)N + nR - (n + 1) \\ &\geq (n + 1)N + n(R - 1). \end{aligned}$$

Proposition 4.11 from [11] says that whenever $\ell(v) \geq (n + 2)N + (n + 1)(R - 1)$, $|\text{lk}_\downarrow(v)|$ is n -connected. This implies that $|\text{lk}_\downarrow(v_1)|$ is $(n - 1)$ -connected. Thus $|\text{st}([\Psi_{i_1}]) \cap \dots \cap \text{st}([\Psi_{i_m}])|$ is $(n - 1)$ -connected for any choice of vertices $[\Psi_{i_1}]$, \dots , $[\Psi_{i_m}]$ which overlap with $[\Phi]$. Part ii) now follows by applying the Nerve Theorem to the cover $\{\text{st}([\Psi_i]) \mid i = 1, \dots, p\}$. Note that the nerve of this cover is a simplex. \square

Now, since $\ell(\langle w_n \rangle_b) \rightarrow \infty$ and $\ell(\langle w_n \rangle_a) \rightarrow \infty$ as $n \rightarrow \infty$, it follows that $\mathcal{D}_b(\mathcal{P}, w)$ and $\mathcal{D}_a(\mathcal{P}, w)$ are of type F_∞ , by Propositions 2, 3, and 7. Theorem 1 follows.

4. THE RATIONAL HOMOLOGY OF PICTURE GROUPS

4.1. Vanishing of rational homology in high dimensions. Note that only the picture groups $\mathcal{D}_b(\mathcal{P}, w)$ are under consideration in the remaining sections, and not the annular groups $\mathcal{D}_a(\mathcal{P}, w)$.

The following proposition will be fundamental to the homology calculations in the remainder of the paper. I am indebted to Ken Brown for the proof, which was relayed to me by the referee.

Proposition 8. *If X is a contractible G -complex with finite cell stabilizers, then $C_*(X, \mathbb{Q})$, the cellular chain complex of X with coefficients in \mathbb{Q} , is a projective resolution of $\mathbb{Q}G$.*

Proof. Since the complex is clearly a resolution, we only need to make sure that the modules $C_p(X, \mathbb{Q})$ are projective for each dimension p . It follows from formula (7.5) on page 173 in [4] that

$$C_p(X, \mathbb{Q}) = \bigoplus_{\sigma \in X_p} \mathbb{Q}_\sigma \uparrow_{G_\sigma}^G.$$

Here the modules \mathbb{Q}_σ are projective G_σ -modules by Exercise 5 on page 30 in [4]. Since induction takes projectives to projectives, the modules $\mathbb{Q}_\sigma \uparrow_{G_\sigma}^G$ are projective. \square

Now suppose that \mathcal{P} is an arbitrary semigroup presentation. Since $\mathcal{D}_b(\mathcal{P}, w)$ acts with finite cell stabilizers on $\tilde{K}_b(\mathcal{P}, w)$ (see the proof of Theorem 4.9 in [10]), the rational homology of $\mathcal{D}_b(\mathcal{P}, w)$ may be computed from $\mathbb{Q} \otimes_{\mathcal{D}_b(\mathcal{P}, w)} \mathcal{C}_*(\tilde{K}_b(\mathcal{P}, w))$, where $\mathcal{C}_*(\tilde{K}_b(\mathcal{P}, w))$ is the cellular chain complex associated to the cell complex $\tilde{K}_b(\mathcal{P}, w)$.

Proposition 9. *Let C be a cube of $\tilde{K}_b(\mathcal{P}, w)$ in standard form.*

- i) *If there are two white transistors of C labelled by the same relator, then there is $\Delta_1 \in \mathcal{D}_b(\mathcal{P}, w)$ such that Δ_1 maps C to itself and reverses its orientation.*
- ii) *If no two of the white transistors of C are labelled by the same relator and $\Gamma_b(\mathcal{P})$ is terminating, then any element in $\mathcal{D}_b(\mathcal{P}, w)$ that fixes C , fixes C pointwise.*

Proof. i) Let $C = C_b(\Delta, \Psi)$, where Ψ is in standard form and the bottom label of Δ is w_1 ; suppose that T_1 and T_2 are white transistors of Ψ which are labelled by the same relation. Assume that T_1 (thus also T_2) has m top contacts. Throughout the following argument, “the i th top contact” means “the i th top contact from the left”. The *number* of a contact is its number when counting from the left along the top or bottom of a given transistor (or along the frame).

If a wire w meets the k th top contact of T_1 , let i_k be the number of the top contact of Ψ at the other end of w . Similarly, if a wire w meets the k th top contact of T_2 , let j_k be the number of the top contact of Ψ at the other end of w . Note that the i_k th and j_k th top contacts of Ψ necessarily have the same label.

