ISSN 1472-2739 (on-line) 1472-2747 (printed)

Algebraic & Geometric Topology
Volume 5 (2005) 1075–1109
Published: 31 August 2005

Discrete Morse theory and graph braid groups

Daniel Farley
Lucas Sabalka

Abstract If Γ is any finite graph, then the *unlabelled configuration space of n points on* Γ, denoted $UC^n\Gamma$, is the space of n-element subsets of Γ. The *braid group of* Γ *on n strands* is the fundamental group of $UC^n\Gamma$.

We apply a discrete version of Morse theory to these $UC^n\Gamma$, for any n and any Γ, and provide a clear description of the critical cells in every case. As a result, we can calculate a presentation for the braid group of any tree, for any number of strands. We also give a simple proof of a theorem due to Ghrist: the space $UC^n\Gamma$ strong deformation retracts onto a CW complex of dimension at most k, where k is the number of vertices in Γ of degree at least 3 (and k is thus independent of n).

AMS Classification 20F65, 20F36; 57M15, 57Q05, 55R80

Keywords Graph braid groups, configuration spaces, discrete Morse theory

1 Introduction

Let Γ be a finite graph, and fix a natural number n. The *labelled configuration space* $C^n\Gamma$ 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* $UC^n\Gamma$ is the quotient of $C^n\Gamma$ by the action of the symmetric group S_n, where the action permutes the factors. The fundamental group of $UC^n\Gamma$ (respectively, $C^n\Gamma$) is the *braid group* (respectively, the *pure braid group*) *of* Γ *on n strands*, denoted $B_n\Gamma$ (respectively, $PB_n\Gamma$).

Configuration spaces of graphs occur naturally in certain motion planning problems. An element of $C^n\Gamma$ or $UC^n\Gamma$ may be regarded as the positions of n robots on a given track (graph) within a factory. A path in $C^n\Gamma$ represents a collision-free movement of the robots. Ghrist and Koditschek [15] used a repelling vector field to define efficient and safe movements for the fundamentally important case of two robots on a Y-shaped graph. Farber [11] has recently shown that

$TC(\mathcal{C}^n\Gamma) = 2k+1$, where "$TC$" denotes topological complexity, Γ is any tree, k is the number of vertices of Γ having degree at least 3, and $n \geq 2k$.

In [14], Ghrist showed that $\mathcal{C}^n\Gamma$ strong deformation retracts onto a complex X of dimension at most k, where k is the number of vertices having degree at least three in Γ (and thus is independent of n). If Γ is a *radial tree*, i.e., if Γ is a tree having only one vertex of degree more than 2, then $\mathcal{C}^n\Gamma$ strong deformation retracts on a graph. By computing the Euler characteristic of this graph, Ghrist computes the rank of the pure braid group on Γ as a free group. (See also [13], where the Euler characteristic of any configuration space of a simplicial complex is computed.)

Ghrist also conjectures in [14] that every braid group $B_n\Gamma$ is a *right-angled Artin group* – i.e. a group which has a presentation in which all defining relations are commutators of the generators. Abrams [2] revised the conjecture to apply only to planar graphs. These conjectures were one of our original reasons to investigate graph braid groups. However, Abrams [1] has shown that some graph braid groups, including a planar graph and a tree braid group, are not Artin right-angled. There are positive results: Crisp and Wiest [10] have recently shown that every graph braid group embeds in a right-angled Artin group.

We use Forman's discrete Morse theory [12] (see also [7]) in order to simplify any space $U\mathcal{C}^n\Gamma$ within its homotopy type. We are able to give a very clear description of the critical cells of a discretized version of $U\mathcal{C}^n\Gamma$ with respect to a certain discrete gradient vector field W (see Section 2 for definitions). By a theorem of Forman's ([12], page 107), a CW complex X endowed with a discrete gradient vector field W is homotopy equivalent to a complex with m_p cells of dimension p, where m_p is the number of cells of dimension p that are critical with respect to W. (A similar theorem was proved by Brown earlier; see page 140 in [7].)

Our classification of critical cells leads to a simple proof of Ghrist's theorem that the braid group of any radial tree is free, and we compute the rank of this free group as a function of the degree of the central vertex and the number of strands. This computation features an unusual application of the formula for describing the number of ways to distribute indistinguishable balls into distinguishable boxes. We also prove a somewhat strengthened version of Ghrist's strong deformation retraction theorem by a simple application of Morse-theoretic methods. Our dimension bounds resemble those of [20], where the homological dimension of the braid groups $B_n\Gamma$ was estimated. The strong deformation retract in our theorem has an explicit description in terms of so-called collapsible, redundant, and critical cells (see Section 2 for definitions).

The main theorem of the paper describes how to compute a presentation of $B_n\Gamma$, where n is an arbitrary natural number and Γ is any tree. We deduce our theorem with the aid of a particular discrete gradient vector field W and the "redundant 1-cells lemma", which greatly simplifies calculations. The generators in our presentation correspond to critical 1-cells of $UC^n\Gamma$, and the relators correspond to the critical 2-cells. The relators are all commutators, so that the rank of the abelianization of $B_n\Gamma$ is equal to the number m_1 of critical 1-cells. It follows easily that any CW complex homotopy equivalent to $UC^n\Gamma$ must have at least m_1 1-cells. It seems reasonable to guess that, more generally, any CW complex homotopy equivalent to $UC^n\Gamma$ must have at least m_p p-cells, where m_p is the number of p-cells that are critical with respect to W, although we don't prove this. The presentations that we obtain for $UC^n\Gamma$ do not look like presentations of right-angled Artin groups, but we don't know how to show that tree braid groups are not right-angled Artin groups by our methods.

It is possible to calculate group presentations for any graph braid group using the techniques of this paper, but we offer no general theorem here, since the resulting presentations are somewhat less than optimal – e.g., our calculations in some cases yield presentations of the free group which contain non-trivial relators. We hope to improve these presentations in the near future.

Our presentation is mostly self-contained. Section 2 contains a short exposition of Forman's Morse theory, which is sufficient for all of the applications in later sections. The central idea here is that of a discrete gradient vector field, which induces a classification of the cells into one of three mutually exclusive types: collapsible, redundant, and critical. Section 3 describes how to define a discrete gradient vector field on the configuration space of any graph, and gives a description of the collapsible, redundant, and critical cells with respect to the given discrete gradient vector field. Section 4 contains a calculation of the braid group of a radial tree, and a refined version of Ghrist's strong deformation retraction theorem. Section 5 contains the main theorem, about presentations of tree braid groups, with examples.

We would like to thank Ilya Kapovich and Robert Ghrist for participating in numerous discussions relating to this work. We thank Kim Whittlesey, Aaron Abrams, James Slack and the referee for reading an earlier version of the manuscript, and suggesting corrections. The first author thanks Tadeusz Januszkiewicz for telling him of the references [13] and [20].

2 Preliminary material

2.1 Discrete Morse theory

The fundamental idea of discrete Morse theory is that of a *collapse*. We take the definition from [8], page 14. If (X, Y) is a finite CW pair then X *collapses to* Y *by an elementary collapse* – denoted $X \searrow^{e} Y$ – if and only if

(1) $X = Y \cup e^{n-1} \cup e^{n}$ where e^{n} and e^{n-1} are open cells of dimension n and $n-1$, respectively, which are not in Y, and

(2) there exists a ball pair $(Q^{n}, Q^{n-1}) \approx (I^{n}, I^{n-1})$ and a map $\phi : Q^{n} \to X$ such that

 (a) ϕ is a characteristic map for e^{n}

 (b) $\phi \mid Q^{n-1}$ is a characteristic map for e^{n-1}

 (c) $\phi(P^{n-1}) \subseteq Y^{n-1}$, where $P^{n-1} = cl(\partial Q^{n} - Q^{n-1})$.

We say that X *collapses to* Y, and write $X \searrow Y$, if Y may be obtained from X by a sequence of elementary collapses.

Let X be a finite CW complex. Let K denote the set of open cells of X with K_p the set of open p-dimensional cells of X. For an open cell $\sigma \in K$, we write $\sigma^{(p)}$ to indicate that σ is of dimension p. We write $\sigma < \tau$ if $\sigma \neq \tau$ and $\sigma \subseteq \overline{\tau}$ (where $\overline{\tau}$ denotes the closure of τ). We write $\sigma \leq \tau$ if $\sigma = \tau$ or $\sigma < \tau$.

We will need to work with a special type of CW complex. From now on, every CW complex X we consider will have the following property: if $\sigma^{(p)}$ and $\tau^{(p+1)}$ are open cells of X and $\sigma^{(p)} \cap \overline{\tau^{(p+1)}} \neq \emptyset$, then $\sigma^{(p)} < \tau^{(p+1)}$. The importance of this assumption will become apparent in Proposition 2.2.

Suppose that $\sigma^{(p)}$ is a face of $\tau^{(p+1)}$ $(\sigma < \tau)$. Let B be a closed ball of dimension $p+1$, and let $h : B \to X$ be a characteristic map for τ. The cell σ is a *regular face* of τ if $h : h^{-1}\sigma \to \sigma$ is a homeomorphism, and $\overline{h^{-1}(\sigma)}$ is a closed p-ball.

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

(i) Each W_i is injective

(ii) If $W_i(\sigma) = \tau$, then σ is a regular face of τ.

(iii) $\mathrm{im}(W_i) \cap \mathrm{dom}(W_{i+1}) = \emptyset$.

This definition of a discrete vector field differs very slightly from Forman's ([12], page 130-131). Note that a *partial function* from a set A to a set B is simply a function defined on a subset of A.

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

Given any discrete gradient vector field W on X, there is an associated classification of cells in X into 3 types: redundant, collapsible, and critical (this terminology is partially borrowed from [12] as well as from Ken Brown, [7]). A cell $\sigma \in K$ is *redundant* if $\sigma \in \operatorname{dom} W$, *collapsible* if $\sigma \in \operatorname{im} W$, and *critical* otherwise. The *rank* of a cell c with respect to a discrete gradient vector field W is the length of the longest W-path $c = c_1, \ldots, c_r$ having the property that $c_i \neq c_j$ if $i \neq j$. Critical and collapsible cells all have rank 1, and redundant cells are of rank at least 2. If $c' < W(c)$ and $c \neq c'$, then clearly $rank(c') < rank(c)$.

In [12] (page 131), Forman shows that discrete gradient vector fields are precisely the discrete vector fields which arise from discrete Morse functions (in a manner he describes). We will work directly with discrete gradient vector fields, and never with an explicit discrete Morse function.

2.2 Monoid presentations and rewrite systems

An *alphabet* is simply a set Σ. The *free monoid on* Σ, denoted Σ^*, is the collection of all positive words in the generators Σ, together with the empty word, endowed with the operation of concatenation.

A *monoid presentation*, denoted $\langle \Sigma \mid \mathcal{R} \rangle$, consists of an alphabet Σ together with a collection $\mathcal{R}$ of ordered pairs of elements in Σ^*. An element of $\mathcal{R}$ should be regarded as an equality between words in Σ^*, but, in what follows, the order will matter.

A *rewrite system* Γ is an oriented graph. The vertices of Γ are called *objects* and the positive edges are called *moves*. If v_1 is the initial vertex of some positive edge in Γ and v_2 is the terminal vertex, then write $v_1 \to_\Gamma v_2$, or $v_1 \to v_2$ if the name of the specific rewrite system is clear from the context. An object is called *reduced* if it is not the initial vertex of any positive edge (move). The reflexive, transitive closure of $\to$ is denoted $\dot{\to}$.

