DIMENSION OF THE TORELLI GROUP FOR $\mathrm{Out}(F_n)$

MLADEN BESTVINA, KAI-UWE BUX, AND DAN MARGALIT

May 18, 2007

Abstract. Let $\mathcal{T}_n$ be the kernel of the natural map $\mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$. We use combinatorial Morse theory to prove that $\mathcal{T}_n$ has an Eilenberg–MacLane space which is $(2n-4)$-dimensional and that $H_{2n-4}(\mathcal{T}_n, \mathbb{Z})$ is not finitely generated ($n \geq 3$). In particular, this shows that the cohomological dimension of $\mathcal{T}_n$ is equal to $2n-4$ and recovers the result of Krstić–McCool that $\mathcal{T}_3$ is not finitely presented. We also give a new proof of the fact, due to Magnus, that $\mathcal{T}_n$ is finitely generated.

1. Introduction

There is a natural homomorphism from $\mathrm{Out}(F_n)$, the group of outer automorphisms of the free group on n generators, to $\mathrm{GL}_n(\mathbb{Z})$, given by abelianizing the free group F_n. It is a theorem of Nielsen that this map is surjective [15]. We call its kernel the *Torelli subgroup* of $\mathrm{Out}(F_n)$, and we denote it by $\mathcal{T}_n$:

$$1 \to \mathcal{T}_n \to \mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z}) \to 1$$

Main Theorem. *Let $n \geq 3$.*

1. *$\mathcal{T}_n$ has a $(2n-4)$-dimensional Eilenberg–MacLane space.*
2. *$H_{2n-4}(\mathcal{T}_n, \mathbb{Z})$ is infinitely generated.*
3. *$\mathcal{T}_n$ is finitely generated.*

Let cd and vcd denote cohomological dimension and virtual cohomological dimension, respectively; see [16] for background. The main theorem implies that $\mathrm{cd}(\mathcal{T}_n) = 2n-4$.

Part (3) of the main theorem is due to Magnus; we give our own proof in Section 5. We remark that $\mathcal{T}_1$ is obviously trivial and $\mathcal{T}_2$ is trivial by a classical result of Nielsen [15] (we give a new proof of the latter fact in Section 5).

The group $\mathcal{T}_n$, like any torsion free subgroup of $\mathrm{Out}(F_n)$, acts freely on the spine for Outer space (see Section 2), and therefore has an Eilenberg–MacLane space of

2000 *Mathematics Subject Classification.* Primary: 20F36; Secondary: 20F28.
Key words and phrases. *Out*(F_n), Torelli group, cohomological dimension.
The first and third authors gratefully acknowledge support by the National Science Foundation.

dimension $2n-3$, the dimension of this spine. Our theorem improves this upper bound on the dimension and shows that $2n-4$ is sharp.

When $n=3$, we obtain that $H_2(\mathcal{T}_3,\mathbb{Z})$ is not finitely generated, and this immediately implies the result of Krstić–McCool that $\mathcal{T}_3$ is not finitely presented [10].

Historical background. The question of whether $H_k(\mathcal{T}_n,\mathbb{Z})$ is finitely generated, for various values of k and n, is a long standing problem with few solutions. This question was explicitly asked by Vogtmann in her survey article [19]. We now give a brief history of related results, most of which are recovered by our main theorem.

Nielsen proved in 1924 that $\mathcal{T}_3$ is finitely generated [15]. Ten years later, Magnus proved that $\mathcal{T}_n$ is finitely generated for every n [11].

In 1986, Culler–Vogtmann proved that $\mathrm{vcd}(\mathrm{Out}(F_n)) = 2n-3$ when $n \geq 2$ (and so $\mathrm{cd}(\mathcal{T}_n) \leq 2n-3$). Then, Smillie–Vogtmann proved in 1987 that, if $2 < n < 100$ or $n > 2$ is even, then $H_\star(\mathcal{T}_n,\mathbb{Z})$ is not finitely generated [17] [18]. Their method is to consider the rational Euler characteristics of the groups in the short exact sequence defining $\mathcal{T}_n$ (see [19]).

The Krstić–McCool result that $\mathcal{T}_3$ is not finitely presented was proven in 1997, via completely algebraic methods [10]. It is a general fact that if the second homology of a group is not finitely generated, then the group is not finitely presented.

Large abelian subgroups. It follows from the second part of the main theorem that the first part is sharp; i.e., $\mathcal{T}_n$ does not have an Eilenberg–MacLane space of dimension less than $2n-4$. A simpler proof that $\mathrm{cd}(\mathcal{T}_n) \geq 2n-4$ is to exhibit an embedding of $\mathbb{Z}^{2n-4}$ into $\mathcal{T}_n$. There is a subgroup $G \cong \mathbb{Z}^{2n-4} < \mathcal{T}_n$ consisting of elements with representative automorphisms given by

$$\begin{array}{rcl} x_1 & \mapsto & x_1 \\ x_2 & \mapsto & x_2 \\ x_3 & \mapsto & [x_1,x_2]^{p_3} x_3 [x_1,x_2]^{q_3} \\ & \vdots & \\ x_n & \mapsto & [x_1,x_2]^{p_n} x_n [x_1,x_2]^{q_n} \end{array}$$

where p_i and q_i range over all integers (the x_i are generators for F_n).

In Section 7, we prove that specific conjugates of G represent independent classes in $H_{2n-4}(\mathcal{T}_n,\mathbb{Z})$, thus proving the second part of the main theorem. These conjugates generate a subgroup of $H_{2n-4}(\mathcal{M}_n,\mathbb{Z})$, where $\mathcal{M}_n$, called the "toy model", is a particularly simple subcomplex of the Eilenberg–MacLane space $\mathcal{Y}_n$ defined in Section 2. In Section 7, we prove that this infinitely generated subgroup of $H_{2n-4}(\mathcal{M}_n,\mathbb{Z})$ injects into $H_{2n-4}(\mathcal{Y}_n,\mathbb{Z})$. Since $H_k(\mathcal{M}_n,\mathbb{Z})$ is not finitely generated for $2 \leq k \leq 2n-4$, we are led to the following question.

Question. *Does $H_*(\mathcal{M}_n,\mathbb{Z})$ inject into $H_*(\mathcal{Y}_n,\mathbb{Z})$?*

Mapping class groups. The term "Torelli group" comes from the theory of mapping class groups. Let Σ_g be a closed surface of genus $g \geq 1$. The *mapping class group* of Σ_g, denoted $\mathrm{Mod}(\Sigma_g)$, is the group of isotopy classes of orientation preserving homeomorphisms of Σ_g. The *Torelli group*, $\mathcal{I}_g$, is the subgroup of $\mathrm{Mod}(\Sigma_g)$

acting trivially on the homology of Σ_g. As $\mathrm{Mod}(\Sigma_g)$ acts on $H_1(\Sigma_g, \mathbb{Z})$ by symplectic automorphisms, $\mathcal{I}_g$ is defined by:

$$1 \to \mathcal{I}_g \to \mathrm{Mod}(\Sigma_g) \to \mathrm{Sp}_{2g}(\mathbb{Z}) \to 1$$

It is a classical theorem of Dehn, Nielsen, and Baer that the natural map $\mathrm{Mod}(\Sigma_g) \to \mathrm{Out}(\pi_1(\Sigma_g))$ is an isomorphism. In this sense, $\mathcal{T}_n$ is the direct analog of $\mathcal{I}_g$.