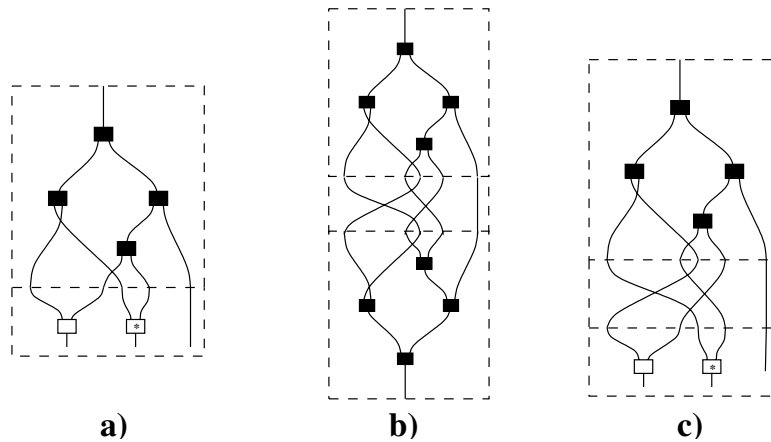


FIGURE 6.

Let Φ be the (w_1, w_1) permutation picture such that the i_k th top contact is connected by a wire to the j_k th bottom contact and the j_k th top contact is connected by a wire to the i_k th bottom contact. All other wires in Φ should connect the ℓ th top contact with the ℓ th bottom contact.

Now consider the action of $\Delta_1 = \Delta \circ \Phi \circ \Delta^{-1}$ on $C = C_b(\Delta, \Psi)$:

$$(\Delta \circ \Phi \circ \Delta^{-1}) \cdot C = C_b(\Delta \circ \Phi, \Psi).$$

Thus, by the definition of Φ , the action of Δ_1 simply interchanges the transistors T_1 and T_2 , leaving the cube C fixed. The operation of interchanging T_1 and T_2 corresponds to an orientation-reversing homeomorphism of $|C|$.

The idea of part i) is illustrated in Figure 6. The top picture in a) is Δ and the bottom picture is (the equivalence class of) Ψ , with one white transistor marked to indicate the orientation. The picture in b) is $\Delta_1 = \Delta \circ \Phi \circ \Delta^{-1}$ for the choice of Δ and Ψ in a). The cube depicted in c) is $C_b(\Delta \circ \Phi, \Psi)$. Note that the cubes in a) and c) are the same, but any isomorphism between the two underlying pictures must carry the unmarked transistor to the marked transistor. This implies that Δ_1 reverses the orientation.

ii) Since $\Gamma_b(\mathcal{P})$ is terminating, the cube C has a unique vertex $[\Delta]$ with the property that $\langle w_1 \rangle_b$, the bottom label of $[\Delta]$, is larger than the bottom labels of all of the other vertices of C in the partial order on b -words. (In fact, the vertex

$[\Delta]$ may be obtained by removing the white transistors of C .) Since the action of $\mathcal{D}_b(\mathcal{P}, w)$ on $\tilde{K}_b(\mathcal{P}, w)$ preserves the bottom labels of vertices, any group element Δ' which fixes the cube C must map the vertex $[\Delta]$ to itself. Thus $\Delta' \circ \Delta = \Delta \circ \Theta_1$, where Θ_1 is a permutation picture, so $\Delta' = \Delta \circ \Theta_1 \circ \Delta^{-1}$.

Now choose any vertex of C that is adjacent to $[\Delta]$. Such a vertex has the form $[\Delta \circ \Psi_1]$, where Ψ_1 is a thin picture with just one transistor. One wants to show that the vertex $[\Delta \circ \Psi_1]$ is fixed by Δ' . If not, then $(\Delta \circ \Theta_1 \circ \Delta^{-1}) \circ \Delta \circ \Psi_1 = \Delta \circ \Psi_2 \circ \Theta_2$, for some thin picture Ψ_2 with just one transistor and some permutation picture Θ_2 . Note that, by the assumption on C , the transistor in Ψ_1 has a different label from the transistor in Ψ_2 . But then $\Theta_1 \circ \Psi_1 = \Psi_2 \circ \Theta_2$, and the pictures $\Theta_1 \circ \Psi_1$ and $\Psi_2 \circ \Theta_2$ each contain just one transistor, and these transistors are labelled by different relations. This is a contradiction. It follows that the vertex $[\Delta]$ and each vertex of C adjacent to $[\Delta]$ is fixed by the action of Δ' . This means that Δ' must fix the whole cube. \square

Lemma 10. *Let $C_b(\Delta, \Psi)$ be an oriented cube of $\tilde{K}_b(\mathcal{P}, w)$, where Ψ is in standard form. If two (or more) transistors of Ψ are labelled by the same relation, then $1 \otimes_{\mathcal{D}_b(\mathcal{P}, w)} C_b(\Delta, \Psi) = 0$ in $\mathbb{Q} \otimes_{\mathcal{D}_b(\mathcal{P}, w)} \mathcal{C}_*(\tilde{K}_b(\mathcal{P}, w))$.*