A rewrite system is called *terminating* if every sequence $a_1 \to a_2 \to a_3 \to \ldots$ is finite. A rewrite system is called *confluent* if, whenever $a \dot{\to} b$ and $a \dot{\to} c$, there is an object d such that $b \dot{\to} d$ and $c \dot{\to} d$. A rewrite system is *locally confluent* if when $a \to b$ and $a \to c$, then there is d such that $b \dot{\to} d$ and $c \dot{\to} d$.

Lemma 2.1 [17] *Every terminating locally confluent rewrite system is confluent.*

A rewrite system is called *complete* if it is both terminating and confluent. For a complete rewrite system, it is not difficult to argue that every equivalence class of the equivalence relation generated by $\to$ has a unique reduced object.

Every monoid presentation $\langle \Sigma \mid \mathcal{R} \rangle$ has a natural rewrite system, called a *string rewriting system*, associated to it. The set of objects of this string rewriting system is the free monoid Σ^*. There is a move from $w_1 \in \Sigma^*$ to $w_2 \in \Sigma^*$ if $w_1 = ur_1v$ and $w_2 = ur_2v$ in Σ^*, where $u, v \in \Sigma^*$, $(r_1, r_2) \in \mathcal{R}$.

2.3 Discrete Morse theory and group presentations

Assume in this section that X is a finite connected CW complex with a discrete gradient vector field W. Let X'_n be the subcomplex of X consisting of the n-skeleton, but with the redundant n-cells removed. Let X''_n consist of the n-skeleton of X, but with the redundant and critical n-cells removed. The following proposition was essentially proved by Brown ([7], page 140) in the case of semi-simplicial complexes.

Proposition 2.2 *Consider the following filtration of X:*

$$\emptyset = X''_0 \subseteq X'_0 \subseteq \ldots \subseteq X''_n \subseteq X'_n \subseteq X''_{n+1} \subseteq \ldots.$$

(1) *For any n, X'_n is obtained from X''_n by attaching m_n n-cells to X''_n along their boundaries, where m_n is the number of critical n-cells of the discrete gradient vector field W.*

(2) *For any n, $X''_{n+1} \searrow X'_n$.*

Proof (1) This is obvious.

(2) Let $X_{n,k}$ be the subcomplex of X consisting of the entire $(n-1)$-skeleton, together with all n-cells of rank at most k, and all $(n+1)$-cells that are the images of such n-cells under the function W. Thus, for example, $X'_n = X_{n,1}$.

We claim that $X_{n,i+1} \searrow X_{n,i}$, for $i \in \mathbb{N}$.

Let c be an open n-cell of rank exactly $i+1$. Since W is injective, $W(c)$ cannot be the image under W of a cell of rank less than or equal to i, and so it lies outside of $X_{n,i}$. If $c \cap X_{n,i}$ is nonempty, it can only be because there is some open collapsible $(n+1)$-cell c' such that $c \cap \overline{c'}$ is nonempty and $c' = W(c'')$, for some open n-cell c'' of rank less than or equal to i. Given our standing assumption about the CW complex X (from Subsection 2.2), we thus know that $c < c'$. Now if $c = c_1, c_2, c_3, \ldots c_{i+1}$ is a W-path without any repeated cells (which exists because $rank(c) = i+1$) we have that $c'', c_1, c_2, \ldots, c_{i+1}$ is also a W-path without repetitions, since clearly $c'' \neq c_1$ and there are no non-stationary closed W-paths. This implies that the rank of c'' is at least $i+2$, a contradiction. It follows that the first part of the definition of a collapse is satisfied for the pair $(X_{n,i+1}, X_{n,i+1} - (c \cup W(c)))$.

Since c is a regular face of $W(c)$, there is a characteristic map $\phi : B^{n+1} \to X_{n,i+1}$ for $W(c)$ such that $\phi : \phi^{-1}(c) \to c$ is an homeomorphism and $\overline{\phi^{-1}(c)}$ is a closed ball. It follows easily that the second part of the definition of a collapse is satisfied for the same pair.

Repeating this argument for every $(n+1)$-cell of $X_{n,i+1}$, we eventually conclude that $X_{n,i+1} \searrow X_{n,i}$, since the individual collapses are compatible.

It follows that $X_{n,k} \searrow X'_n$ for any $k \in \mathbb{N}$. For k sufficiently large, $X_{n,k}$ consists of the entire n-skeleton, plus the collapsible $n+1$-cells. That is, $X''_{n+1} = X_{n,k} \searrow X'_n$. □

We collect a number of corollaries in the following proposition:

Proposition 2.3

1. *The inclusion of X'_n into X induces an isomorphism from $\pi_{n-1}(X'_n)$ to $\pi_{n-1}(X)$.*
2. *If X has no critical cells of dimension greater than k, then $X \searrow X'_k$.*
3. ([12], page 107) *X is homotopy equivalent to a CW complex with m_n cells of dimension n, where m_n is the number of critical n-cells in X.*
4. *The subcomplex of X generated by the collapsible and critical edges is connected.*
5. *The subcomplex of X generated by the collapsible edges and the 0-skeleton of X is a maximal forest.*

(6) *If there is only one critical* 0*-cell, then the graph consisting of (the closures of) the collapsible edges is a maximal tree in* X*.*

Proof

(1) Note that the map $\pi_{n-1}(X'_n) \to \pi_{n-1}(X)$ factors as

$$\pi_{n-1}(X'_n) \to \pi_{n-1}(X''_{n+1}) \to \pi_{n-1}(X).$$

The first map is an isomorphism, because $X''_{n+1} \searrow X'_n$. The second map is also an isomorphism, since X is obtained from X''_{n+1} by attaching cells of dimension greater than or equal to $n+1$, which have no effect on π_{n-1}.

(2) We have the sequence

$$X'_k \subseteq X''_{k+1} = X'_{k+1} \subseteq X''_{k+2} = X'_{k+2} \subseteq \ldots .$$

Each complex in this sequence collapses onto the one before it (sometimes trivially, where equality holds). Since the sequence terminates at X, (2) is proved.

(3) This follows easily from the previous proposition.

(4) In fact, the subcomplex in question is X'_1, and $\pi_0(X'_1) \to \pi_0(X)$ is a bijection by (1). Since X is connected, so is X'_1.

(5) Since $X''_1 \searrow X'_0$, each component of X''_1 is contractible, and so X''_1 is a forest. It is true by definition that X''_1 contains the whole 0-skeleton, so X''_1 is also maximal.

(6) Since $X''_1 \searrow X'_0$, and X'_0 is a singleton set, X''_1 is connected. By (5), we know that X''_1 is also a maximal forest, so it must in fact be a maximal tree. □

Choose a maximal tree T of X consisting of all of the collapsible edges in X, and additional critical edges, as needed.

Define a monoid presentation $\mathcal{MP}_{W,T}$ as follows: Generators are oriented edges in X, both positive and negative, so that there are two oriented edges for each geometric edge in X. If e denotes a particular oriented edge, let $\overline{e}$ denote the edge with the opposite orientation. If w denotes a sequence of oriented edges $e_1 \ldots e_m$, let $\overline{w}$ denote sequence of oriented edges $\overline{e_m} \ldots \overline{e_1}$.

A *boundary word* of a 2-cell c is simply one of the possible relations determined by an attaching map for c (cf. [19], page 139); if w_1 and w_2 are two boundary words for a cell c, then w_1 can be obtained from w_2 by the operations of inverting and taking cyclic shifts.

There are several types of relations.

(1) For a given oriented edge e in T, introduce the relations $(e,1)$ and $(\overline{e},1)$.

(2) For any oriented edge e, introduce relations $(e\overline{e},1)$ and $(\overline{e}e,1)$.

(3) For a collapsible 2-cell c, consider the (unique) geometric 1-cell e such that $e = W^{-1}(c)$. Suppose that a boundary word of c is ew. In this case, the word w contains no occurrence of e or $\overline{e}$, since the geometric edge corresponding to e is a regular face of c. Introduce the relations $(e,\overline{w})$ and $(\overline{e},w)$.

Proposition 2.4 *The rewrite system associated to the monoid presentation $\mathcal{MP}_{W,T}$ is complete.*

Proof If $w \to w_1$ and $w \to w_2$ correspond to disjoint applications of relations in $\mathcal{MP}_{W,T}$, that is, if $w = rs_1tu_1v$, $w_1 = rs_2tu_1v$, and $w_2 = rs_1tu_2v$, where r, s_1, s_2, t, u_1, u_2, and v are words in the free monoid on oriented edges in X, and (s_1,s_2), (u_1,u_2) are relations in $\mathcal{MP}_{W,T}$, then $w_1 \to w_3$ and $w_2 \to w_3$, where $w_3 = rs_2tu_2v$.

Thus we need only consider the cases in which the moves $w \to w_1$ and $w \to w_2$ are not disjoint. Checking definitions, we see that it is not possible for a move of the first type to overlap with a move of the third type, since the left side of a relation of type 1 is a word of length one involving only an edge in T, and the left side of a relation of type 3 is another word of length one, but involving a redundant edge, and redundant edges lie outside of T. Similarly easy arguments show that the moves $w \to w_1$ and $w \to w_2$ can only overlap if at least one of these moves is of the second type.

Suppose, without loss of generality, that $w = te\overline{e}v$ and $w_1 = tv$. If, say, $w_2 = t\overline{e}v$, and thus the move $w \to w_2$ involves the application of a relation of type 1, then it must be that the edge corresponding to e is in T, so that $w_2 \to w_1$. If $w \to w_2$ involves the application of a relation of type 3, say $w_2 = tewv$, then $w_2 \to t\overline{w}wv \dot{\to} tv = w_1$.

The case in which both moves $w \to w_1$ and $w \to w_2$ involve relations of the second type is left as an easy exercise. □

In view of Proposition 2.4 (and the remarks after Lemma 2.1), there is a unique reduced word in each equivalence class modulo the presentation $\mathcal{MP}_{W,T}$. If w is any word in the generators of $\mathcal{MP}_{W,T}$, let $r(w)$ denote the unique reduced word that is equivalent to w.

Theorem 2.5 *Let X be a finite connected CW complex with a discrete gradient vector field W. Then:*

$$\pi_1(X) \cong \langle \Sigma \mid \mathcal{R} \rangle,$$

where Σ is the set of positive critical 1-cells that aren't contained in T, and $\mathcal{R} = \{r(w) | w$ is the boundary word of a critical 2-cell$\}$.

Proof According to Proposition 2.3(1), $\pi_1(X'_2) \cong \pi_1(X)$. The usual edge-path presentation of the fundamental group (see [19], page 139) of X'_2 says that

$$\pi_1\left(X'_2\right) \cong \langle \widehat{\Sigma} \mid \widehat{\mathcal{R}} \rangle,$$

where $\widehat{\Sigma}$ is the set of positively oriented 1-cells, and the relations $\widehat{\mathcal{R}}$ are of three types: first, there are the boundary words of critical 2-cells; second, there are the boundary words of collapsible 2-cells; third, there are words of length one corresponding to edges in T.

Any collapsible 2-cell is necessarily $W(c)$ for some redundant 1-cell c. Choose c to have the largest rank of all 1-cells. There is a boundary word of $W(c)$ with the form cw, where w is a word involving only occurrences of 1-cells having smaller rank than c. In the presentation of $\pi_1(X'_2)$, replace every occurrence of c or $\overline{c}$ other than the occurrence in the boundary word of $W(c)$ with $\overline{w}$ or w, respectively. In the new presentation, neither c nor $\overline{c}$ occurs in any relation except the boundary word of $W(c)$, where c or $\overline{c}$ appears, but not both. In fact, the operation of replacing c (respectively, $\overline{c}$) with $\overline{w}$ (respectively, w) can change only the relations of the first type, by the assumption about the rank of c and because c is not collapsible. Thus we can remove the generator c from the new presentation for $\pi_1(X'_2)$ along with the boundary word for $W(c)$ to obtain another presentation of the same group. Notice that the effect on the boundary words of the critical 2-cells has been to perform a sequence of reductions of type 3.