Our (lack of) knowledge of the finiteness properties of $\mathcal{I}_g$ mirrors that for $\mathcal{T}_n$. Using the fact that $\mathrm{Mod}(\Sigma_1) \cong \mathrm{SL}_2(\mathbb{Z}) = \mathrm{Sp}_2(\mathbb{Z})$, it is obvious that $\mathcal{I}_1$ is trivial. Johnson showed in 1983 that $\mathcal{I}_g$ is finitely generated for $g \geq 3$ [9]. In 1986, McCullough–Miller showed that $\mathcal{I}_2$ is not finitely generated [12], and Mess improved on this in 1992 by showing that $\mathcal{I}_2$ is a free group of infinite rank [14]. At the same time, Mess further showed that $H_3(\mathcal{I}_3, \mathbb{Z})$ is not finitely generated. In another paper, Mess proved that $\mathrm{cd}(\mathcal{I}_g) \geq 3g - 5$ [13]. Akita showed that $H_\star(\mathcal{I}_g, \mathbb{Z})$ is not finitely generated for $g \geq 7$ [2]. In Kirby's problem list, Mess asked about finiteness properties in higher genus [1].

Our main theorem has the following corollary for $\mathcal{I}_g$.

Corollary 1.1. *For $g \geq 1$, $\mathrm{cd}(\mathcal{I}_g) < \mathrm{vcd}(\mathrm{Mod}(\Sigma_g))$.*

Since $\mathrm{vcd}(\mathrm{Mod}(\Sigma_g)) = 4g - 5$, we have $\mathrm{cd}(\mathcal{I}_g) \leq 4g - 6$. Compare with the lower bound of $3g - 5$ given by Mess.

Actually, the immediate corollary of the main theorem is that $\mathrm{cd}(\mathcal{I}_{g,1})$ is strictly less than $\mathrm{vcd}(\mathrm{Mod}(\Sigma_{g,1}))$, where $\Sigma_{g,1}$ is a once punctured surface, and $\mathcal{I}_{g,1}$ is the corresponding Torelli group. The keys are that $\mathrm{vcd}(\mathrm{Mod}(\Sigma_{g,1})) = \mathrm{vcd}(\mathrm{Out}(F_{2g}))$ and $\mathcal{I}_{g,1} < \mathcal{T}_{2g}$ (the identification is made by retracting $\Sigma_{g,1}$ to a 1-dimensional spine) [6, 7]. To pass from $\Sigma_{g,1}$ to Σ_g one needs the fact that the dimensions of both the Torelli group and the mapping class group increase by 2 when the puncture is added, c.f. [13].

Automorphisms vs. outer automorphisms. Strictly speaking, Magnus and Krstić–McCool study the group $\mathcal{K}_n$, by which we mean the kernel of $\mathrm{Aut}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$, where $\mathrm{Aut}(F_n)$ is the automorphism group of the free group. By considering the short exact sequence

$$1 \to F_n \to \mathcal{K}_n \to \mathcal{T}_n \to 1$$

we see that $\mathcal{K}_n$ is finitely generated if and only if $\mathcal{T}_n$ is finitely generated. Moreover, it follows from our main theorem and the spectral sequence associated to this short exact sequence that $H_{2n-3}(\mathcal{K}_n, \mathbb{Z})$ is not finitely generated and if k is the smallest index so that $H_k(\mathcal{T}_n, \mathbb{Z})$ is not finitely generated, then $H_k(\mathcal{K}_n, \mathbb{Z})$ is not finitely generated.

From the topological point of view, $\mathcal{T}_n$ is the more natural group to study.

In the literature, $\mathcal{T}_n$ is sometimes denoted by IA_n for "identity on abelianization" (see, e.g. [19]). However, since Krstić–McCool use IA_n to denote the kernel of $\mathrm{Aut}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$, we avoid this notation to eliminate the confusion. The notation $\mathcal{K}_n$ comes from Magnus [11].

Acknowledgements. We would like to thank Bob Bell, Misha Gromov, Jon McCammond, Kevin Wortman, and the referee for helpful comments. We are especially grateful to Karen Vogtmann for explaining her unpublished work.

2. An Eilenberg–MacLane space

In Section 2.1, we recall the definition of Culler–Vogtmann's spine for Outer space. Then, in Section 2.2, we describe the quotient of this space by $\mathcal{T}_n$. For $n \geq 2$, this quotient is a $(2n-3)$-dimensional Eilenberg–MacLane space for $\mathcal{T}_n$.

A *rose* is a graph with one vertex. The *standard rose* in rank n, denoted R_n, is a particular rose which is fixed once and for all. We denote the standard generators of $F_n \cong \pi_1(R_n)$ by $x_1, \ldots, x_n$.

2.1. Spine for Outer space.

Culler–Vogtmann introduced the *spine for Outer space*, which we denote by $\mathcal{X}_n$, as a tool for studying $\mathrm{Out}(F_n)$ [6]. This is a simplicial complex defined in terms of marked graphs.

A *metric* on a graph Γ is a function from the edges of Γ to $(0, \infty)$, and an *isometry* between metric graphs is a homeomorphism which preserves the metric.

A *marked graph* is a pair (Γ, g), where Γ is a finite metric graph with no separating edges and no vertices of valence less than 3 and $g : R_n \to \Gamma$ is a homotopy equivalence (g is called the *marking*). We say that two marked graphs (Γ, g) and (Γ', g') are *equivalent* if $g' \circ g^{-1}$ is homotopic to an isometry, where g^{-1} is any homotopy inverse of g. We will denote the equivalence class $[(\Gamma, g)]$ by (Γ, g).

The vertices of $\mathcal{X}_n$ are equivalence classes of marked graphs where all edges have length 1. A set of vertices

$$\{(\Gamma_1, g_1), \ldots, (\Gamma_k, g_k)\}$$

is said to span a simplex if Γ_{i+1} is obtained from Γ_i by collapsing a forest in Γ_i, and g_{i+1} is the marking obtained from g_i via this operation.

We can think of arbitrary points of $\mathcal{X}_n$ as marked metric graphs: for instance, as we move along a 1-cell between two vertices in $\mathcal{X}_n$, the length of some edge in the corresponding graphs (more generally, the lengths of the edges in a forest) varies between 0 and 1. Thus, the metric of an arbitrary point of $\mathcal{X}_n$ has the following description: outside of some maximal tree, all edges have length 1, and edges in the maximal tree have length less than or equal to 1.

There is a natural right action of $\mathrm{Out}(F_n)$ on $\mathcal{X}_n$. Namely, given $\phi \in \mathrm{Out}(F_n)$ and $(\Gamma, g) \in \mathcal{X}_n$, the action is given by:

$$(\Gamma, g) \cdot \phi = (\Gamma, g \circ \phi)$$

(here we are using the fact that every element ϕ of $\mathrm{Out}(F_n)$ can be realized by a homotopy equivalence $R_n \to R_n$, also denoted ϕ, uniquely up to homotopy).