Proof. By Proposition 9, there is some $\Delta_1 \in \mathcal{D}_b(\mathcal{P}, w)$ such that Δ_1 reverses the orientation of $C_b(\Delta, \Psi)$. Write \otimes in place of $\otimes_{\mathcal{D}_b(\mathcal{P}, w)}$:

$$(1 \otimes C_b(\Delta, \Psi)) = ((1 \cdot \Delta_1) \otimes C_b(\Delta, \Psi)) = (1 \otimes (\Delta_1 \cdot C_b(\Delta, \Psi))) = (1 \otimes -C_b(\Delta, \Psi)).$$

From this, it follows that $(1 \otimes C_b(\Delta, \Psi)) = 0$. \square

One has this corollary:

Theorem 11. *If the semigroup presentation \mathcal{P} has only n relations, then $H_m(\mathcal{D}_b(\mathcal{P}, w); \mathbb{Q}) = 0$ when $m > n$.*

Proof. If $m > n$, then any chain in $\mathbb{Q} \otimes_{\mathcal{D}_b(\mathcal{P}, w)} \mathcal{C}_m(\tilde{K}_b(\mathcal{P}, w))$ can be represented as a linear combination of elements of the form $1 \otimes C_b(\Delta, \Psi)$, where one can assume that Ψ is in standard form, and thus two or more of the transistors of Ψ are labelled by the same relation. It follows from Lemma 10 that all such elements are 0, so $\mathbb{Q} \otimes_{\mathcal{D}_b(\mathcal{P}, w)} \mathcal{C}_m(\tilde{K}_b(\mathcal{P}, w)) = 0$. \square

The rest of the paper is devoted to showing that the above theorem is also true for $m = n$, in case the rewrite system $\Gamma_b(\mathcal{P})$ is terminating.

4.2. Morse theory. This section contains a simple version of combinatorial Morse theory, based on the work of Forman [13] and Brown’s idea of a collapsing scheme from [7]. The idea here is to algebraicize their geometric arguments so that the theory will apply to the chain complex from the previous section.

Let $(\{C_n\}_{n \geq 0}, \partial_n)$ be a chain complex of free R -modules, where R is a commutative ring with 1. (In this paper, $R = \mathbb{Q}$.) Call the chain map $f : C_* \rightarrow C_*$ a *discrete flow* on $(\{C_n\}_{n \geq 0}, \partial_n)$ if:

- i) f is chain homotopic to the identity, and
- ii) for any chain $c \in C_*$, there is $n \in \mathbb{N}$ such that $f^n(c) = f^{n+1}(c)$.

If f is a discrete flow, then property ii) implies that $f^\infty(c)$ is well-defined for any chain c . The map $f^\infty : C_* \rightarrow C_*$ is a chain map. Let $I_f(C)_*$ denote the R -module of f -invariant chains in C_* . It is clear that $I_f(C)_* = f^\infty(C_*)$.

Lemma 12. *The map $f^\infty : C_* \rightarrow I_f(C)_*$ induces an isomorphism on homology.*

Proof. Since f^∞ is a chain map, it maps cycles to cycles and boundaries to boundaries. The map f^∞ is also surjective on cycles, since, for any cycle $c \in I_f(C)_*$, c is also a cycle in C_* , and $f^\infty(c) = c$.

Finally, one needs to show that a chain c is mapped to a boundary if and only if c is itself a boundary. If $f^\infty(c) = \partial d$, then $f^\infty(c) = f^\infty \circ f^\infty(c) = \partial f^\infty(d)$. There is some $m \in \mathbf{N}$ such that $f^\infty(c) = f^m(c)$ and $f^\infty(d) = f^m(d)$. This implies that $f^m(c) = \partial f^m(d)$, so $f^m(c) = 0$ on the level of homology. Since f^m induces an isomorphism on homology, c is a boundary. \square

Choose a basis B_n for each C_n . Let $\pm B_n = B_n \cup -B_n \cup \{0\}$.

Let $V : B_* \rightarrow \pm B_{*+1}$ satisfy the following conditions:

- a) If $V(b_1) = \pm V(b_2) \neq 0$, then $b_1 = b_2$;
- b) $VV = 0$;
- c) for any $b \in B_n (n \geq 0)$, either $V(b) = 0$ OR b occurs in $\partial_{n+1}V_n b$ with the coefficient -1 ;
- d) for $b_1, b_2 \in B_n$, 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 $b_1 \neq b_2$. One requires:
 - d.1) that for each $b_0 \in B_n$, there are only finitely many $b_1 \in B_n$ satisfying $b_0 > b_1$, i.e., satisfying condition d), and
 - d.2) that the irreflexive, transitive closure of $>$, also denoted $>$, be a strict partial order with no infinite descending chains.