One continues in the same way, inductively removing redundant 1-cells of the largest remaining rank, until all of the redundant 1-cells have been removed.

Next, remove the 1-cells occurring in T from the list of generators, along with the corresponding relations (which are words of length one), and remove all occurrences of the 1-cells of T from the remaining relations. This procedure results in a presentation of the same group. The final result after freely reducing is the presentation $\langle \Sigma \mid \mathcal{R} \rangle$, since every alteration to the boundary words of the critical 2-cells has been a move in the rewrite system corresponding to the monoid presentation $\mathcal{MP}_{W,T}$, the remaining words are reduced modulo

$\mathcal{MP}_{W,T}$, and the remaining generators are the positively-oriented 1-cells lying outside of T. □

In case there is just one critical 0-cell, the discrete gradient vector field W completely determines the maximal tree T, and we denote the presentation $\langle \Sigma \mid \mathcal{R} \rangle$ from the previous theorem $\mathcal{P}_W$, where the oriented CW complex is understood. The presentation $\mathcal{P}_W$ depends only on the choice of the boundary words for the critical 2-cells, since the string rewriting system associated to the monoid presentation $\mathcal{MP}_{W,T}$ is complete and by the previous theorem.

3 Discrete gradient vector fields and graph braid groups

Let Γ be a graph, and fix a natural number n. The *labelled configuration space* of Γ on n points is the space

$$\left(\prod^n \Gamma\right) - \Delta,$$

where Δ is the set of all points $(x_1, \ldots, x_n) \in \prod^n \Gamma$ such that $x_i = x_j$ for some $i \neq j$. The *unlabelled configuration space* of Γ on n points is the quotient of the labelled configuration space by the action of the symmetric group S_n, where the action permutes the factors. The *braid group* of Γ on n strands, denoted $B_n\Gamma$, is the fundamental group of the unlabelled configuration space of Γ on n strands. The *pure braid group*, denoted $PB_n\Gamma$, is the fundamental group of the labelled configuration space.

The set of vertices of Γ will be denoted by $V(\Gamma)$, and the degree of a vertex $v \in V(\Gamma)$ is denoted $d(v)$. If a vertex v is such that $d(v) \geq 3$, v is called *essential*.

Let Δ' denote the union of those open cells of $\prod^n \Gamma$ whose closures intersect the diagonal Δ. Let $\mathcal{D}^n\Gamma$ denote the space $\prod^n \Gamma - \Delta'$. Note that $\mathcal{D}^n\Gamma$ inherits a CW complex structure from the Cartesian product, and that a cell in $\mathcal{D}^n\Gamma$ has the form $c_1 \times \cdots \times c_n$ where each c_i is either a vertex or the interior of an edge whose closure is disjoint from the closure of c_j for $i \neq j$. We also let $U\mathcal{D}^n\Gamma$ denote the quotient of $\mathcal{D}^n\Gamma$ by the action of the symmetric group S_n by permuting the coordinates. Thus, an open cell in $U\mathcal{D}^n\Gamma$ may be written $\{c_1, \ldots, c_n\}$ where each c_i is either a vertex or the interior of an edge whose closure is disjoint from the closure of c_j for $i \neq j$. The set notation is used to indicate that order does not matter.

Under most circumstances, the labelled (respectively, unlabelled) configuration space of Γ is homotopy equivalent to $\mathcal{D}^n\Gamma$ (respectively, $U\mathcal{D}^n\Gamma$). Specifically:

Theorem 3.1 [2] *For any $n > 1$ and any graph Γ with at least n vertices, the labelled (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 homotopically nontrivial path from a vertex to itself passes through at least $n+1$ edges.*

A graph Γ satisfying the conditions of this theorem for a given n is called *sufficiently subdivided* for this n. It is clear that every graph is homeomorphic to a sufficiently subdivided graph, no matter what n may be. Throughout the rest of the paper, we work exclusively with the space $U\mathcal{D}^n\Gamma$ where Γ is sufficiently subdivided for n. Also from now on, "edge" and "cell" will refer to closed objects.

We define a discrete gradient vector field W using a maximal tree and specific order on the vertices of Γ. Choose a maximal tree T in Γ. Edges outside of T are called *deleted edges*. Pick a vertex $*$ of valence 1 in T to be the root of T. Define an operation $\wedge$ on vertices of T to take two vertices v_1 and v_2 and yield the vertex of T which is the endpoint of the geodesic segment $[*, v_1] \cap [*, v_2]$ other than $*$.

For a given vertex $v \in \Gamma$, call the edges adjacent to v the *directions from* v. For each v, fix a total ordering of the directions from v by labelling each direction with a number between 0 and $d(v)-1$, assigning the number 0 to the direction leading from v back to $*$. (The single direction from $*$ is given the number 1.) Also, define a function $g : V(\Gamma) \times V(\Gamma) \to \mathbb{Z}$ such that $g(v_1, v_2)$ is the label of the direction from v_1 that lies on the unique geodesic connecting v_1 to v_2 within T, or $g(v_1, v_2) = 0$ if $v_1 = v_2$.

Now order the vertices of T as follows. Let v_1 and v_2 be two vertices, and define the vertex $v_3 := v_1 \wedge v_2$. Then $v_1 \leq v_2$ if and only if $v_3 = v_1$, or $v_3 \neq v_1$ and $g(v_3, v_1) < g(v_3, v_2)$.

For a given edge e, let $\iota(e)$ and $\tau(e)$ denote the endpoints of e, where $\iota(e) \geq \tau(e)$ in the ordering on vertices. For a vertex $v \neq *$ we let $e(v)$ denote the (unique) edge of T such that $\iota(e(v)) = v$ – i.e. the edge of T incident with v and closest to $*$. The following lemma is easy to prove:

Lemma 3.2 (Order Lemma) *The ordering $\leq$ on vertices of Γ is a linear order, with the following additional properties:*

(i) *if $v_2 \in [*, v_1]$ then $v_2 \leq v_1$*

(ii) *if $v \in V(\Gamma)$ and e is an edge of T such that $e(v) \cap e = \tau(e)$ and $v < \iota(e)$, then $\tau(e)$ is an essential vertex and $0 < g(\tau(e), v) < g(\tau(e), \iota(e))$ (and thus $\tau(e) < v < \iota(e)$).* □

Example 3.3 Consider the labelled tree Γ in Figure 1, which is sufficiently subdivided for $n = 4$. This tree is especially interesting since it is the smallest tree for which $B_n\Gamma$ $(n \geq 4)$ appears not to be a right-angled Artin group. See Example 5.4 for a presentation of $B_4\Gamma$, and Example 5.5 for a discussion of the right-angled Artin property.

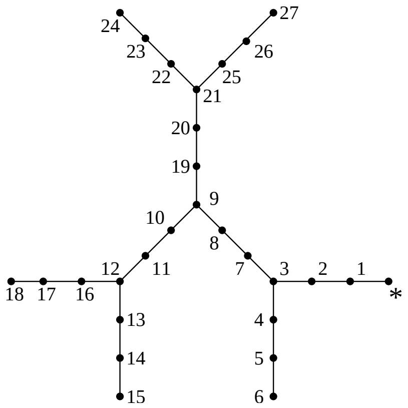

Figure 1: A sufficiently subdivided tree

Let us write v_i for the vertex labelled i. The numbering induces an obvious linear order on the vertices: $v_i < v_j$ if $i < j$. This order is a special instance of the one mentioned above, for a particular choice of g which we now describe. It is enough to describe how to number the directions from any given vertex v. The direction from v to $*$ is numbered 0, as is required, and the other directions are numbered $1, 2, \ldots, d(v) - 1$ consecutively in the clockwise order. It is meaningful to speak of a clockwise ordering because we have specified an embedding in the plane by the picture. The numbering of directions depends only on the location of $*$ and the choice of the embedding.

Thus, for example, $g(v_3, v_4) = 1$, $g(v_3, v_7) = 2$, $g(v_9, v_{13}) = 1$, and so forth.

A heuristic way to describe the numbering (or order) of the vertices is as follows. Begin with an embedded tree, having a specified basepoint $*$ of degree one. Now walk along the tree, following the leftmost branch at any given intersection, and consecutively number the vertices in the order in which they are first encountered. (When you reach a vertex of degree one, turn around.) Note that this numbering depends only on the choice of $*$ and the embedding of the tree.

Any ordering of the kind mentioned in the lemma may be realized by an embedding and a choice of $*$, although we will not need to make use of this fact.

3.1 The function W is a discrete vector field

Let $c = \{c_1, \ldots, c_{n-1}, v\}$ be a cell of $U\mathcal{D}^n\Gamma$ containing the vertex v. If $e(v) \cap c_i = \emptyset$ for $i = 1, \ldots, n-1$, then define the cell $\{c_1, \ldots, c_{n-1}, e(v)\} \subset U\mathcal{D}^n\Gamma$ to be the *elementary reduction* of c from v, and we say v is *unblocked* in c. Otherwise, there exists some $c_i \in c$ with $e(v) \cap c_i \neq \emptyset$, and we say v is *blocked by c_i in c*. If v is the smallest unblocked vertex of c in the sense of the order on vertices, then the reduction from v is *principal.*

Define a function W on $U\mathcal{D}^n\Gamma$ inductively. If c is a 0-cell, let $W_0(c)$ be the principal reduction if it exists. For $i > 0$, let $W_i(c)$ of an i-cell c be its principal reduction if it exists *and* $c \notin \operatorname{im} W_{i-1}$, and undefined otherwise.

Let c be a cell, and let $e \in c$ be an edge of Γ. The edge e is said to be *order-respecting in c* if $e \subseteq T$ and, for every vertex $v \in c$, $e(v) \cap e = \tau(e)$ implies that $v > \iota(e)$ in the order on vertices.

For example, consider the tree in Example 3.3 (Figure 1). For any vertex v_i, let $e_i := e(v_i)$. For instance, e_{19} connects v_{19} to v_9. In the 1-cell $\{v_{10}, e_{19}, v_{12}, v_{16}\}$, e_{19} is not order-respecting, since $e_{10} \cap e_{19} = v_9 = \tau(e_{19})$. On the other hand, e_{19} is order-respecting in $\{e_{19}, v_{20}, v_{21}, v_{22}\}$. The edge e_{10} is order-respecting in any cell of $C^4\Gamma$.

An edge e that is order-respecting in c is the *minimal order-respecting edge in c* if $\iota(e)$ is minimal in the order on vertices among the initial vertices of the order-respecting edges in c.

Lemma 3.4 (Order-respecting edges lemma)

(i) *If $\{c_1, \ldots, c_{n-1}, e\}$ is obtained from $\{c_1, \ldots, c_{n-1}, \iota(e)\}$ by principal reduction, then e is an order-respecting edge.*

(ii) *Let* $c = \{c_1, \ldots, c_{n-1}, e\}$ *where* e *is an edge contained in* T*. Then an edge* e' *in* $\{c_1, \ldots, c_{n-1}, \iota(e)\}$ *is order-respecting if and only if it is order-respecting in* c*.*

Proof

(i) Suppose that e is not an order-respecting edge. Clearly $e \subseteq T$, so it must be that there is $v \in \{c_1, \ldots, c_{n-1}, e\}$ such that $e(v) \cap e = \tau(e)$ and $v < \iota(e)$. The elementary reduction from v is thus well-defined in $\{c_1, \ldots, c_{n-1}, \iota(e)\}$, so the principal elementary reduction of $\{c_1, \ldots, c_{n-1}, \iota(e)\}$ cannot be the elementary reduction from $\iota(e)$, since $v < \iota(e)$. We have a contradiction.