Culler–Vogtmann proved the following result [6]:

Theorem 2.1. *The space $\mathcal{X}_n$ is contractible.*

This theorem has the consequence that, for $n \geq 2$, the virtual cohomological dimension of $\mathrm{Out}(F_n)$ is equal to $2n-3$, the dimension of $\mathcal{X}_n$.

The *star of a rose* in $\mathcal{X}_n$ is the union of the closed simplices containing the vertex corresponding to a rose. The key idea for Theorem 2.1 is to think of $\mathcal{X}_n$ as the union of stars of vertices corresponding to marked roses. We take an analogous approach in this paper.

2.2. **The quotient.** Baumslag–Taylor proved that $\mathcal{T}_n$ is torsion free [3]. Also, the action of $\mathcal{T}_n$ on $\mathcal{X}_n$ is free: by the definition of the action, point stabilizers are graph isometries, and hence finite. Finally, the action is simplicial, and so it follows that the quotient of $\mathcal{X}_n$ by $\mathcal{T}_n$ is an Eilenberg–MacLane space for $\mathcal{T}_n$:

$$\mathcal{Y}_n = \mathcal{X}_n/\mathcal{T}_n$$

Homology markings. Since $\mathrm{Out}(F_n)$ identifies every pair of isometric graphs of $\mathcal{X}_n$, points of $\mathcal{Y}_n$ can be thought of as equivalence classes of pairs (Γ, g), where Γ is a metric graph (as before), and g is a *homology marking*; that is, g is an equivalence class of homotopy equivalences $R_n \to \Gamma$, where two homotopy equivalences are equivalent if (up to isometries of Γ) they induce the same map $H_1(R_n, \mathbb{Z}) \to H_1(\Gamma, \mathbb{Z})$.

Via the marking g, we can think of the (oriented) edges of Γ as elements of $H^1(R_n) \cong \mathbb{Z}^n$ (if e is an edge and x is a simplicial 1-chain, then $e(x)$ is the number of times e appears in x). As such, if we think of the generators $x_1, \ldots, x_n$ of $\pi_1(R_n)$ as elements of $H_1(R_n, \mathbb{Z})$, then we can *label* each oriented edge a of Γ by the corresponding row vector:

$$(a(g(x_1)), \ldots, a(g(x_n)))$$

where $a(g(x_i))$ is the number of times $g(x_i)$ runs over a homologically, with sign.

In this way, a point of $\mathcal{Y}_n$ is given by a labelled graph, and two such graphs represent the same point in $\mathcal{Y}_n$ if and only if there is a label preserving graph isometry between them. See Figure 1 for an example of a labelled graph. We remark that this example exhibits the fact that $\mathcal{Y}_n$ is not a simplicial complex—there are two edge collapses (and hence two 1-cells in $\mathcal{Y}_n$) taking this vertex to the rose with the identity marking.

Motion. An intuitive way to think about the above union of 1-cells is to imagine the loop labelled $(0,1,0)$ as "moving" around the loop labelled $(0,0,1)$ (in this case, we can also think of the loop labelled $(1,0,0)$ as the one in motion).

When convenient, we will confuse the points of $\mathcal{Y}_n$ with the corresponding marked graphs and the pictorial representations as labelled graphs.

We will make use of the following generalities about marked graphs in $\mathcal{Y}_n$.

Proposition 2.2. *Let (Γ, g) be a marked graph.*

(1) *If an edge of Γ is collapsed, the labels of the remaining edges do not change.*

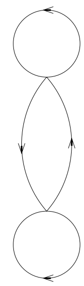

FIGURE 1. An example of a labelled graph.

(2) *Any two edges of* Γ *with the same label (up to sign) are* parallel *in the sense that the union of their interiors disconnects* Γ.
(3) *The sum of the labels of the (oriented) edges coming into a vertex of* Γ *is equal to the sum of the labels of the edges leaving the vertex.*

We leave the proofs to the reader.

Roses. Let (Γ, g) be a rose in $\mathcal{Y}_n$. Up to isometries of Γ, the marking g gives an element of $\mathrm{GL}_n(\mathbb{Z})$, called the *marking matrix*; the rows are exactly the labels of the edges.

Since all edges have length 1, the isometry group of Γ is generated by swapping edges and by reversing the orientations of edges; the former operation has the effect of switching rows of the matrix, and the latter corresponds to changing signs of rows. Thus, in this case, (Γ, g) gives rise to an element of $W\backslash \mathrm{GL}_n(\mathbb{Z})$, where $W = W_n$ is the signed permutation subgroup of $\mathrm{GL}_n(\mathbb{Z})$, acting on the left. In fact, this gives a bijection between roses in $\mathcal{Y}_n$ and elements of $W\backslash \mathrm{GL}_n(\mathbb{Z})$, as $\mathrm{Out}(F_n)$ acts transitively on the roses of $\mathcal{X}_n$.

The right action of $\mathrm{Out}(F_n)$ on $\mathcal{X}_n$ descends to a right action of $\mathrm{GL}_n(\mathbb{Z})$ on $\mathcal{Y}_n$. In particular, the action on roses is given by the right action of $\mathrm{GL}_n(\mathbb{Z})$ on $W\backslash \mathrm{GL}_n(\mathbb{Z})$.

3. STARS OF ROSES

As with $\mathcal{X}_n$, we like to think of the quotient $\mathcal{Y}_n$ as the union of stars of roses. By definition, the *star of a rose* in $\mathcal{Y}_n$ is the image of the star of a rose in $\mathcal{X}_n$. Thus, it consists of graphs which can be collapsed to a particular rose. We now discuss some of the basic properties of the star of a rose.

3.1. **Labels in the star of a rose.**

We will need several observations about the behavior of labels in the star of a rose. The proofs of the various parts of the propositions are straightforward and are left to the reader. In each of the statements, let ρ be a rose in $\mathcal{Y}_n$ represented by a marked graph (Γ, g). Say that its edges $a_1, \ldots, a_n$ are labelled by $v_1, \ldots, v_n \in \mathbb{Z}^n$. If the label of an edge is $\sum k_i v_i$ with $k_j \neq 0$, we say that the label *contains* v_j.

Proposition 3.1. *If* (Γ', g') *is a marked graph in* $\mathrm{St}(\rho)$, *we have:*

(1) *For each i, there is an edge of Γ' labelled $\pm v_i$.*
(2) *The edges not labelled $\pm v_i$ form a forest.*
(3) *The union of edges with label containing v_i (for any particular i) is a topological circle (see, e.g., Figure 1).*
(4) *The label of any edge of Γ' is of the form*

$$\sum_{i=1}^{n} k_i v_i$$

where $k_i \in \{-1, 0, 1\}$.

We have the following converse to the first two parts of the previous proposition.