The map V can be extended to a map from C_* to C_{*+1} by linearity.

Define the height $h(b)$ of a basis element $b \in B_n$ to be the length of the largest strictly descending chain $b > b_1 > \dots > b_n$. Thus, if b is minimal in the partial order $>$, then $h(b) = 1$. The height of a basis element is well defined because of the following lemma, which was provided by the referee.

Lemma 13. *Let M be a set with a binary relation \rightarrow satisfying the following two conditions:*

- (1) *There is no infinite chain $m_1 \rightarrow m_2 \rightarrow m_3 \rightarrow \dots$*
- (2) *For every $m \in M$, there are only finitely many $m' \in M$ with $m \rightarrow m'$.*

Then the irreflexive transitive closure of \rightarrow is a strict partial order on M such that every $m \in M$ has a well-defined height

$$h(m) := \max\{n \mid \text{there is a chain } m = m_1 \rightarrow \dots \rightarrow m_n\}.$$

Proof. The first hypothesis implies that the transitive closure of \rightarrow is a strict partial order on M . Thus, we only need to argue that heights are well defined and finite.

Let T_m be the rooted tree whose vertices are finite chains starting at m and such that the parent of the chain

$$m = m_1 \rightarrow m_2 \rightarrow \dots \rightarrow m_{n-1} \rightarrow m_n$$

is the chain

$$m = m_1 \rightarrow m_2 \rightarrow \dots \rightarrow m_{n-1}.$$

The second hypothesis implies that T_m is locally finite. By the first hypothesis, T_m does not contain any infinite path. Thus, by König's Lemma, the tree T_m is finite. \square

If $c \in C_*$ is any non-zero chain, let $h(c) = \max\{h(b) \mid b \in B_* \text{ occurs in } c \text{ with a non-zero coefficient}\}$. Call a basis element $b \in B_n$ *critical* if $V(b) = 0$ and $b \notin \text{Im } V$. Call a basis element $b \in B_n$ *collapsible* if $b \in \text{Im } V$. Call a basis element $b \in B_n$

redundant if $V(b) \neq 0$. It is clear from the definitions that any given basis element is critical, collapsible, or redundant, but not more than one of these.

Let $g = \partial V + V\partial$ and $f = 1 + g$.

Lemma 14. *Let $b \in B_n$.*

i) *If b is critical, then $g(b) = 0$ or $h(g(b)) < h(b)$, and every basis element occurring in $g(b)$ is collapsible.*

ii) *If b is collapsible, then $f(b) = 0$ or $h(f(b)) < h(b)$ and every basis element occurring in $f(b)$ is collapsible.*

iii) *If b is redundant, then $f(b) = 0$ or $h(f(b)) < h(b)$.*

Proof. i) If b is critical, then $g(b) = V\partial b$. Assume $g(b) \neq 0$. All of the basis elements occurring in $V\partial b$ are of course collapsible. Thus, if $b_1 \in B_*$ occurs in $V\partial b$, $b \neq b_1$ and, by definition, $b > b_1$. It follows easily that $h(b_1) < h(b)$.

ii) Assume $f(b) \neq 0$. If b is collapsible, then $b = V(b_1)$ for some $b_1 \in \pm B_{n-1}$. It follows that b_1 occurs in $\partial V(b_1)$ with a coefficient of -1 . Now

$$\begin{aligned} f(b) &= (1 + \partial V + V\partial)(V(b_1)) = V(b_1) + V\partial V(b_1) \\ &= V(b_1) + V([\partial V b_1 + b_1] - b_1) = V[\partial V b_1 + b_1], \end{aligned}$$

where $\partial V b_1 + b_1$ contains no occurrence of b_1 after cancelling. By property a) of V , b therefore does not occur in $f(b)$. Each basis element occurring in $f(b)$ therefore occurs in $V\partial(b)$ so that, for any basis element b_2 occurring in $f(b)$, $b > b_2$. It follows that $h(b) > h(f(b))$. By the above calculation, $f(b) \in \text{Im } V$, so each basis element in $f(b)$ is collapsible.

iii) Assume $f(b) \neq 0$. If b is redundant, then b occurs in $\partial V b$ with a coefficient of -1 . Thus the sum $b + \partial V b$ contains no occurrence of the basis element b after collecting terms, and each basis element b_1 that does occur is such that $b > b_1$. It follows that $h(b) > h(b + \partial V b)$.

Each basis element occurring in $V\partial b$ is collapsible, so b does not occur in $V\partial b$. If b_1 occurs in $V\partial b$, then $b > b_1$ by the definition of $>$, and it follows that $h(b) > h(V\partial b)$. Therefore $h(b) > h(f(b))$. \square

Corollary 15. *An f -invariant chain contains no occurrence of any redundant basis element.* \square

Proposition 16. *If $f = 1 + \partial V + V\partial$ (V as above), then f is a discrete flow.*

Proof. Since f is clearly chain homotopic to the identity, it is sufficient to show that, for any basis element b , there is $n \in \mathbb{N}$ such that $f^n(b) = f^{n+1}(b)$.