(ii) We can assume that e' is contained in T, for otherwise e' fails to be order-respecting in every cell, and there is nothing to prove.

($\Rightarrow$) If e' is not order-respecting in $\{c_1, \ldots, c_{n-1}, e\}$, then there is $v \in \{c_1, \ldots, c_{n-1}, e\}$ such that $e(v) \cap e' = \tau(e')$ and $v < \iota(e')$. Then $v \in \{c_1, \ldots, c_{n-1}\} \subseteq \{c_1, \ldots, c_{n-1}, \iota(e)\}$ and thus e' is not order-respecting in $\{c_1, \ldots, c_{n-1}, \iota(e)\}$.

($\Leftarrow$) Suppose without loss of generality that $e' = c_{n-1}$. If e' is not order-respecting in $c' = \{c_1, \ldots, c_{n-2}, e', \iota(e)\}$, then there is $v \in c'$ such that $e(v) \cap e' = \tau(e')$ and $v < \iota(e')$. In this case $v \neq \iota(e)$ since, otherwise, $e(v) = e$ so $e \cap e' \neq \emptyset$, a contradiction. Thus $v \in \{c_1, \ldots, c_{n-2}\} \subseteq \{c_1, \ldots, c_{n-2}, e', e\}$, and so e' is not order-respecting in $\{c_1, \ldots, c_{n-1}, e\}$. □

Lemma 3.5 (Classification lemma)

(i) *If a cell* c *contains no order-respecting edge, then it is critical if every vertex of* c *is blocked, and redundant otherwise.*

(ii) *Suppose* c *contains an order-respecting edge, and let* e *denote the minimal order-respecting edge in* c*. If there is an unblocked vertex* $v \in c$ *such that* $v < \iota(e)$*, then* c *is redundant. If there is no such vertex, then* c *is collapsible.*

Proof

(i) The previous lemma implies that c cannot be in the image of W. If every vertex in c is blocked, then the principal elementary reduction of c is undefined, and thus c is critical. Otherwise, the principal elementary reduction of c is defined, and c is redundant.

(ii) Suppose first that there is an unblocked vertex $v \in c$ such that $v < \iota(e)$. We claim that $c = \{c_1, \ldots, c_n\}$ is not in the image of W. If it is, then there is some e' (without loss of generality, $e' = c_n$) such that $\{c_1, \ldots, c_{n-1}, e'\}$ is the principal elementary reduction of $\{c_1, \ldots, c_{n-1},\ \iota(e')\}$. It follows that e' is an order-respecting edge in $\{c_1, \ldots, c_{n-1}, e'\}$, so that $\iota(e') \geq \iota(e) > v$. Suppose, again without loss of generality, that $v = c_{n-1}$. Now, since elementary reduction from v is defined for $\{c_1, \ldots, c_{n-2}, v, \iota(e')\}$, it follows that $\{c_1, \ldots, c_{n-1}, e'\}$ is not the principal elementary reduction of $\{c_1, \ldots, c_{n-1}, \iota(e')\}$. This is a contradiction. Since c is not in the image of W and it has unblocked vertices, it must be redundant.

Now suppose that there is no unblocked vertex $v \in c$ satisfying $v < \iota(e)$. Suppose, without loss of generality, that $c = \{c_1, \ldots, c_{n-1}, e\}$. We claim that $\{c_1, \ldots, c_{n-1}, e\}$ is the principal reduction from $\{c_1, \ldots, c_{n-1},\ \iota(e)\})$. Let v be the smallest unblocked vertex of $\{c_1, \ldots, c_{n-1}, \iota(e)\}$. Clearly $v \leq \iota(e)$, since $\iota(e)$ is unblocked. If $v < \iota(e)$, then v is unblocked in $\{c_1, \ldots, c_{n-1}, \iota(e)\}$ but blocked in $\{c_1, \ldots, c_{n-1}, e\}$. This can only be because $e(v) \cap e = \tau(e)$. Since e is order-respecting, we have $v > \iota(e)$, which is a contradiction. Thus the principal elementary reduction of $\{c_1, \ldots, c_{n-1}, \iota(e)\}$ is $\{c_1, \ldots, c_{n-1}, e\}$.

It remains to be shown that $\{c_1, \ldots, c_{n-1}, \iota(e)\}$ is not collapsible. If it is, then there is some edge $e' \in \{c_1, \ldots, c_{n-1}, \iota(e)\}$ (without loss of generality, $e' = c_{n-1}$) such that $\{c_1, \ldots, c_{n-2}, e', \iota(e)\}$ is the principal elementary reduction of $\{c_1, \ldots, c_{n-2}, \iota(e'), \iota(e)\}$. Since $\iota(e')$ and $\iota(e)$ are both unblocked in $\{c_1, \ldots, c_{n-2}, \iota(e'), \iota(e)\}$, $\iota(e') < \iota(e)$. We claim that e' is an order-respecting edge of $\{c_1, \ldots, c_{n-2}, e', e\}$. Certainly $e' \subseteq T$. If e' is not an order-respecting edge of $\{c_1, \ldots, c_{n-2}, e', e\}$, then there is some vertex $v \in \{c_1, \ldots, c_{n-2}, e', e\}$ (without loss of generality, $v = c_{n-2}$) such that $e(v) \cap e' = \tau(e')$ and $v < \iota(e') < \iota(e)$. Thus v is unblocked in $\{c_1, \ldots, c_{n-3}, v, \iota(e'), \iota(e)\}$ and $v < \iota(e')$, so that the elementary reduction of $\{c_1, \ldots, c_{n-3}, v, \iota(e'), \iota(e)\}$ from $\iota(e')$ is not principal, a contradiction. This proves that e' is an order-respecting edge of $\{c_1, \ldots, c_{n-2}, e', e\}$.

We now reach the contradiction that e is not the minimal order-respecting edge of c. This completes the proof. □

Theorem 3.6 (Classification Theorem)

(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 a minimal order-respecting edge e) 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 (and thus a minimal order-respecting edge e) and, for any $v < \iota(e)$, v is blocked.*

Proof This follows logically from the previous lemma. □

Theorem 3.7 *The function W is one-to-one.*

Proof Suppose that $c = \{c_1, \ldots, c_n\}$ is collapsible. Thus there is a minimal order-respecting edge e. Since c is in the image of W, there must exist some edge e' (without loss of generality, assume $c_n = e'$) such that

$$W(\{c_1, \ldots, c_{n-1}, \iota(e')\}) = \{c_1, \ldots, c_{n-1}, e'\}.$$

A previous lemma implies that e' is order-respecting in c. We claim that $e' = e$.

If not, then $\iota(e') > \iota(e)$, $\{c_1, \ldots c_{n-1}, \iota(e')\}$ is redundant, and e is order-respecting in $\{c_1, \ldots, c_{n-1}, \iota(e')\}$. By the previous theorem, there is an unblocked vertex $v \in \{c_1, \ldots c_{n-1}, \iota(e')\}$ such that $v < \iota(e) < \iota(e')$. Now since $v \neq \iota(e')$, $v \in c$, in which it must be blocked, since c is collapsible. It follows that $e(v) \cap e' = \tau(e')$ and $v < \iota(e')$. Since $e' \subseteq T$, we have that e' is not order-respecting in c, a contradiction.

It now follows that W is one-to-one, since we can solve for the preimage of any collapsible cell. □

3.2 Proof that there are no non-stationary closed W-paths

For each vertex v of Γ, define a function f_v from the cells of $UD^n\Gamma$ to $\mathbb{Z}$, setting $f_v(c)$ equal to the number of $c_i \in c$ such that c_i is a subset of the geodesic in T connecting $*$ to v.

Each function f_v has the following properties:

(1) For any redundant cell c, $f_v(c) = f_v(W(c))$.

(2) If a cell c' is obtained from a cell c by replacing $e \subseteq T$ with $\iota(e)$, then $f_v(c') = f_v(c)$.

(3) If a cell c' is obtained from a cell c by replacing $e \subseteq T$ with $\tau(e)$, then $f_v(c') = f_v(c)$ unless $v \wedge \iota(e) = \tau(e)$, in which case $f_v(c') = f_v(c) + 1$.

(4) If a cell c' is obtained from a cell c by replacing $e \not\subseteq T$ with $\tau(e)$, then $f_v(c') = f_v(c)$ unless $v \wedge \tau(e) = \tau(e)$, in which case $f_v(c') = f_v(c) + 1$.

(5) If a cell c' is obtained from a cell c by replacing $e \not\subseteq T$ with $\iota(e)$, then $f_v(c') = f_v(c)$ unless $v \wedge \iota(e) = \iota(e)$, in which case $f_v(c') = f_v(c) + 1$.

Theorem 3.8 *W has no non-stationary closed paths.*

Proof Suppose that $\sigma_0, \ldots, \sigma_r$ is a minimal non-stationary closed path, so that no repetitions occur among the subsequence $\sigma_0, \ldots, \sigma_{r-1}$, $\sigma_0 = \sigma_r$, and $r > 1$. Since, for any vertex v in Γ, $f_v(\sigma_0) \leq f_v(\sigma_1) \leq \ldots \leq f_v(\sigma_r) = f_v(\sigma_0)$, equality must hold throughout. By considering different choices for v, it is clear that σ_{i+1} may not be obtained from $W(\sigma_i)$ using rules (3), (4), or (5). Thus, σ_{i+1} is obtained from $W(\sigma_i)$ by replacing some edge $e' \in W(\sigma_i)$ with $\iota(e')$ (and never with $\tau(e')$), where e' is necessarily contained in T. Note also that each of the cells $\sigma_0, \sigma_1, \ldots, \sigma_r = \sigma_0$ must be redundant.

We claim that if σ_{i+1} is obtained from $W(\sigma_i)$ by replacing some $e \in W(\sigma_i)$ with $\iota(e)$, then $\iota(e_{i+1}) < \iota(e_i)$, where e_i and e_{i+1} are the minimal order-respecting edges in σ_i and σ_{i+1}, respectively. (If e_i or e_{i+1} doesn't exist, then $\iota(e_i)$ or $\iota(e_{i+1})$, respectively, is ∞.)

Consider first the case in which σ_i has no order-respecting edges. Since $W(\sigma_i)$ is collapsible, it has an order-respecting edge e, and $W(\sigma_i)$ must be obtained from σ_i by replacing $\iota(e)$ with e (see, for instance, the proof that W is injective). Since $\sigma_{i+1} \neq \sigma_i$ and σ_{i+1} must be obtained from $W(\sigma_i)$ by replacing an edge e' from $W(\sigma_i)$ with $\iota(e')$, then $e' \neq e$. Thus e is the only order-respecting edge in σ_{i+1}, and so $e_{i+1} = e$ and $\iota(e_{i+1}) = \iota(e) < \infty = \iota(e_i)$. This establishes the claim if σ_i has no order-respecting edges.

Now suppose that σ_i has a minimal order-respecting edge e_i. Since σ_i is redundant, the minimal unblocked vertex v of σ_i satisfies $v < \iota(e_i)$. The cell $W(\sigma_i)$ is obtained from σ_i by replacing v with $e(v)$. Since $\sigma_{i+1} \neq \sigma_i$, σ_{i+1} must be obtained from $W(\sigma_i)$ by replacing some edge $e' \neq e(v)$ with $\iota(e')$. This implies that $e(v)$ is an order-respecting edge in σ_{i+1} since $e(v)$ is clearly order-respecting in $W(\sigma_i)$, and thus $\iota(e_{i+1}) \leq v < \iota(e_i)$. This proves the claim.