Proposition 3.2. *If (Γ', g') is a marked graph which has, for each i, at least one edge of length 1 labelled $\pm v_i$, then (Γ', g') is in $\mathrm{St}(\rho)$.*

We also have a criterion for when a marked graph is in the frontier of the star of a rose.

Proposition 3.3. *A marked graph (Γ', g') in $\mathrm{St}(\rho)$ is in the frontier of $\mathrm{St}(\rho)$ if and only if it has at least one edge of length 1 whose label is not $\pm v_i$ for any i. In this case, the given label is a label for some rose whose star contains (Γ', g').*

3.2. **Ideal edges.** Let $\rho = (\Gamma, g)$ be a rose whose edges $a_1, \ldots, a_n$ are labelled $v_1, \ldots, v_n$, as above. An *ideal edge* is any formal sum:

$$\sum_{i=1}^{n} k_i a_i$$

where $k_i \in \{-1, 0, 1\}$, and at least two of the k_i are nonzero. An ideal edge is a "direction" in $\mathrm{St}(\rho)$ in the following sense: for any ideal edge, we can find a marked graph (Γ', g') in the frontier of $\mathrm{St}(\rho)$ where one of the edges of (Γ', g') has the label $\sum k_i v_i$. If a marked graph in $\mathrm{St}(\rho)$ has an edge of length 1 with label $\sum k_i v_i$, we say that the marked graph *realizes* the ideal edge $\sum k_i a_i$.

Lemma 3.4. *Given any ideal edge for a particular rose, there is a 1-edge blowup of ρ in the frontier of $\mathrm{St}(\rho)$ which realizes that ideal edge.*

Proof. As above, let ρ be a rose with edges $\{a_i\}$ labelled $\{v_i\}$ and consider the ideal edge $\sum k_i a_i$. We construct the desired graph explicitly, starting with a single oriented edge e. For $1 \leq i \leq n$, we glue an oriented edge e_i labelled v_i to e according to the value of k_i. If $k_i = 0$, we glue both endpoints of e_i to a single vertex of e; if $k_i = 1$, we glue the original vertex of e_i to the terminal vertex of e and the terminal vertex of e_i to the original vertex of e, and if $k_1 = -1$, we do the reverse. By Proposition 2.2(3) the label of e is $\sum k_i v_i$. See Figure 2 for a picture of a 1-edge blowup realizing the ideal edge $a_1 - a_3 + a_4$ in rank 5. $\square$

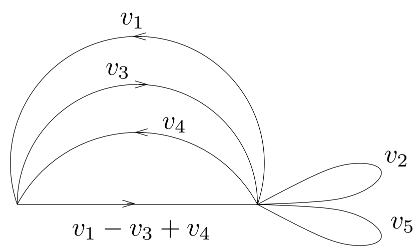

FIGURE 2. A 1-edge blowup realizing the ideal edge $a_1 - a_3 + a_4$.

Our notion of an ideal edge is simply the homological version of the ideal edges of Culler–Vogtmann [6].

We say that an ideal edge ι' is *subordinate* to the ideal edge $\iota = \sum k_i a_i$ if ι' is obtained by changing some of the k_i to zero. A *2-letter ideal edge* is an ideal edge of the form $k_i a_i + k_j a_j$. Two ideal edges are said to be *opposite* if one can be obtained from the other by changing the sign of exactly one nonzero coefficient. The following facts are used in Section 5:

Lemma 3.5. *Let $\rho = (\Gamma, g)$ be a rose whose edges $a_1, \ldots, a_n$ are labelled $v_1, \ldots, v_n$. Suppose that ι and ι' are ideal edges and that either*

1. *ι' is subordinate to ι, or*
2. *ι and ι' are 2-letter ideal edges which are not opposite.*

In either case, there is a marked graph (Γ', g') in $\mathrm{St}(\rho)$ which simultaneously realizes ι and ι'.

Proof. In each case, we can explicitly describe the desired graph. If ι' is subordinate to the ideal edge $\iota = \sum k_i a_i$, we start with a 1-edge blowup realizing ι (Lemma 3.4), and then blow up another edge to separate the edges which appear in ι' from those which do not. Figure 3 (left hand side) demonstrates this for $\iota = v_1 + v_2 + v_3 + v_4$ and $\iota' = v_1 + v_2$ in rank 4.

For the case of two 2-letter ideal edges which are not opposite, without loss of generality it suffices to demonstrate marked graphs which simultaneously realize $v_1 + v_2$ with $-v_1 - v_2$, $v_2 + v_3$, or $v_3 + v_4$ (the arbitrary case is obtained by renaming/reorienting edges and by attaching extra 1-cells to any vertex). See Figure 3 (right hand side) for a demonstration. One can use Proposition 2.2(3) to verify the labels. □

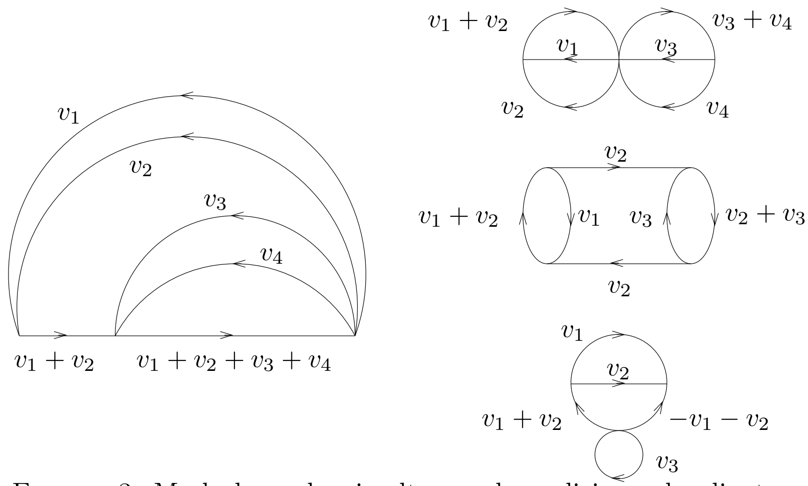

FIGURE 3. Marked graphs simultaneously realizing subordinate ideal edges (left) and 2-letter ideal edges which are not opposite (right).

The reader may verify that opposite ideal edges are never simultaneously realized.

We remark that, in the framework established by Culler–Vogtmann, one can think of this lemma in terms of compatibility of partitions, in which case the proof is immediate; see [6].

3.3. **Homotopy type.** In the remainder of this section, we prove that the star of any rose retracts onto the subcomplex consisting of "cactus graphs", and this subcomplex is homeomorphic to a union of $(n-2)$-tori.

We define a *rank n cactus graph* inductively as follows. When $n \leq 2$, a rank n cactus graph is a graph with one vertex and n edge(s). For $n \geq 3$, a rank n cactus graph is a point of $\mathcal{Y}_n$ obtained as follows: we start with a rank $n-1$ cactus graph, choose any point on that graph, declare this point to be a vertex v (by subdividing) if it is not already a vertex, take a rank 1 cactus graph, and glue its vertex to the vertex v. We note that a rank n cactus graph has exactly n embedded circles, and every edge belongs to exactly one embedded circle (Figure 1 is an example).