The proof is by induction on $h(b)$. If $h(b) = 1$, then, by Lemma 14, $f(b) = 0$ if b is redundant or collapsible, or $f(b) = b$ if b is critical. In either case $f(b) = f(f(b))$. If $h(b) > 1$ and b is redundant or collapsible, then, again by Lemma 14, $h(f(b)) < h(b)$ and the conclusion that $f^n(b) = f^{n+1}(b)$ for some n follows easily by induction.

The only case left to consider occurs when $h(b) > 1$ and b is critical.

$$f^2(b) - f(b) = f(f(b) - b) = fg(b).$$

By Lemma 14, only collapsible basis elements occur in $fg(b)$. It then follows easily, again from Lemma 14, that $f^m(fg(b)) = 0$, for some m . Thus $f^{m+2}(b) = f^{m+1}(b)$. \square

Let $M_f(C)_n$ be the free R -submodule of C_n generated by the critical basis elements of C_n . Let $\Pi_{M_f(C)_*}$ denote the projection onto $M_f(C)_*$.

Proposition 17. $f^\infty : M_f(C)_* \rightarrow I_f(C)_*$ and $\Pi_{M_f(C)_*} : I_f(C)_* \rightarrow M_f(C)_*$ are mutually inverse isomorphisms.

Proof. Choose a basis element $b \in B_n$ in $M_f(C)_*$. Lemma 14 implies that $f^\infty(b) = b+$ collapsible elements, so $\Pi_{M_f(C)_*} f^\infty(b) = b$.

Conversely, choose a chain $c \in I_f(C)_*$. The chain c contains no occurrences of redundant basis elements by Corollary 15. Thus $\Pi_{M_f(C)_*}(c) - c$ contains occurrences only of collapsible basis elements. It follows from Lemma 14 that

$$f^\infty(\Pi_{M_f(C)_*}(c) - c) = 0.$$

Therefore $f^\infty \Pi_{M_f(C)_*}(c) - f^\infty(c) = 0$, so $f^\infty \Pi_{M_f(C)_*}(c) = f^\infty(c) = c$. □

Now identify $M_f(C)_*$ with $I_f(C)_*$ by the isomorphism f^∞ . This identification gives the R -modules $M_f(C)_*$ the structure of a chain complex $(M_f(C)_*, \tilde{\partial}_*)$ in which $\tilde{\partial}_* = \Pi_{M_f(C)_{*-1}} \partial f^\infty$. This complex is called the *Morse complex associated to the flow f* . It is now easy to prove:

Theorem 18. *The Morse complex $(M_f(C)_*, \tilde{\partial}_*)$ and (C_n, ∂_n) have isomorphic homology groups.* □

4.3. Terminating semigroup presentations and vanishing rational homology. One needs an explicit description of the chain complex $\mathbb{Q} \otimes_{\mathcal{D}_b(\mathcal{P}, w)} \mathcal{C}_*(\tilde{K}_b(\mathcal{P}, w))$.

Lemma 19. *Suppose that $\Gamma_b(\mathcal{P})$ is terminating. Let $C_b(\Delta_1, \Psi_1)$ and $C_b(\Delta_2, \Psi_2)$ be cubes of $\tilde{K}_b(\mathcal{P}, w)$, where Ψ_1 and Ψ_2 are in standard form. These cubes are in the same orbit under the action of $\mathcal{D}_b(\mathcal{P}, w)$ if and only if there are permutation pictures Φ_1 and Φ_2 such that $\Phi_1 \circ \Psi_1 \circ \Phi_2 = \Psi_2$.*

Proof. Suppose first that there are permutation pictures Φ_1 and Φ_2 such that $\Phi_1 \circ \Psi_1 \circ \Phi_2 = \Psi_2$. Let $\Delta_3 = \Delta_2 \circ \Phi_1 \circ \Delta_1^{-1}$.

$$\begin{aligned} \Delta_3 \cdot C_b(\Delta_1, \Psi_1) &= \Delta_3 \cdot C_b(\Delta_1 \circ \Phi_1^{-1}, \Phi_1 \circ \Psi_1) \\ &= \Delta_3 \cdot C_b(\Delta_1 \circ \Phi_1^{-1}, \Phi_1 \circ \Psi_1 \circ \Phi_2) \\ &= C_b(\Delta_2, \Psi_2). \end{aligned}$$