We now reach a contradiction, because $\iota(e_0) > \iota(e_1) > \ldots > \iota(e_r)$, but $\sigma_0 = \sigma_r$. □

3.3 Visualizing critical cells of W

Using Theorem 3.6, it is not difficult to visualize the critical cells of $U\mathcal{D}^n\Gamma$, for any n and any Γ. If c is a critical cell in $U\mathcal{D}^n\Gamma$, then every vertex v in c is blocked, and every edge e in c has the property that either: (i) e is a deleted edge, or (ii) $\tau(e)$ is an essential vertex, and there is some vertex v of c that is adjacent to $\tau(e)$ satisfying $\tau(e) < v < \iota(e)$.

Consider Γ the tree given in Example 3.3. A critical 1-cell in $\mathcal{C}^4\Gamma$ is depicted in Figure 2:

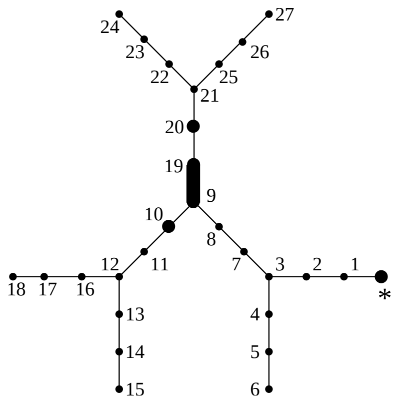

Figure 2: An example of a critical 1-cell

This is the 1-cell $\{e_{19}, v_{10}, v_{20}, *\}$, by the numbering of Figure 1. Notice that, since the order on vertices depends only on the embedding and the location of $*$, it would be immediately clear that $\tau(e_{19}) < v_{10} < \iota(e_{19})$ even in the absence of an explicit numbering. It is convenient to introduce a notation for certain cells (critical cells among them) which doesn't refer to a numbering of the vertices. The notation we introduce will also be "independent of subdivision" in a sense which will be specified.

Fix an alphabet $A, B, C, D, \ldots$. Assign to each essential vertex of Γ a letter from the alphabet. For this example, we will do this so that the smallest essential vertex is assigned the letter A, the next smallest is assigned the letter B, and so on. Thus, $A = v_3$, $B = v_9$, $C = v_{12}$, and $D = v_{21}$. If $\vec{a} = (a_1, \ldots, a_{d(A)-1})$ is a vector with $d(A) - 1$ entries, all in $\mathbb{N} \cup \{0\}$, $1 \leq m \leq d(A) - 1$, and $a_m \neq 0$, we define the notation $A_m[\vec{a}]$ to represent the following subconfiguration. In the mth direction from A, there is a single edge e with $\tau(e) = A$, and $a_m - 1$ vertices being blocked, stacked up behind e. In any other direction $i > 0$, $i \neq m$ from A, there are a_i blocked vertices stacked up behind e at A. For example, $A_2[(1,2)]$ refers to the collection $\{v_4, e_7, v_8\}$, and $A_1[(3,0)]$ refers to $\{v_6, v_5, e_4\}$. We also need a notation for collections of blocked vertices clustered around $*$. Let $k*$ denote a collection of k vertices blocked at $*$. Thus $1*$ refers to $\{*\}$, $2*$ refers to $\{v_1, *\}$, and so forth. We combine these new expressions using additive notation. Thus, the cell $\{v_{10}, e_{19}, v_{20}, *\}$ would be expressed as $B_2[(1,2)] + 1*$. The critical 2-cell $\{v_{13}, e_{16}, v_{10}, e_{19}\}$ would be expressed as $B_2[(1,1)] + C_2[(1,1)]$. If a cell c contains deleted edges, these can be simply listed, while still using the additive notation. For example, one would write $e + k*$ to represent a cell consisting of the deleted edge e together with a collection of k vertices clustered at $*$.

We must mention one last convention. The expression $B_1[(4,0)]$ could refer either to $\{v_{13}, v_{12}, v_{11}, e_{10}\}$ or $\{v_{16}, v_{12}, v_{11}, e_{10}\}$. It is clear that if we subdivide Γ further, then this ambiguity disappears. In arguments using our new notation, we extend the notion of "sufficiently subdivided for n" in such a way that all of our new notations are unambiguous for Γ when they refer to collections of n or fewer cells. It is clearly possible to do this.

Now our notation is "subdivision invariant": if Γ is sufficiently subdivided for n, and the expression $A_m[\vec{a}] + B_n[\vec{b}] + \ldots + k*$ satisfies $a_1 + \ldots + a_{d(A)-1} + b_1 + \ldots + b_{d(B)-1} + \ldots + k \leq n$, then $A_m[\vec{a}] + B_n[\vec{b}] + \ldots + k*$ specifies a unique cell of $U\mathcal{D}^n\Gamma$, and does so no matter how many times we subdivide.

With these conventions, every $U\mathcal{D}^n\Gamma$ has a unique critical 0-cell, namely $n*$. The discrete gradient vector field W thus determines a presentation of $B_n\Gamma$ which is unique up to the choices of the boundary words of critical 2-cells in $U\mathcal{D}^n\Gamma$.

We record for future reference a description of critical cells, in terms of the new notation:

Proposition 3.9 *Let Γ be a sufficiently subdivided graph, with a maximal subtree T and basepoint $*$. Suppose that T has been embedded in the plane, so that there is a natural order on the vertices. Assume also that the endpoint of each deleted edge has degree 1 in T. With notation as above, a cell described by a formal sum*

$$A_{l_A}[\vec{a}] + B_{l_B}[\vec{b}] + \ldots + e_1 + e_2 + \ldots + k*,$$

where each e_i is a deleted edge, is critical provided that, for $X = A, B, \ldots$, some component x_j of $\vec{x} = (x_1, x_2, \ldots, x_{d(X)-1})$ is non-zero, for $j < l_X$. Conversely, every critical cell can be described by such a sum.

If, for some essential vertex X, $x_j = 0$ for all $j < l_X$, then the above cell is collapsible. □

We refer to each term of the sum in Proposition 3.9 as a *subconfiguration.* If the subconfiguratin itself represents a critical cell (with fewer strands), we call it *critical.*

4 Corollaries

Theorem 4.1 *Let Γ be a graph with the maximal tree T. If $T \neq \Gamma$, then assume that the endpoints of every deleted edge have degree 1 in the tree T (and furthermore that T is sufficiently subdivided). Fix the discrete gradient vector field W as in the previous section. Let D be the number of deleted edges. Then $\mathcal{P}_W$ has*

$$D + \sum_{\substack{v \in V(T) \\ essential}} \sum_{i=2}^{d(v)-1} \left[\binom{n+d(v)-2}{n-1} - \binom{n+d(v)-i-1}{n-1} \right]$$

generators.

Proof If c is a critical 1-cell, then c contains exactly one edge e, which must either be a deleted edge or there is some v in c such that $e(v) \cap e = \tau(e)$. In the latter case, $\tau(e)$ is an essential vertex, and $0 < g(\tau(e), v) < g(\tau(e), \iota(e))$. It follows that $2 \leq g(\tau(e), \iota(e)) \leq d(\tau(e)) - 1$.

Now let us count the critical 1-cells c. If the unique edge e in c is a deleted edge, then c is uniquely determined by e, so there are exactly D critical 1-cells of this description. If the edge e is not a deleted edge, then $\tau(e)$ is essential and $2 \leq g(\tau(e), \iota(e)) \leq d(\tau(e)) - 1$. Note that since every vertex of c is necessarily blocked, the critical cell c is determined by the numbers of vertices of c that are in each of the $d(\tau(e))$ connected components of $T - e$. There are $\binom{n+d(\tau(e))-2}{n-1}$ ways to assign $n-1$ vertices to the $d(\tau(e))$ connected components of $T - e$. (This is the number of ways to assign $n-1$ indistinguishable balls to $d(\tau(e))$ distinguishable boxes.) Not every such assignment results in a critical 1-cell, however. The condition that $0 < g(\tau(e), v) < g(\tau(e), \iota(e))$, for some $v \in c$, won't be satisfied if, for each $v \in c$, either $g(\tau(e), \iota(e)) \leq g(\tau(e), v) \leq d(\tau(e)) - 1$ or $g(\tau(e), v) = 0$. There are $\binom{d(\tau(e))+n-g(\tau(e),\iota(e))-1}{n-1}$ such "illegal" assignments. Subtracting these from the total, we get

$$\binom{n+d(\tau(e))-2}{n-1} - \binom{n+d(\tau(e))-g(\tau(e),\iota(e))-1}{n-1}$$

different critical 1-cells for a fixed edge e. Letting the edge e of c vary over all possibilities, we obtain the sum in the statement of the theorem. □

Corollary 4.2 [14] *If Γ is a* radial *tree – i.e. has exactly one essential vertex, v – then $B_n\Gamma$ is free of rank*

$$\sum_{i=2}^{d(v)-1} \left[\binom{n+d(v)-2}{n-1} - \binom{n+d(v)-i-1}{n-1} \right].$$

Proof There are no critical cells of dimension greater than 1 by the classification of critical cells, since each blocking edge must be at its own essential vertex. Thus the presentation $\mathcal{P}_W$ has no relations. By Theorem 4.1, the presentation $\mathcal{P}_W$ has the given number of generators. $\square$

Theorem 4.3 *Let Γ be a tree and c a critical cell of $UD^n\Gamma$. Let*

$$k := \min\left\{\left\lfloor \frac{n}{2} \right\rfloor, \#\{v \in \Gamma^0 \mid v \text{ is essential}\}\right\}.$$

Then $\dim c \leq k$. In particular, $UD^n\Gamma$ strong deformation retracts on $(UD^n\Gamma)'_k$.

Proof For every edge e in c, there is some vertex v in c such that $e(v) \cap e = \tau(e)$. If e_1 and e_2 are two edges in c, and the vertices $v_1, v_2 \in c$ satisfy $e(v_i) \cap e_i = \tau(e_i)$, for $i = 1, 2$, then certainly $v_1 \neq v_2$, since $e_1 \cap e_2 = \emptyset$. It follows that there are at least as many vertices as edges in c. Since the dimension of c is equal to the number of edges in c, and the total number of cells in c is n, we have that the dimension of c is less than or equal to $n/2$.

Since $\tau(e)$ must be an essential vertex of Γ for each e in c, and the edges contained in c must be disjoint, we have that the dimension of c is bounded above by the number of essential vertices of Γ.

The final statement now follows from Proposition 2.3(2). $\square$

The following result was proven independently for the tree case by Carl Mautner, an REU student working under Aaron Abrams [1].

Theorem 4.4 *Let Γ be a sufficiently subdivided graph, and let $\chi(\Gamma)$ denote the Euler characteristic of Γ. Then $UD^n\Gamma$ strong deformation retracts onto a CW-complex of dimension at most k, where*

$$k := \min\left\{\left\lfloor \frac{n+1-\chi(\Gamma)}{2} \right\rfloor, \#\{v \in \Gamma^0 \mid v \text{ is essential}\}\right\}.$$

Proof We construct a maximal subtree T of Γ whose deleted edges all neighbor essential vertices in Γ. We note that the existence of a *connected* maximal subtree with this property is not clear a priori.

Let T' be any maximal subtree for Γ. Let T be a maximal subtree of Γ constructed as follows. For every deleted edge $e \in \Gamma - T'$, there are two essential vertices of Γ nearest e. For each deleted edge e, let v_e a choice of be such a vertex. Construct T by removing from Γ, for every deleted edge e of T', the unique edge adjacent to v_e on the simple path from v_e to e which does not cross any essential vertices. Then T is connected and is a maximal tree since T' is.