Let $\mathcal{C}(\rho)$ denote the space of cactus graphs in $\mathrm{St}(\rho)$. Given any rose ρ', there is a canonical homeomorphism $\mathcal{C}(\rho) \to \mathcal{C}(\rho')$, once we choose orderings of the edges of ρ and ρ'. Thus, we can unambiguously use $\mathcal{C}_n$ to denote the space of cactus graphs in the star of a rose in rank n.

In the remainder, assume that $\rho = (\Gamma, g)$ is a rose in $\mathcal{Y}_n$ with edges $a_1, \ldots, a_n$ labelled $v_1, \ldots, v_n$.

Lemma 3.6. *$\mathrm{St}(\rho)$ strongly deformation retracts onto $\mathcal{C}(\rho)$.*

Proof. For every marked graph in $\mathrm{St}(\rho)$, the set of edges whose label is not $\pm v_i$ is a forest (Proposition 3.1(2)). We perform a strong deformation retraction of $\mathrm{St}(\rho)$ by shrinking the edges of each such forest in each marked graph in $\mathrm{St}(\rho)$.

Consider any marked graph (Γ', g') in the image of the retraction. By Propositions 3.1(1) and 3.1(3), there is a circle of edges labelled $\pm v_i$ for each i. We consider the "dual graph" obtained by assigning a vertex to each such circle (the *circle vertices*) and each intersection point (the *point vertices*) and we connect a point vertex to a circle vertex if the point is contained in the circle. It follows from Proposition 2.2(2) that this graph is a tree, and hence (Γ', g') is a cactus graph. □

Corollary 3.7. *For $n \geq 2$, the star of any rose $\mathrm{St}(\rho)$ in $\mathcal{Y}_n$ is homotopy equivalent to a complex of dimension $n-2$.*

Proof. By the definition of cactus graphs, we can see that the dimension of $\mathcal{C}_n$ increases with slope 1 with respect to n, starting at $n = 2$. Since $\mathcal{C}_2$ is a point, $\mathcal{C}_n$ is a complex of dimension $n-2$. An application of Lemma 3.6 completes the proof. □

We can filter $\mathcal{C}(\rho)$ by subsets according to the number of vertices in the cactus graphs:

$$\{\rho\} = V_0 \subset V_1 \subset \cdots \subset V_{n-2} = \mathcal{C}(\rho)$$

Each V_k consists of cactus graphs with at most $k+1$ vertices.

Our goal now is to give a generating set for $\pi_1(\mathcal{C}(\rho))$. Since V_2 is simple to understand, the following proposition will make it easy to do this.

Proposition 3.8. *There is a cell structure on $\mathcal{C}(\rho)$ so that the k-skeleton is exactly V_k.*

Proof. We declare the 0-skeleton to be the single point $V_0 = \{\rho\}$. Given any $k > 1$, and any marked graph (Γ', g') of $V_k - V_{k-1}$, we define a unique embedding of a k-dimensional product of simplices C into V_k, so that the boundary of C maps to V_{k-1}.

For each i, let n_i be the number of edges of Γ' labelled $\pm v_i$. Let C_i be the standard $(n_i - 1)$-simplex in $\mathbb{R}^{n_i}$ given by $\sum x_i = 1$. We obtain an embedding of C_i into V_k by sending the point $(x_1, \ldots, x_{n_i})$ to the graph obtained from Γ' by assigning the length $x_j / \max\{x_m\}$ to the j^{th} edge of Γ' labelled $\pm v_i$.

We denote the product of the C_i by C. By combining the maps defined on the C_i, we get an embedding of C into V_k. The dimension of C is $\sum(n_i - 1)$, which, by a straightforward count, is equal to k.

The boundary of C is the set of points where some edge is assigned length 0; this must be an edge labelled $\pm v_i$ with $n_i > 1$. Clearly, $\partial C \subset V_{k-1}$, and so the proposition follows. □

For the remainder of this section, we use the cell structure for $\mathcal{C}(\rho)$ given by Proposition 3.8, which is different from the cell structure inherited from $\mathcal{Y}_n$.

Let V_1^1 be the subset of V_1 consisting of graphs with a vertex of valence 4 and a vertex of valence $2n-2$ (i.e. only a single loop is "moving" around another loop). We will see in Section 5 that the obvious generators for $\pi_1(V_1^1)$ correspond to one of the two types of Magnus generators for $\mathcal{T}_n$.

Proposition 3.9. *The map $\pi_1(V_1^1) \to \pi_1(\mathcal{C}(\rho))$ induced by the inclusion of spaces is a surjection.*

Proof. First, each of $\mathcal{C}_1$ and $\mathcal{C}_2$ is a single point. In rank 3, $V_1^1 = V_1$. Thus, in all of these cases, the proposition is vacuously true. For the remainder, assume $n \geq 4$.

As per Proposition 3.8, V_1 can be thought of as the 1-skeleton of the cell complex $\mathcal{C}(\rho)$. This subcomplex has 1 vertex (the rose ρ) and a 1-cell for each combinatorial type of labelled graph with 2 vertices. We now need to show that any such *standard loop* α in V_1 can be written in $\pi_1(\mathcal{C}(\rho))$ as a product of loops in V_1^1.

At an interior point of a standard loop α, there is a central circle with two vertices, and the two vertices have valence, say, $p = p(\alpha)$ and $q = q(\alpha)$. By definition of V_1, we have that p and q are even and at least 4; say $p \leq q$. We thus have a filtration of V_1: α is in V_1^k if $(p-2)/2 \leq k$. The number $(p-2)/2$ is the number of loops glued to that vertex, other than the central circle.

Now, suppose that α is a standard loop of V_1^k for some $k \geq 2$. At any interior point of α, we perform a blowup so that we end up with a graph in $V_2 - V_1$ with a "central circle", as in the left hand side of Figure 4 (the other possibility is shown in the right hand side). Moreover, we choose the blowup so that (at least) one of the vertices has valence 4.

In the cell structure of Proposition 3.8, this blowup corresponds to a triangle. One edge of the triangle is the original standard loop α, one is a loop in V_1^1, and one is a loop in V_1^{k-1}. Therefore, in $\pi_1(\mathcal{C}(\rho))$, α is a product of loops in V_1^{k-1}. By induction, α can be written as a product of loops in V_1^1. □

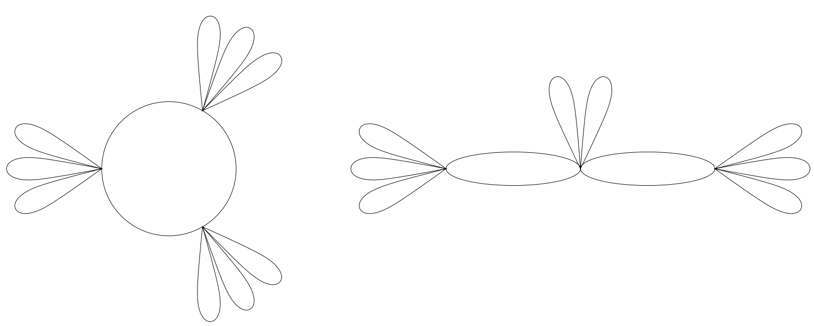

FIGURE 4. Two types of graphs in V_2.