Conversely, if $C_b(\Delta_1, \Psi_1)$ and $C_b(\Delta_2, \Psi_2)$ are in the same orbit, and Ψ_i is in standard form, for $i = 1, 2$, then Ψ_1 and Ψ_2 must have the same number of transistors of any given label. (Recall that the label of a transistor T is a pair $(u_T, v_T) \in \Sigma^+$, where u_T is the top label and v_T is the bottom label.) Moreover, for $i = 1, 2$, $C_b(\Delta_i, \Psi_i)$ has a unique vertex, namely $[\Delta_i]$, such that the bottom label $\langle u_i \rangle_b$ of $[\Delta_i]$ is larger in the partial order on b -words than the bottom label of any other vertex in $C_b(\Delta_i, \Psi_i)$. There is also, for $i = 1, 2$, a unique vertex, namely $[\Delta_i \circ \Psi_i]$ of $C_b(\Delta_i, \Psi_i)$ with a bottom label $\langle v_i \rangle_b$ that is smaller than that of any other vertex in $C_b(\Delta_i, \Psi_i)$. Since the action of $\mathcal{D}_b(\mathcal{P}, w)$ on $\tilde{K}_b(\mathcal{P}, w)$ preserves the bottom label of any vertex, it follows that $\langle u_1 \rangle_b = \langle u_2 \rangle_b$ and $\langle v_1 \rangle_b = \langle v_2 \rangle_b$.

All of this information together implies that the transistors in Ψ_1 and Ψ_2 have the same labels (and these labels occur with the same multiplicities), and the same is true of the wires of Ψ_1 and Ψ_2 . Thus, the thin pictures Ψ_1 and Ψ_2 are completely determined up to the attaching maps of the wires to the frame. It follows that $\Phi_1 \circ \Psi_1 \circ \Phi_2 = \Psi_2$, for appropriate permutation diagrams Φ_1 and Φ_2 . □

Now consider the effect of tensoring $\mathcal{C}_*(\tilde{K}_b(\mathcal{P}, w))$ with \mathbb{Q} over $\mathcal{D}_b(\mathcal{P}, w)$. By Lemma 10, the only (potentially) non-zero terms of $\mathbb{Q} \otimes_{\mathcal{D}_b(\mathcal{P}, w)} \mathcal{C}_*(\tilde{K}_b(\mathcal{P}, w))$ (hereafter \mathcal{C}_*) are sums of rational multiples of elements of the form $1 \otimes C_b(\Delta, \Psi)$, where $C_b(\Delta, \Psi)$ is an oriented cube, Ψ is in standard form, and no two transistors of Ψ are labelled by the same relation. Lemma 19 shows that two such cubes $C_b(\Delta_1, \Psi_1)$ and $C_b(\Delta_2, \Psi_2)$ are in the same orbit if and only if $\Phi_1 \circ \Psi_1 \circ \Phi_2 = \Psi_2$. Proposition 9 ii) shows that the effect of taking the tensor product is simply to identify the basis elements corresponding to two such cubes, without, for instance, making either of these basis elements 0. The conclusion is that a basis for \mathcal{C}_* consists of $\{1 \otimes C_b(\Delta_i, \Psi_i)\}$, where the $C_b(\Delta_i, \Psi_i)$ are a choice of oriented cubes, any two of which are from different orbits, with the additional property that the Ψ_i are in standard form, and no two transistors of a given Ψ_i are labelled by the same relation.

If $\Gamma_b(\mathcal{P})$ is terminating, define the *standard basis* of \mathcal{C}_* to be the set of equivalence classes, denoted $\langle \rangle$, of thin pictures in standard form, where $\langle \Psi_1 \rangle = \langle \Psi_2 \rangle$ if there are permutation pictures Φ_1 and Φ_2 such that $\Phi_1 \Psi_1 \Phi_2 = \Psi_2$, and no two transistors of a given Ψ_i are labelled by the same relation. A member $\langle \Psi \rangle$ of the standard basis can also be thought of as a subset of $\mathcal{R} \subseteq (\Sigma^+ \times \Sigma^+)$, which represents the labelling of the transistors in Ψ , together with a b -word $\langle u \rangle_b$, which represents the labellings of the wires that do not touch the transistors. There is a one-to-one correspondence f between the basis elements $(1 \otimes C_b(\Delta, \Psi))$ mentioned in the previous paragraph and the standard basis, given by $f(1 \otimes C_b(\Delta, \Psi)) = \langle \Psi \rangle$. Identify the two sets by the map f . With this identification, the boundary of $\langle \Psi \rangle$, denoted $\partial \langle \Psi \rangle$, is a well-defined notion. In the remainder of the argument, it will never be necessary to compute the boundary of $\langle \Psi \rangle$ explicitly; it is sufficient to know which terms occur in $\partial \langle \Psi \rangle$ with a non-zero coefficient. (In fact, the coefficient of any term will be ± 1 or 0.) These terms may all be described as follows: Suppose that $L = \{(u_1, v_1), \dots, (u_k, v_k)\}$ is the set of labels of the transistors of Ψ , and the remaining wires are represented by the b -word $\langle u \rangle_b$. Remove one of the relations (u_i, v_i) from L , and replace the b -word $\langle u \rangle_b$ by $\langle uu_i \rangle_b$ or $\langle uv_i \rangle_b$. The element $\langle \Psi' \rangle$ represented by the new data occurs with a non-zero coefficient in $\partial \langle \Psi \rangle$.