Any embedding of the tree T induces a discrete gradient vector field W with the property that *every* edge in a critical cell contains an essential vertex. Thus, the number of essential vertices bounds the dimension of the cells of $U\mathcal{D}^n\Gamma$ that are critical with respect to W.

Let D be the number of deleted edges with respect to T. Thus $D = 1 - \chi(\Gamma)$. Since a critical subconfiguration involves either one strand on a deleted edge or at leasttwo strands about an essential vertex, the dimension of any critical cell is also bounded by $D + \lfloor \frac{n-D}{2} \rfloor = \lfloor \frac{n+D}{2} \rfloor$.

The theorem now follows from Proposition 2.3(2). □

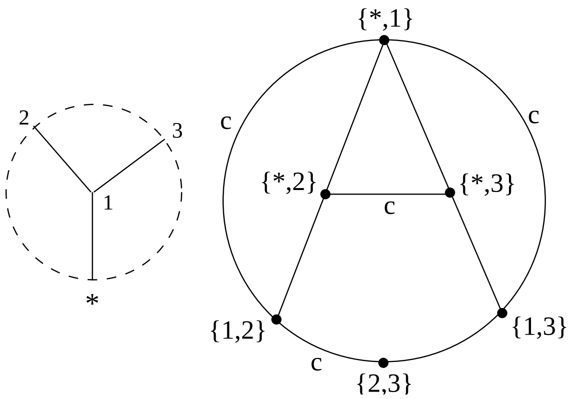

Figure 3: On the left, we have K_4 with the choice of maximal tree depicted. On the right, we have the complex X_1' onto which $\mathcal{UD}^2K_4$ deformation retracts. The critical 1-cells are indicated with a lower case "c"; the critical 0-cell is $\{*, 1\}$.

Note that this bound is not sharp, as two deleted edges for T adjacent to the same essential vertex may not both be involved in a critical configuration. The real thing to count (which varies by choice of T) is the number of essential vertices in Γ which touch deleted edges, not deleted edges themselves.

Example 4.5 Consider $\mathcal{UD}^2K_4$, where K_4 is the complete graph on four vertices. We choose a radial tree as our maximal tree. It is not difficult to check that there are no critical 2-cells in $\mathcal{UD}^2K_4$ with respect to this choice of maximal tree, so the subcomplex $X_1' \subseteq \mathcal{UD}^2K_4$ is a strong deformation retract of $\mathcal{UD}^2K_4$ (see Figure 3).

5 Presentations of tree braid groups

Let v and v' be vertices of the graph Γ. Let $(v,v') = \{v'' \in V(\Gamma) \mid v < v'' < v'\}$; in other words, (v,v') denotes the set of vertices "between" v and v' in the ordering on vertices. The following lemma is extremely useful in computing presentations of graph braid groups.

Lemma 5.1 (Redundant 1-cells lemma) *Let c be a redundant 1-cell in $UD^n\Gamma$, where Γ is an arbitrary finite graph. Suppose $c = \{c_1, \ldots, c_{n-2}, v, e\}$, where v is the smallest unblocked vertex of c and e is the unique edge of c. Let $c' := \{c_1, \ldots, c_{n-2}, \tau(e(v)), e\}$. If $(\tau(e(v)), v) \cap \{c_1, c_2, \ldots, c_{n-2}, \iota(e), \tau(e)\} = \emptyset$, then $c \dot{\to} c'$.*

Here $\dot{\to}$ refers to a sequence of moves over the complete rewrite system $\mathcal{MP}_{W,T}$ (see Proposition 2.4 and the discussion preceding). In particular, c and c' are words (of length 1) in the alphabet of oriented 1-cells.

Proof Let $c_\iota := \{c_1, \ldots, c_{n-2}, e(v), \iota(e)\}$ and $c_\tau := \{c_1, \ldots, c_{n-2}, e(v), \tau(e)\}$. If we apply W to c, then we get the following 2-cell, where the arrows indicate orientation:

$$\begin{tikzcd} \bullet \arrow[r, "c"] \arrow[d, "c_\iota"'] & \bullet \arrow[d, "c_\tau"] \\ \bullet \arrow[r, "c'"'] & \bullet \end{tikzcd}$$

It follows that $c \to c_\iota c' \overline{c_\tau}$; we need to show that $c_\iota \dot{\to} 1$ and $c_\tau \dot{\to} 1$.

Consider the following condition on 1-cells $\hat{c}$ of $UD^n\Gamma$:

(1) If v is a vertex and $v \in \hat{c}$, then $v \notin (\tau(\hat{e}), \iota(\hat{e}))$, where $\hat{e} \subseteq T$ is the unique edge of $\hat{c}$.

We claim: first, that if a 1-cell $\hat{c}$ satisfies (1), then so does any 1-cell in a W-path starting with $\hat{c}$; second, that a 1-cell $\hat{c}$ satisfying (1) cannot be critical. Since c_ι and c_τ both satisfy (1) by hypothesis, this will prove the lemma by induction on the rank of c_ι and c_τ.

We begin with the latter claim. Suppose $\hat{c}$ satisfies (1). If $\hat{c}$ is critical, then there is $v \in \hat{c}$ such that $e(v) \cap \hat{e} = \tau(\hat{e})$, and $0 < g(\tau(\hat{e}), v) < g(\tau(\hat{e}), \iota(\hat{e}))$. But then $\tau(\hat{e}) < v < \iota(\hat{e})$, a contradiction.

Now for the first claim. If $\hat{c}$ is collapsible, then any W-path starting with $\hat{c}$ consists of $\hat{c}$ alone. We may therefore assume that $\hat{c}$ is redundant. Since $\hat{c}$ is redundant, there is $v \in \hat{c}$ such that $v < \iota(\hat{e})$ and v is unblocked. We may choose v to be the minimal vertex with this property. By assumption, $v < \tau(\hat{e})$ as well. The two edges in $W(\hat{c})$, namely $e(v)$ and $\hat{e}$, therefore have the property that $\tau(e(v)) < \iota(e(v)) < \tau(\hat{e}) < \iota(\hat{e})$. Moreover, there is no vertex in $W(\hat{c}) \cap (\tau(e(v)), \iota(e(v)))$, for a minimal such vertex v' would be an unblocked vertex in $\hat{c}$ satisfying $v' < v$, which is impossible. The immediate successor $\hat{c}_1$ of $\hat{c}$ in a W-path is obtained by replacing an edge of $W(\hat{c})$ with either its initial or terminal vertex. It thus follows that $\hat{c}_1$ satisfies (1). $\square$

In order to prove our theorem on the presentations of tree braid groups, we need to introduce new notation. Let $\dot{A}[\vec{a}]$ denote the collection of vertices consisting of A itself together with a_i vertices arranged consecutively in the ith direction from A, so that every vertex in the collection is blocked except for A. Thus, to use the tree from Figure 1, the expression $\dot{B}[(2,1)]$ would refer to the collection $\{v_9, v_{10}, v_{11}, v_{19}\}$. Let $A[\vec{a}]$ denote the same collection as does $\dot{A}[\vec{a}]$, but excluding the vertex A itself. Thus, $B[(2,1)]$ refers to $\{v_{10}, v_{11}, v_{19}\}$ in our favorite tree. We combine these new notations with the old ones (namely, $A_m[\vec{a}]$ and $k*$) additively, as before.

A word about notation: we will need to describe the boundaries of critical 2-cells in $U\mathcal{D}^n\Gamma$, and these boundaries sometimes consist of certain 1-cells which cannot be expressed in terms of our notation. For instance, consider $B_2[(1,2)] + D_2[(1,1)]$ in relation to Figure 1. This is the 2-cell $c = \{v_{10}, e_{19}, v_{20}, v_{22}, e_{25}\}$. The 1-cells forming the boundary are obtained by replacing the edges $e_i \in c$, for $i = 19$ or 25, with $\iota(e_i)$ or $\tau(e_i)$. Thus one of the 1-cells on the boundary of c is $\{v_{10}, v_9, v_{20}, v_{22}, e_{25}\}$. Note that this configuration isn't covered by our notation, since the inessential vertex v_{20} is unblocked.

Rather than introduce more notation to deal with these extra configurations, we introduce the idea of a *slide.* A 1-cell c' is obtained from c by a slide (or *c slides to c'*) if there is some unblocked vertex $v \in c$ such that the endpoints, say v and v', of $e(v)$ are consecutive in the order on vertices, and c' is obtained from c by replacing v with v'. For example, $\{v_{10}, v_9, v_{19}, v_{22}, e_{25}\}$ (which is $\dot{B}[(1,1)] + D_2[(1,1)]$) is obtained from $\{v_{10}, v_9, v_{20}, v_{22}, e_{25}\}$ by a slide.

If c slides to c', then we may use them interchangeably in our calculations. For if $c = \{c_1, \ldots, c_{n-2}, v, e\}$ and $c' = \{c_1, \ldots, c_{n-2}, v', e\}$, where $v' = \tau(e(v))$, then c and c' are parallel sides of the square $\{c_1, \ldots, c_{n-2}, e(v), e\}$, and the other sides are $c_\iota = \{c_1, \ldots, c_{n-2}, e(v), \iota(e)\}$ and $c_\tau = \{c_1, \ldots, c_{n-2}, e(v), \tau(e)\}$. It is clear that c_ι and c_τ both satisfy the condition (1) from the proof of the redundant 1-cells lemma, so $c_\iota \dot{\to} 1$ and $c_\tau \dot{\to} 1$. It then follows that the oriented 1-cells c and c' both represent the same element in the usual edge-path presentation of $\pi_1(U\mathcal{D}^n\Gamma)$. In fact, more is true: if we add all relations corresponding to slides (c, c') to the monoid presentation $\mathcal{MP}_{W,T}$, the associated string rewriting system is still complete, and has the same reduced objects. We leave the verification of this fact as an exercise. (Note that the square $\{c_1, \ldots, c_{n-2}, e(v), e\}$ may be redundant, so the proof is not entirely trivial.) In our calculations, we will use slides without further notice.

Let δ_m denote a vector such that the mth component is 1 and every other component is 0. The length of δ_m will be clear from the context. If $\vec{v}$ is a vector having entries in the set of non-negative integers, let $\vec{v} - 1$ be the vector obtained from $\vec{v}$ by subtracting 1 from the first non-zero entry. This last notation must be used carefully to avoid ambiguity – note for instance that $\delta_1 + (\delta_2 - 1) \neq (\delta_1 + \delta_2) - 1$. If $\vec{v}$ is any vector, let $|\vec{v}|$ denote the sum of the entries of $\vec{v}$.

Lemma 5.2 *Let A and B be essential vertices of T, a maximal tree in Γ.*

(1) *Suppose that $A \wedge B = C$ where C is an essential vertex distinct from both A and B. Let $g(C, A) = i$ and $g(C, B) = j$, where $i < j$. Then*

(a) $\left(A[\vec{a}] + B_l[\vec{b}] + p*\right) \dot{\to} \left(B_l[\vec{b}] + (p + |\vec{a}|)*\right)$,

(b) $\left(\dot{A}[\vec{a}] + B_l[\vec{b}] + p*\right) \dot{\to} \left(B_l[\vec{b}] + (p + 1 + |\vec{a}|)*\right)$,

(c) $\left(A_k[\vec{a}] + B[\vec{b}] + p*\right) \dot{\to} w_1 \left(A_k[\vec{a}] + (p + |\vec{b}|)*\right) w_1^{-1}$, *and*

(d) $\left(A_k[\vec{a}] + \dot{B}[\vec{b}] + p*\right) \dot{\to} w_2 \left(A_k[\vec{a}] + (1 + p + |\vec{b}|)*\right) w_2^{-1}$,

where

$$w_1 = \prod_{\alpha=0}^{|\vec{b}|-1} \left(C_j[|\vec{a}|\delta_i + (|\vec{b}| - \alpha)\delta_j] + (p + \alpha)*\right),$$

and w_2 is the same as w_1, but with $|\vec{b}| + 1$ in place of $|\vec{b}|$.