Remark. We mention that the entire space $\mathcal{C}_n$ can be thought of as a union of $(n-2)$-tori, and the intersection between any two of these tori is a lower dimensional torus which is a product of subtori. For example, the triangle in the previous proof is part of a torus, obtained by fixing one vertex on the central circle and letting the other two vertices "move" around independently on the central circle.

4. Cohomological Dimension

We now give the argument for the first part of the main theorem, that $\mathcal{Y}_n$ is homotopically $(2n-4)$-dimensional. The basic strategy is to put an ordering on the stars of roses of $\mathcal{Y}_n$ (we think of the ordering as a Morse function) and then to glue the stars of roses together in the prescribed order. This is in the same spirit as the proof of Culler–Vogtmann that $\mathcal{X}_n$ is contractible.

4.1. **Morse function.** The ordering on roses will come from an ordering on matrices. We start with vectors. By the *norm* of an element $v = (a_1, \cdots, a_n)$ of $\mathbb{Z}^n$, we mean

$$|v| = (|a_1|, \cdots, |a_n|) \in \mathbb{Z}^n_+$$

where the elements of $\mathbb{Z}^n_+$ are ordered lexicographically. Consider the matrix

$$M = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix}$$

The *norm* of M is

$$|M| = (|v_n|, \ldots, |v_1|) \in (\mathbb{Z}^n_+)^n$$

where $(\mathbb{Z}^n_+)^n$ has the lexicographic ordering on the n factors. We say that M is a *standard representative* for an element of $W \backslash \mathrm{GL}_n(\mathbb{Z})$ if $|v_n| < \cdots < |v_1|$ (i.e. if it is a representative with smallest norm). Note that two rows of a matrix in $\mathrm{GL}_n(\mathbb{Z})$ cannot have the same norm, for otherwise these two rows would be equal after reducing modulo 2, and the resulting matrix would not be invertible.

We declare the norm of an element of $W \backslash \mathrm{GL}_n(\mathbb{Z})$ to be the norm of a standard representative, and the norm of a rose in $\mathcal{Y}_n$ to be the norm of the corresponding element of $W \backslash \mathrm{GL}_n(\mathbb{Z})$.

In what follows, the next fact will be important.

Lemma 4.1. *If the stars of two distinct roses intersect, then the roses have different norms.*

Proof. If M and M' are marking matrices for neighboring roses, then $M' = NM$, where each entry of N is either -1, 0, or $+1$ (apply Proposition 3.1(4)). Then, if $|M| = |M'|$, it follows that N is the identity modulo 2, and so $N \in W$. □

4.2. **The induction.** By arbitrarily breaking ties, the norm on roses turns the set of roses into a well-ordered set. Let A be the resulting index set, and define an *initial segment* of $\mathcal{Y}_n$ to be a union of stars of a set of roses which is closed under taking smaller roses (i.e. a sublevel set of the "Morse function" given on stars of roses). We denote an initial segment of $\mathcal{Y}_n$ by $\mathcal{Y}_n^\alpha$, where $\alpha \in A$; this is the union of roses with index less than α. Note that, in general, an initial segment consists of infinitely many roses. Our argument for the dimension of $\mathcal{Y}_n$ will be a transfinite induction on initial segments.

To this end, we define the *descending link* of a rose in $\mathcal{Y}_n$ to be the intersection of its star with the union of all stars of roses of strictly smaller norm (by Lemma 4.1, we need not worry about roses of equal norm). The descending link of a rose ρ, denoted $\mathrm{Lk}_<(\rho)$, is a subset of the frontier of its star. We will prove the following in Section 6.

Proposition 4.2. *For $n \geq 3$, descending links are homotopically $(2n-5)$-dimensional.*

Given this, we can prove the first part of the main theorem.

Theorem 4.3. *For $n \geq 3$, the complex $\mathcal{Y}_n$ is homotopy equivalent to a complex of dimension at most $2n-4$.*

Proof. The inductive statement, call it $P(\alpha)$, is the following: there exist $(2n-4)$-dimensional complexes $\{Z_n^\beta\}_{\beta<\alpha}$ and homotopy equivalences $\{h_\beta : \mathcal{Y}_n^\beta \to Z_n^\beta\}_{\beta<\alpha}$ with the properties that when $\gamma < \beta$, we have $Z_n^\gamma \subset Z_n^\beta$ and $h_\beta|_{\mathcal{Y}_n^\gamma} = h_\gamma$. The base step is Corollary 3.7. We need to show that if $P(\beta)$ is true for all $\beta < \alpha$, then $P(\alpha)$ is true.

If α is a limiting ordinal, there is nothing to do. So suppose that α has an immediate predecessor β. This means that $\mathcal{Y}_n^\alpha$ is obtained from $\mathcal{Y}_n^\beta$ by gluing the star of a rose $\mathrm{St}(\rho)$ to $\mathcal{Y}_n^\beta$ along its descending link. We can think of this operation as a diagram of spaces:

$$\mathcal{Y}_n^\beta \leftarrow \mathrm{Lk}_<(\rho) \rightarrow \mathrm{St}(\rho)$$

By the inductive hypothesis, $\mathcal{Y}_n^\beta$ is homotopy equivalent to a $(2n-4)$-dimensional complex Z_n^β. Denote by $\mathrm{St}(\rho)'$ the $(n-2)$-complex homotopy equivalent to $\mathrm{St}(\rho)$ given by Proposition 3.7. By Proposition 4.2, the descending link $\mathrm{Lk}_<(\rho)$ is homotopy equivalent to a $(2n-5)$-dimensional space $\mathrm{Lk}_<(\rho)'$. We choose maps $\mathrm{Lk}_<(\rho)' \to Z_n^\beta$ and $\mathrm{Lk}_<(\rho)' \to \mathrm{St}(\rho)'$ so that the following diagram commutes up to homotopy.

$$\begin{array}{ccc}
\mathrm{St}(\rho) & \longleftrightarrow & \mathrm{St}(\rho)' \\
\uparrow & & \uparrow \\
\mathrm{Lk}_<(\rho) & \longleftrightarrow & \mathrm{Lk}_<(\rho)' \\
\downarrow & & \downarrow \\
\mathcal{Y}_n^\beta & \longleftrightarrow & Z_n^\beta
\end{array}$$

It follows that the colimit of the diagram of spaces in the left column is homotopy equivalent to the colimit of the diagram of spaces in the right column (see e.g. [8, Proposition 4G.1]). The former, call it Z_n^α, is $(2n-4)$-dimensional (consider the double mapping cylinder), and the latter is $\mathcal{Y}_n^\alpha$. By construction, the homotopy equivalence $\mathcal{Y}_n^\alpha \to Z_n^\alpha$ extends the previous homotopy equivalences $\mathcal{Y}_n^\gamma \to Z_n^\gamma$, where $\gamma < \alpha$, as per the inductive hypothesis. It follows that the induced map $h : \mathcal{Y}_n \to Z_n$, where Z_n is the direct limit of the Z_n^α, is a homotopy equivalence (see, e.g., the discussion following [8, Proposition 4G.1]). Since Z_n has dimension at most $2n - 4$ (by construction), we are done. □

Remark. If one wants to avoid transfinite induction, it is possible to alter the Morse function so that it is the same locally (i.e. Proposition 4.2 and its proof do not change) but the image of the Morse function is order isomorphic to the positive integers.