Let Ψ be a member of the standard basis. (The brackets \langle and \rangle will be omitted hereafter.) The relation $u \rightarrow v$ is *applicable* to Ψ if there is some word \bar{w} in the b -word associated to Ψ such that $\bar{w} = w_1 u w_2$ in Σ^+ , where the words w_1 and w_2 may be empty. To *apply* this relation to Ψ , add $u \rightarrow v$ onto the list of defining relations for Ψ and let $\langle w_1 w_2 \rangle_b$ be the new b -word associated to the new standard basis element Ψ' . The inverse operation is called *removing* a relation.

Choose some injective function $rk : \mathcal{R} \rightarrow \mathbb{N}$. The *rank* of a relation r is $rk(r)$. The rank of a transistor is the rank of its label.

If there is some relation r such that r is applicable to Ψ and $rk(r) < rk(s)$ for any relation s already present in Ψ , then choose the minimal relation r with this property and apply it to Ψ to get $\pm V(\Psi)$. (It is always possible to choose the sign of $V(\Psi)$ so that condition c) of the previous subsection is satisfied, so it will be dropped hereafter.) Otherwise, let $V(\Psi) = 0$.

Lemma 20. *V satisfies properties a), b), c), and d).*

Proof. a) Suppose that Ψ_1 and Ψ_2 are elements of the standard basis, and $V(\Psi_1) = V(\Psi_2) \neq 0$. It follows from the definition of V that, for $i = 1, 2$, Ψ_i may be obtained from $V(\Psi_i)$ by removing the transistor of lowest rank in $V(\Psi_i)$. Since

$V(\Psi_1) = V(\Psi_2)$, the results of removing the lowest-rank transistors from each must be the same. Thus $\Psi_1 = \Psi_2$.

b) is easy to check, and c) is true by definition.

d.1) It is enough for this argument to consider only the primitive relation $>$ (as opposed to its transitive closure). Consider a sequence $\Psi_1 > \Psi_2 > \Psi_3 > \Psi_4 > \dots$ of elements of the standard basis. It is clear that $V(\Psi_i)$ and Ψ_i have the same top label. There are two possibilities for the top label of Ψ_{i+1} : 1. the top label of Ψ_{i+1} is the same as the top label of Ψ_i , or 2. the equivalence class of the top label of Ψ_{i+1} is less than that of Ψ_i in the partial order on b -words. In any infinite, strictly descending sequence $\Psi_1 > \Psi_2 > \dots > \Psi_n > \dots$ of members of the standard basis, at most finitely many b -words may occur as top labels of the Ψ_i s, since the rewrite system $\Gamma_b(\mathcal{P})$ is terminating. Thus, to rule out any such sequence, it is enough to consider the case in which all standard basis elements occurring in $\Psi_1 > \dots > \Psi_n > \dots$ have the same top label.

In this case, each Ψ_i is determined by the ranks of its relations. Thus each Ψ_i may be identified with an m -tuple $(a_1, a_2, \dots, a_m) \in \mathbb{N}$, for some m , where the a_j s are the ranks of the relations of Ψ_i , arranged in ascending order. (Note that each Ψ_i has the same number of relations, namely m .) Now, if $\Psi_i > \Psi_{i+1}$, the first coordinate of the m -tuple for Ψ_i is strictly larger than that of Ψ_{i+1} . This obviously implies that there can be no infinite descending chain, since the entries of these m -tuples are natural numbers.

d.2) Since the semigroup presentation \mathcal{P} has only finitely many relations and each picture Ψ contains only finitely many wires and transistors, this condition follows easily from the definition of $<$. \square

Theorem 21. *If $\Gamma_b(\mathcal{P})$ is a terminating rewrite system and the presentation \mathcal{P} has n relations, then $H_j(D_b(\mathcal{P}, w); \mathbb{Q}) = 0$, for $j \geq n$.*

Proof. Under the given hypotheses, an element Ψ of the n -dimensional standard basis involves all of the relations in \mathcal{P} . Consider the formal vertex Ψ' formed by removing from Ψ the transistor corresponding to the relation of lowest rank in \mathcal{P} . By the definition of V , $V(\Psi') = \Psi$, which implies that Ψ is not critical. It follows from Theorem 18 that $H_n(D_b(\mathcal{P}, w); \mathbb{Q}) = 0$.