(2) *Suppose that $A \wedge B = A$ and $g(A, B) = i$. Then*

(a) $\left(A_k[\vec{a}] + B[\vec{b}] + p*\right) \dot{\to} \left(A_k[\vec{a} + |\vec{b}|\delta_i] + p*\right)$,

(b) $\left(A_k[\vec{a}] + \dot{B}[\vec{b}] + p*\right) \dot{\to} \left(A_k[\vec{a} + (1 + |\vec{b}|)\delta_i] + p*\right)$,

(c) $\left(A[\vec{a}] + B_l[\vec{b}] + p*\right) \dot{\to} w_3 \left(B_l[\vec{b}] + (p + |\vec{a}|)*\right) w_3^{-1}$, *and*

(d) $\left(\dot{A}[\vec{a}] + B_l[\vec{b}] + p*\right) \dot{\to} \left(A[\vec{a}] + B_l[\vec{b}] + (p+1)*\right)$,

where

$$w_3 = \prod_{\alpha=0}^{|\vec{a}|-1} \left(A_\beta[|\vec{b}|\delta_i + (\vec{a}-\alpha)] + (p+\alpha)*\right),$$

and β is the smallest coordinate of $\vec{a}-\alpha$ that is non-zero. Here a factor in w_3 is considered trivial if $\beta \le i$.

Proof

(1a) Under the given assumptions, the smallest vertex of the subconfiguration $A[\vec{a}]$ may be moved until it is blocked at $*$, by repeated applications of the redundant 1-cells lemma. That is,

$$\left(A[\vec{a}] + B_l[\vec{b}] + p*\right) \dot{\to} \left(A[\vec{a}-1] + B_l[\vec{b}] + (p+1)*\right).$$

After repeated applications of the above identity, we eventually arrive at the statement (a).

(1b) This is similar to (a).

(1c) We begin by applying the redundant 1-cells lemma to the smallest vertex of the subconfiguration $B[\vec{b}]$. This smallest vertex can be moved freely, until it occupies the place adjacent to C, and lying in the jth direction from C. That is,

$$\left(A_k[\vec{a}] + B[\vec{b}] + p*\right) \dot{\to} \left(A_k[\vec{a}] + B[\vec{b}-1] + C[\delta_j] + p*\right).$$

At this point, the redundant 1-cells lemma no longer applies, since all of the vertices in the subconfiguration $A_k[\vec{a}]$ lie between C and the vertex (say v) which is adjacent to C and which lies in the jth direction. We must therefore appeal to the definition of W:

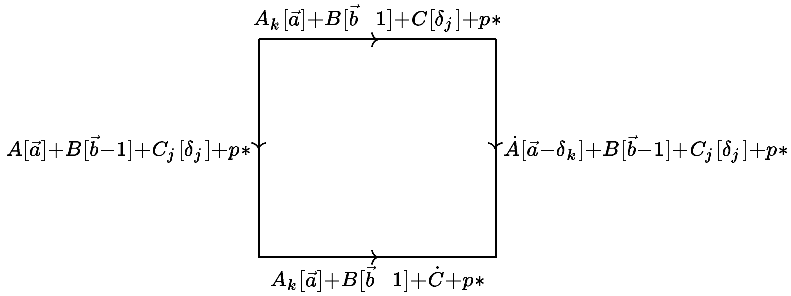

The 2-cell depicted above is $A_k[\vec{a}] + B[\vec{b}-1] + C_j[\delta_j] + p*$, the image under W of the 1-cell $A_k[\vec{a}] + B[\vec{b}-1] + C[\delta_j] + p*$ (located at the top). (Note: the label of the "source" vertex on the upper left can be computed by replacing each of the edges in the 2-cell $A_k[\vec{a}]+B[\vec{b}-1]+C_j[\delta_j]+p*$ with its initial vertex. The "sink" vertex on the bottom right is obtained from the same 2-cell by replacing each edge with its terminal vertex. The other two vertices are determined by replacing one of the edges with its initial vertex, and the other with its terminal vertex. Furthermore, if we specify the identity of any one of the 1-cells on the boundary of the 2-cell, then the labels of the other three 1-cells are uniquely determined.)

Now consider the 1-cell $A[\vec{a}]+B[\vec{b}-1]+C_j[\delta_j]+p*$. The redundant 1-cells lemma applies to the vertices in the subconfiguration $A[\vec{a}]$. These may move until they are blocked at C. Since there are $|\vec{a}|$ vertices in $A[\vec{a}]$, and each lies in the direction i from C, we have

$$\left(A[\vec{a}]+B[\vec{b}-1]+C_j[\delta_j]+p*\right)\dot{\to}\left(B[\vec{b}-1]+C_j[\delta_j+|\vec{a}|\delta_i]+p*\right).$$

If we let the vertices in the subconfiguration $B[\vec{b}-1]$ move, then by the redundant 1-cells lemma, these vertices will flow until they are blocked at C. We get

$$\left(B[\vec{b}-1]+C_j[\delta_j+|\vec{a}|\delta_i]+p*\right)\dot{\to}\left(C_j[|\vec{b}|\delta_j+|\vec{a}|\delta_i]+p*\right).$$

This last 1-cell is critical. The 1-cell on the right side of the square pictured above flows to the same 1-cell, by similar reasoning.

Finally, the 1-cell $A_k[\vec{a}]+B[\vec{b}-1]+\dot{C}+p*$ flows to $A_k[\vec{a}]+B[\vec{b}-1]+(p+1)*$ by the redundant 1-cells lemma. It follows that

$$\left(A_k[\vec{a}]+B[\vec{b}]+p*\right)\dot{\to}w\left(A_k[\vec{a}]+B[\vec{b}-1]+(p+1)*\right)w^{-1},$$

where $w=(C_j[|\vec{a}|\delta_i+|\vec{b}|\delta_j]+p*)$. Part (c) now follows by repeatedly applying the above identity.

(1d) This is similar to (c).

(2a) In this case, the vertices in the subconfiguration $B[\vec{b}]$, beginning with the smallest, can flow by the redundant 1-cells lemma until they are blocked at the essential vertex A. Since the vertices in $B[\vec{b}]$ lie in the direction i from A, when these vertices are blocked, the resulting configuration is $A_k[\vec{a}+|\vec{b}|\delta_i]+p*$. This proves (a).

(2b) This is similar to (a).

(2c) We begin by applying W to the configuration $A[\vec{a}]+B_l[\vec{b}]+p*$. The result is $A_\beta[\vec{a}]+B_l[\vec{b}]+p*$, where β is the smallest subscript for which a_β is non-zero (here $\vec{a} = (a_1, a_2, \ldots, a_{d(A)-1})$):

$$\begin{tikzcd}[column sep=huge, row sep=huge] \cdot \arrow[r, "A[\vec{a}]+B_l[\vec{b}]+p*"] \arrow[d, "A_\beta[\vec{a}]+B[\vec{b}]+p*"'] & \cdot \arrow[d, "A_\beta[\vec{a}]+\dot{B}[\vec{b}-\delta_l]+p*"] \\ \cdot \arrow[r, "\dot{A}[\vec{a}-1]+B_l[\vec{b}]+p*"'] & \cdot \end{tikzcd}$$

Now it is either the case that both of the vertical 1-cells are collapsible (if $\beta \le i$), or these vertical 1-cells both flow to

$$\left(A_\beta[\vec{a} + |\vec{b}|\delta_i] + p*\right).$$

The bottom 1-cell flows to

$$\left(A[\vec{a} - 1] + B_l[\vec{b}] + (p+1)*\right).$$

Thus, $\left(A[\vec{a}] + B_l[\vec{b}] + p*\right)$ flows to

$$w\left(A[\vec{a} - 1] + B_l[\vec{b}] + (p+1)*\right) w^{-1},$$

where $w = \left(A_\beta[\vec{a} + |\vec{b}|\delta_i] + p*\right)$ if $\beta > i$, and $w = 1$ if $\beta \le i$. The statement of 2(c) now follows by repeated application of the above identity.

(2d) This is straightforward. □

Theorem 5.3 *Let Γ be a sufficiently subdivided tree with a chosen basepoint $*$. Suppose that an embedding of Γ in the plane is given, so that there is an induced order on the vertices. Then the braid group $B_n\Gamma$ is generated by the collection of critical 1-cells, and the set of relations consists of the reduced forms of the boundary words $w(c)$, where c is any critical 2-cell. For $c = \left(A_k[\vec{a}] + B_l[\vec{b}] + p*\right)$ a critical 2-cell, the reduced form of the boundary word is as follows:*

(1) *If $A \wedge B = C$ with $C \neq A, B$, $g(C,A) = i$, $g(C,B) = j$, and $i < j$, then a reduced form of the boundary word for c is*

$$\left[\left(B_l[\vec{b}] + (p + |\vec{a}|)*\right), w_1\left(A_k[\vec{a}] + (p + |\vec{b}|)*\right) w_1^{-1}\right],$$

where w_1 is as in the previous lemma.

(2) *If $A \wedge B = A$, and $g(A,B) = i$, then a reduced form of the boundary word for c is*

$$\left[w_3^{-1}\left(A_k[\vec{a} + |\vec{b}|\delta_i] + p*\right) w_3', \left(B_l[\vec{b}] + (p + |\vec{a}|)*\right)^{-1}\right],$$

where w_3 is as in the previous lemma, and w_3' is the same as w_3, but with $\vec{a} - \delta_k$ in place of $\vec{a}$ and $p+1$ in place of p.

Proof It follows from Theorem 2.5 that $B_n\Gamma$ is generated by the critical 1-cells, and that the relations are the reduced forms of the boundary words of the critical 2-cells.

Consider the critical 2-cell $\left(A_k[\vec{a}] + B_l[\vec{b}] + p*\right)$:

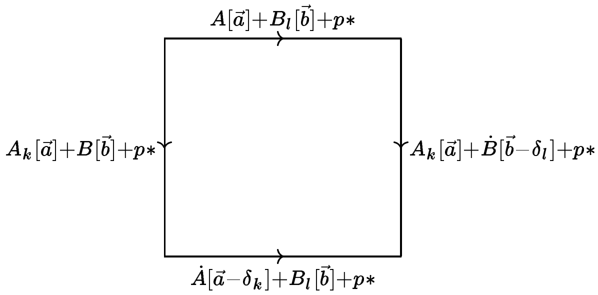

the rest of the statement of the theorem follows by applying the previous lemma to each side. □

Example 5.4 Consider the example of $B_4\Gamma$, where Γ is the tree in Figure 1. This tree is sufficiently subdivided as written. By Theorem 3.6 the critical 1-cells are:

$$(X_2[\vec{v}] + (4 - |\vec{v}|)*),$$

where X is one of the essential vertices A, B, C, or D, and $\vec{v}$ is $(1,1)$, $(1,2)$, $(1,3)$, $(2,1)$, $(2,2)$, or $(3,1)$. Thus, there are a total of 24 critical 1-cells. The critical 2-cells all have the form

$$(X_2[(1,1)] + Y_2[(1,1)]),$$

where X and Y are distinct essential vertices chosen from the set $\{A,B,C,D\}$. Thus there are a total of 6 critical 2-cells.