Remark. In order to illustrate a subtlety of the last proof, we recall the inductive construction of one Eilenberg–MacLane space for $\mathbb{Q}$. For each $m \geq 1$, we take a circle c_m, which we think of as the element $1/m! \in \mathbb{Q}$. For $m \geq 1$, we cross this circle with $[0, 1]$ and for $m \geq 2$ we glue $c_m \times \{0\}$ to $c_{m-1} \times \{1\}$ with a degree m map. The result is the infinite mapping telescope of the sequence

$$\mathbb{Z} \to \mathbb{Z} \to \mathbb{Z} \to \cdots$$

where the bonding maps are multiplication by $2, 3, 4, \ldots$; the direct limit is $\mathbb{Q}$.

At each finite stage, the complex collapses to the last circle—it is homotopically 1-dimensional. However, the full complex is homotopically 2-dimensional. In the proof of Theorem 4.3, the fact that the homotopy models Z_n^α (and homotopy equivalences) are filtered along with the initial segments is what allows us to make the conclusion that $\mathcal{Y}_n$ is homotopically $(2n - 4)$-dimensional.

5. Finite generation

In this section, we recall the definition of the Magnus generating set for $\mathcal{T}_n$, and explain how our point of view recovers the result that these elements do indeed generate $\mathcal{T}_n$ (Theorem 5.6 below).

Throughout the section (and Appendix A), we denote an element ϕ of $\mathrm{Out}(F_n)$ by

$$[\Phi(x_1), \ldots, \Phi(x_n)]$$

where $x_1, \ldots, x_n$ are the generators of F_n, and Φ is a representative automorphism for ϕ.

5.1. Magnus generators.

Magnus proved that $\mathcal{T}_n$ is generated by:

$$\begin{aligned} K_{ik} &= [x_1, \ldots, x_k x_i x_k^{-1}, \ldots, x_n] \\ K_{ikl} &= [x_1, \ldots, x_i [x_k, x_l], \ldots, x_n] \end{aligned}$$

for distinct i, k, and l.

We can see the K_{ik} as loops in the star of a rose in $\mathcal{Y}_n$. Consider the picture in Figure 1. As mentioned in Section 2.2, shrinking either of the parallel edges gives a path leading to the rose with the identity marking, and so this is a loop in the star of that rose in $\mathcal{Y}_n$. By considering what is happening on the level of homotopy (as opposed to homology), we see that this loop is exactly K_{23} (see [6]). By attaching more loops at one of the vertices, and renaming the edges, we see that we can obtain any K_{ik} in the star of the identity rose. In the stars of other roses, the analogously defined loops are conjugates of the K_{ik}. What is more, we have the following.

Proposition 5.1. *The fundamental group of the star of the rose with the identity marking is generated by the* K_{ik}.

The proposition follows immediately from the fact that the loops in the above discussion corresponding to the K_{ik} are exactly the standard generators for $\pi_1(V_1)$ from Proposition 3.9.

5.2. **Proof of finite generation.** Our proof that the Magnus generators generate $\pi_1(Y_n) \cong \mathcal{T}_n$ rests on the following two topological facts about descending links which we prove in Section 6.

Proposition 5.2. *Descending links are nonempty, except for that of the rose with the identity marking.*

Proposition 5.3. *Descending links are connected.*

Combining Propositions 5.1, 5.2, and 5.3 with Van Kampen's theorem and the transitivity of the action of $\mathrm{Out}(F_n)$ on stars of roses, we see that the fundamental group of any initial segment of $\mathcal{Y}_n$ is normally generated by the K_{ik}. By transfinite induction, we have the following.

Proposition 5.4. $\mathcal{T}_n$ *is normally generated by the* K_{ik}.

The group generated by the K_{ik} is not normal in $\mathrm{Out}(F_n)$, as any element of this subgroup is of the form

$$[g_1x_1g_1^{-1}, g_2x_2g_2^{-1}, \ldots, g_nx_ng_n^{-1}]$$

Thus, to find a generating set for $\mathcal{T}_n$, we need to add more elements. We have the following result of Magnus.

Proposition 5.5. *For any n, the group generated by*

$$\{K_{ik}, K_{ikl} : i \neq k < l \neq i\}$$

is normal in $\mathrm{Out}(F_n)$.

It is now easy to prove the following, which is the third part of our main theorem.

Theorem 5.6. $\mathcal{T}_n$ *is finitely generated. In particular, it is generated by* $\{K_{ik}, K_{ikl}\}$.

Proposition 5.5 is also one of the steps in Magnus's proof that the K_{ik} and K_{ikl} generate $\mathcal{K}_n$ [11]. For completeness, we give Magnus's proof of Proposition 5.5 in Appendix A.

5.3. **Proof that $\mathcal{T}_2$ is trivial.** Since there are two ways to blow up a rank 2 rose, it follows that the star of a rose in $\mathcal{Y}_2$ is homeomorphic to an interval and that the frontier is homeomorphic to S^0. If we glue the stars of roses together inductively according to our Morse function as in Section 4, then at each stage we are gluing a contractible space (the star of the new rose) to a contractible space (the previous initial segment is contractible by induction) along a contractible space (Propositions 5.2 and 5.3 and the fact that the frontier is S^0). It follows that each initial segment, and hence all of $\mathcal{Y}_2$, is contractible; hence, $\mathcal{T}_2 = 1$.

It is more illuminating to draw a diagram of $\mathcal{X}_2 = \mathcal{Y}_2$. It is a tree, with (subdivided) 1-cells representing stars of roses. This tree is naturally dual to the classical Farey graph, with the matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ corresponding to the unordered pair $\{\frac{b}{a}, \frac{d}{c}\}$. See Figure 5.

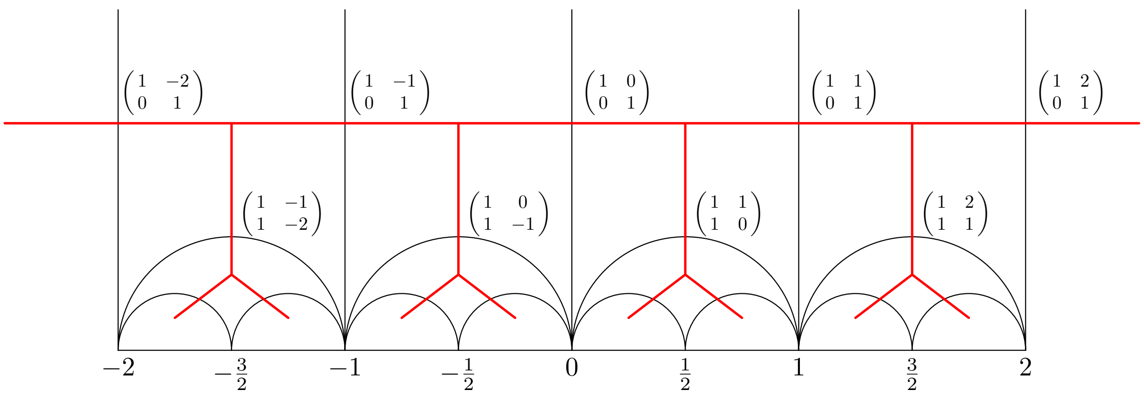

FIGURE 5. A part of the Farey graph and the dual tree $\mathcal{Y}_2$.