For $j > n$, the statement follows from Theorem 11. \square

REFERENCES

- [1] Bestvina, M.; Brady, N. Morse theory and finiteness properties of groups *Invent. Math.* **129** (1997), no. 3, 445-470. MR1465330 (98i:20039)
- [2] Björner, A. Topological Methods in Handbook of Combinatorics, edited by R.L. Graham, M. Grötschel and L. Lovasz, Elsevier Science B.V., Amsterdam, MIT Press, Cambridge, MA, 1995 Vol II pp. i-cii and 1019-2198. MR1373655 (96h:05001)
- [3] Bridson, M. R.; Haefliger, A. Metric Spaces of Non-Positive Curvature, *Springer-Verlag Berlin Heidelberg* 1999 xxi+643 pp. MR1744486 (2000k:53038)
- [4] Brown, K.S. Cohomology of Groups *Springer-Verlag, New York Inc.* 1982 x+306 pp. MR0672956 (83k:20002)
- [5] Brown, K. S. Finiteness properties of groups. Proceedings of the Northwestern conference on cohomology of groups (Evanston, Ill., 1985) *J. Pure Appl. Algebra* **44** (1987), no. 1-3, 45-75. MR0885095 (88m:20110)
- [6] Brown, K. S. The geometry of finitely presented infinite simple groups. Algorithms and classification in combinatorial group theory (Berkeley, CA) 121-136, Math. Sci. Res. Inst. Publ., **23**, Springer, New York, 1992. MR1230631 (94f:20059)

- [7] Brown, K. S. The geometry of rewriting systems: a proof of the Anick-Groves-Squier theorem. Algorithms and classification in combinatorial group theory (Berkeley, CA) 137-163, *Math. Sci. Res. Inst. Publ.*, **23**, Springer, New York, 1992. MR1230632 (94g:20041)
- [8] Brown, K. S.; Geoghegan, R. An infinite-dimensional torsion-free FP_∞ group *Invent. Math.* **77** (1984), no. 2, 367-381. MR0752825 (85m:20073)
- [9] Cannon, J. W.; Floyd, W. J.; Parry, W. R. Introductory notes on Richard Thompson's groups *Enseign. Math. (2)* **42** (1996), no. 3-4, 215-256. MR1426438 (98g:20058)
- [10] Farley, D. S. Actions of Picture Groups on CAT(0) Cubical Complexes *to appear in Geometriae Dedicata*.
- [11] Farley, D. S. Finiteness and CAT(0) properties of Diagram Groups *Topology* **42** (2003), no.5, 1065-1082. MR1978047 (2004b:20057)
- [12] Farley, D. S. Proper Isometric Actions of Thompson's Groups on Hilbert Space *IMRN* 2003, no. 45, 2409-2414. MR2006480 (2004k:22005)
- [13] Forman, R. Morse theory for cell complexes *Adv. Math* **134** (1998), no. 1, 90-145. MR1612391 (99b:57050)
- [14] Greenberg, M.J.; Harper, J.R. Algebraic Topology, *Addison-Wesley Publishing Co.* 1981 xi+311 pp. MR0643101 (83b:55001)
- [15] Guba, V.S.; Sapir, M. V. Diagram groups *Mem. Amer. Math. Soc.* **130** (1997), no. 620, viii+117 pp. MR1396957 (98f:20013)
- [16] Guba, V.S.; Sapir, M.V. Diagram groups and directed 2-complexes: homotopy and homology. *preprint at front.math.ucdavis.edu*.
- [17] Guba, V. S.; Sapir, M. V. On subgroups of R. Thompson's group F and other diagram groups (Russian) *Mat. Sb.* **190** (1999), no. 8, 3-60. MR1725439 (2001m:20045)
- [18] Guba, V. S.; Sapir, M. V. Rigidity properties of diagram groups. International Conference on Geometric and Combinatorial Methods in Group Theory and Semigroup Theory (Lincoln, NE, 2000) *Internat. J. Algebra Comput.* **12** (2002), no. 1-2, 9-17. MR1902358 (2003e:20046)
- [19] Kilibarda, V. On the algebra of semigroup diagrams *Internat. J. Algebra Comput.* **7** (1997), no. 3, 313-338. MR1448329 (98j:20085)
- [20] McKenzie, R.; Thompson, R. J. An elementary construction of unsolvable word problems in group theory. In: Word Problems (Conf., Univ. California, Irvine, 1969, Boone, W.W., Cannonito, F.B., Lyndon, R.C. (eds.), *Studies in Logic and the Foundations of Mathematics*, vol. 71, pp. 457-478. North-Holland, Amsterdam 1973. MR0396769 (53:629)
- [21] Rourke, C.P.; Sanderson, B.J. Introduction to Piecewise-Linear Topology, *Springer-Verlag, Berlin, Heidelberg, New York* 1982 viii+123pp. MR0665919 (83g:57009)
- [22] Spanier, E.H. Algebraic Topology, *McGraw-Hill Book Co., New York-Toronto, Ont.-London* 1966 xiv+528 pp. MR0210112 (35:1007)

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS, URBANA, ILLINOIS 61801