We now describe the resulting relations. We begin with the critical 2-cell $(A_2[(1,1)] + B_2[(1,1)])$. The boundary word, by (2) of the previous theorem with $i = k = l = 2$ and $p = 0$, is

$$\left[w_3^{-1}\left(A_2[(1,3)]\right) w_3', \left(B_2[(1,1)] + 2*\right)\right].$$

The word w_3 in the present case is

$$\prod_{\alpha=0}^{1} \left(A_\beta[(0,2)+((1,1)-\alpha)]+\alpha*\right).$$

Computing, and using the fact that β is the first non-zero entry of $(1,1)-\alpha$, we get

$$w_3=\left(A_1[(1,3)]\right)\left(A_2[(0,3)]+1*\right)=1,$$

since both of the cells in the above expression are collapsible. The word w_3' in the present case is

$$\prod_{\alpha=0}^{0} \left(A_\beta[(0,2)+((1,0)-\alpha)]+(\alpha+1)*\right)=\left(A_1[(1,2)]+1*\right)=1,$$

since the given cell is again collapsible. Summarizing, we get

$$\left[\left(X_2[(1,3)]\right),\left(Y_2[(1,1)]+2*\right)\right],$$

where $(X,Y)=(A,B)$. A completely analogous computation shows that a boundary word for $(X_2[(1,1)]+Y_2[(1,1)])$ is given by the same expression, where $(X,Y)\in\{(A,B),(A,C),(A,D),(B,D)\}$.

Now consider the critical 2-cell $(B_2[(1,1)]+C_2[(1,1)])$. Part (2) of the previous theorem applies again, with B playing the role of A in the theorem, C the role of B, $i=1$, $k=l=2$, and $p=0$:

$$\left[w_3^{-1}\left(B_2[(3,1)]\right)w_3',\left(C_2[(1,1)]+2*\right)\right].$$

In the present case, we get the following expression for w_3:

$$\begin{aligned} w_3 &= \prod_{\alpha=0}^{1}\left(B_\beta[(2,0)+((1,1)-\alpha)]+\alpha*\right)\\ &= \left(B_1[(3,1)]\right)\left(B_2[(2,1)]+1*\right)\\ &= \left(B_2[(2,1)]+1*\right). \end{aligned}$$

We get the following expression for w_3':

$$\begin{aligned} w_3' &= \prod_{\alpha=0}^{0}\left(B_\beta[(2,0)+((1,0)-\alpha)]+(\alpha+1)*\right)\\ &= \left(B_1[(3,0)]+1*\right)\\ &= 1. \end{aligned}$$

We arrive at the following boundary word for $(B_2[(1,1)]+C_2[(1,1)])$:

$$\left[\left(B_2[(2,1)]+1*\right)^{-1}\left(B_2[(3,1)]\right),\left(C_2[(1,1)]+2*\right)\right].$$

Finally, we consider the critical 2-cell $(C_2[(1,1)] + D_2[(1,1)])$. By part (1) of the previous theorem, with C playing the role of A, D playing the role of B, and B playing the role of C, $i = 1$, $j = k = l = 2$, and $p = 0$, a boundary word has the form

$$\left[(D_2[(1,1)] + 2*), w_1 (C_2[(1,1)] + 2*) w_1^{-1}\right],$$

with

$$\begin{aligned} w_1 &= \prod_{\alpha=0}^{1} (B_2[(2, 2-\alpha)] + \alpha *) \\ &= (B_2[(2,2)]) (B_2[(2,1)] + 1*) . \end{aligned}$$

We summarize these calculations in a figure, which also illustrates a useful chalkboard notation for critical 1-cells (see Figure 4).

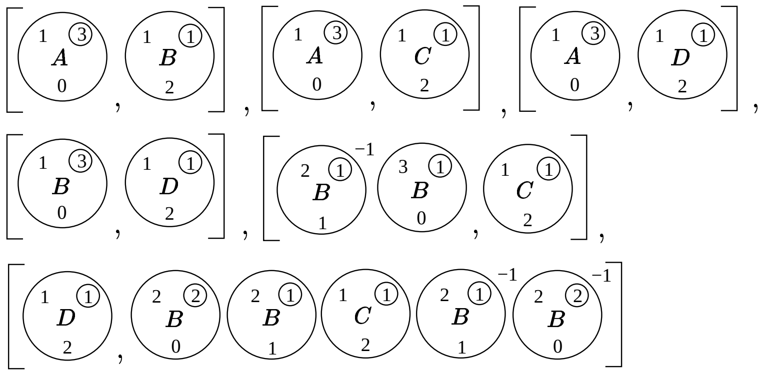

Figure 4: The relations of $B_4\Gamma$ in shorthand notation

For example, the circled letter C with a 1, a circled 1, and a 2 (reading in the clockwise direction from the upper left) represents the element $C_2[(1,1)] + 2*$: the first entry in the upper left refers to the first entry of the vector $(1,1)$, the circled entry in the upper right refers to the second entry of $(1,1)$ and indicates the location of the edge, and the 2 on the bottom refers to the number of vertices clustered near $*$.

A group G is said to be a *right-angled Artin group* provided there is a presentation $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ such that every relation in $\mathcal{R}$ has the form $[a,b]$, where $a, b \in \Sigma$. It is common to describe a right-angled Artin group presentation by a graph $\Gamma_{\mathcal{P}}$, where the vertices of $\Gamma_{\mathcal{P}}$ are in one-to-one correspondence with elements of Σ, and there is an edge connecting two vertices $a, b \in \Sigma$ if and only if $[a,b] \in \mathcal{R}$.

A finitely presented group G is *coherent* if every finitely generated subgroup of G is also finitely presentable.

Example 5.5 Consider the case in which Γ is a tree homeomorphic to the capital letter "H". Let the basepoint $*$ be the vertex on the bottom left, let A be the essential vertex on the left, and let B be the essential vertex on the right.

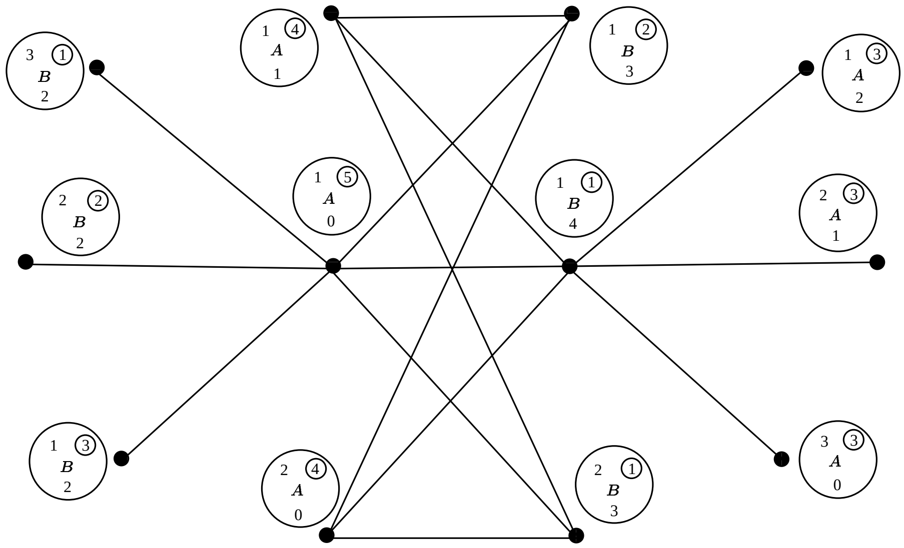

Figure 5: Part of the presentation for B_6H

We claim that $B_n\Gamma$ is a right-angled Artin group, for any n. (See also Connolly and Doig [9], who prove, more generally, that the braid group of any linear tree is a right-angled Artin group. A tree T is *linear* if every essential vertex of T lies along a single embedded arc.) No matter what n is, part (2) of Theorem 5.3 applies, and we must compute the values of the words w_3 and w_3'. In fact, we prove that both of these words will necessarily be trivial. Consider

$$w_3 = \prod_{\alpha=0}^{|\vec{a}|-1} \left(A_\beta \left[|\vec{b}|\delta_i + (\vec{a} - \alpha) \right] + (p + \alpha) * \right).$$

Here $i = g(A, B) = 2$ and β is the subscript of the smallest non-zero entry of $(\vec{a} - \alpha)$. The crucial observation is that β is (in the present case) also the smallest non-zero entry of $|\vec{b}|\delta_i + (\vec{a} - \alpha)$, so, by Proposition 3.9, no term in the above product is a critical 1-cell (in fact, it is easy to see that each is collapsible). The triviality of w_3 follows. Essentially the same reasoning holds for w_3'.

If $n=6$, a routine but lengthy calculation shows that $B_n\Gamma$ is the free product of a free group of rank 18 with the right-angled Artin group in Figure 5.

The copy of $K_{3,3}$ in this graph represents a subgroup isomorphic to $F_3 \times F_3$, the direct product of the free group of rank 3 with itself. Since $F_3 \times F_3$ is not coherent (see, for instance, [5]) and has an unsolvable generalized word problem [16], there are tree braid groups which are not coherent and have an unsolvable generalized word problem.

References

[1] **A Abrams**, *Private Communication*

[2] **A Abrams**, *Configuration spaces of braid groups of graphs*, PhD thesis, UC Berkeley (2000)

[3] **A Abrams**, *Configuration spaces of colored graphs*, Geom. Dedicata 92 (2002) 185–194

[4] **A Abrams**, **R Ghrist**, *Finding topology in a factory: configuration spaces*, Amer. Math. Monthly 109 (2002) 140–150

[5] **M Bestvina**, **N Brady**, *Morse theory and finiteness properties of groups*, Invent. Math. 129 (1997) 445–470

[6] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundlehren series 319, Springer-Verlag, Berlin (1999)

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

[8] **M M Cohen**, *A course in simple-homotopy theory*, Springer-Verlag, New York (1973)

[9] **F Connolly**, **M Doig**, *Braid groups and right angled Artin groups*, `arXiv:math.GT/0411368`

[10] **J Crisp**, **B Wiest**, *Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups*, Algebr. Geom. Topol. 4 (2004) 439–472

[11] **Michael Farber**, *Collision free motion planning on graphs*, e-print `arXiv:math.AT/0406361`

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

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

[14] **R Ghrist**, *Configuration spaces and braid groups on graphs in robotics*, from: "Knots, braids, and mapping class groups—papers dedicated to Joan S. Birman (New York, 1998)", AMS/IP Stud. Adv. Math. 24, Amer. Math. Soc., Providence, RI (2001) 29–40

[15] **R W Ghrist**, **D E Koditschek**, *Safe cooperative robot dynamics on graphs*, SIAM J. Control Optim. 40 (2002) 1556–1575

[16] **K A Mihaĭlova**, *The occurrence problem for direct products of groups*, Mat. Sb. (N.S.) 70 (112) (1966) 241–251

[17] **M H A Newman**, *On theories with a combinatorial definition of "equivalence"*, Ann. of Math. (2) 43 (1942) 223–243

[18] **C P Rourke**, **B J Sanderson**, *Introduction to piecewise-linear topology*, Springer Study Edition, Springer-Verlag, Berlin (1982)

[19] **J C Stillwell**, *Classical topology and combinatorial group theory*, Graduate Texts in Mathematics 72, Springer-Verlag, New York (1980)

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

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

Email: `farley@math.uiuc.edu, sabalka@math.uiuc.edu`

URL: `www.math.uiuc.edu/~farley` and `www.math.uiuc.edu/~sabalka`

Received: 26 October 2004