6. Descending links

Recall that the descending link $\mathrm{Lk}_<(\rho)$ of a rose ρ is the intersection of its star with the union of stars of roses of strictly smaller norm. The goal of this section is to prove Propositions 5.2, 5.3, and 4.2, that descending links are nonempty, connected, and homotopically $(2n-5)$-dimensional.

As in Section 3, let ρ be a rose represented by a marked graph (Γ, g) whose edges a_i are labelled v_i. We assume the a_i are ordered so that the marking matrix

$$M = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix}$$

is a standard representative. Unless specified otherwise, we make use of the cell structure for $\mathcal{Y}_n$ inherited from the simplicial structure of $\mathcal{X}_n$.

6.1. **Descending ideal edges.** An ideal edge for ρ is called *descending* if one of the corresponding 1-edge blowups (Lemma 3.4) lies in $\mathrm{Lk}_<(\rho)$. Every edge of a marked graph in $\mathrm{St}(\rho)$ which is not labelled $\pm v_i$ corresponds to some ideal edge; if the corresponding ideal edge is descending, we may say that the edge is descending.

We now give a criterion for checking whether or not a particular ideal edge is descending.

Lemma 6.1. *Let $\iota = a_{i_1} + \cdots + a_{i_m}$ be an ideal edge with $i_1 = \min\{i_j\}$. The following are equivalent:*

1. *ι is descending*
2. *all of the corresponding 1-edge blowups lie in $\mathrm{Lk}_<(\rho)$*
3. $|v_{i_1} + \cdots + v_{i_m}| < |v_{i_1}|$

Similarly, $\bar{\iota} = -a_{i_1} + a_{i_2} + \cdots + a_{i_m}$ is descending if and only if $|v_{i_1} - (v_{i_2} + \cdots + v_{i_m})| < |v_{i_1}|$.

Proof. Any 1-edge blowup which realizes the ideal edge ι lies in $m+1$ stars of roses (Proposition 3.2). Namely, for each of the v_{i_j}, we get a new marking matrix by replacing that vector with

$$v_{i_1} + \cdots + v_{i_m}$$

and leaving all other row vectors the same. To see if ι is descending, we look at the smallest of these matrices. We claim that the smallest is the matrix we get from M by replacing v_{i_1} with the above sum. Indeed, suppose we had replaced some other row vector, say v_{i_j}, with $v_{i_1} + \cdots + v_{i_m}$, obtaining a matrix N'. Now, forgetting the order of the rows, N and N' share $n-1$ rows, and N has the row vector v_{i_j} whereas N' has the row vector v_{i_1}. By the assumption that M is a standard representative, we have $|v_{i_j}| < |v_{i_1}|$. Now, if we put $|N'|$ in standard form, it is easy to find a representative for the N-coset with smaller norm than the standard representative for N'—simply replace the row of N' consisting of v_{i_1} with the vector v_{i_j}. The norm of N is less than or equal to the norm of this representative, so the claim is proven.

Now both directions are easy: if $|v_{i_1} + \cdots + v_{i_m}| < |v_{i_1}|$ then $|N|$ is obviously strictly less than $|M|$ (the given representative has smaller norm) and so ι is descending; conversely, if $|v_{i_1} + \cdots + v_{i_m}| \geq |v_{i_1}|$, then the standard representative for N either has $v_{i_1 - 1}$ or $v_{i_1} + \cdots + v_{i_m}$ in the i_1 row and so the W-coset of N has norm at least $|M|$. (We remark that the last inequality must be strict by Lemma 4.1.)

The second statement follows by symmetry. □

Corollary 6.2. *A marked graph in $\mathrm{St}(\rho)$ is in $\mathrm{Lk}_<(\rho)$ if and only if it realizes a descending ideal edge.*

As a consequence of Lemma 6.1, we see that there exist pairs of marked graphs which can never be simultaneously descending.

Lemma 6.3. *Let $\iota = a_{i_1} + \cdots + a_{i_m}$ be an ideal edge with $i_1 = \min\{i_j\}$. If ι is descending then the opposite ideal edge $\bar{\iota} = -a_{i_1} + a_{i_2} + \cdots + a_{i_m}$ is not descending.*

Proof. To simplify notation, let $w_j = v_{i_j}$ for all j. We will denote particular entries in each of these row vectors by using double indices; i.e., w_{jk} is the k^{th} entry of the row vector w_j.

Let k be the smallest number so that

$$w_{2k} + \cdots + w_{mk} \neq 0$$

Note that there is such a k, for otherwise, the original matrix M would not be invertible.

We cannot have both

$$|w_{1k} + w_{2k} + w_{3k} + \cdots + w_{mk}| < |w_{1k}|$$

and

$$|w_{1k} - (w_{2k} + w_{3k} + \cdots + w_{mk})| < |w_{1k}|$$

(indeed, the first inequality implies that w_{1k} is nonzero and $w_{2k} + \cdots + w_{mk}$ has opposite sign from w_{1k}, whereas the second inequality implies they have the same sign). An application of Lemma 6.1 completes the proof. □

6.2. **Proof of Propositions 5.2 and 5.3.** As usual, let ρ be a rose represented by a marked graph (Γ, g) with edges $a_1, \ldots, a_n$ labelled by $v_1, \ldots, v_n$, and assume that the edges are ordered so that the marking matrix

$$M = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix}$$

is a standard representative.

We first give the proof that descending links are nonempty.

Proof of Proposition 5.2. Let k be the first column of M which is not a coordinate vector (since M is a standard representative, it follows that the entries in the first $k-1$ column vectors agree with the identity matrix up to sign). If we denote the j^{th} entry of v_i by v_{ij}, then v_{kk} is nonzero. This follows from the fact that M is a standard representative and the fact that M is invertible.

Since the k^{th} column is not a coordinate vector (and since M is invertible), there is a j, different from k, so that v_{jk} is nonzero. If there is a $j > k$ such that $v_{jk} \neq 0$, then, since M is a standard representative, $|v_{jk}| \leq |v_{kk}|$, and $a_k + \epsilon_j a_j$ is a descending ideal edge for some $\epsilon_j = \pm 1$. If $v_{jk} = 0$ for all $j > k$, it follows that $v_{kk} = \pm 1$ (since M is invertible) and there is some $j < k$ so that $v_{jk} \neq 0$. But then, again, $a_j + \epsilon a_k$ is descending for some $\epsilon = \pm 1$. □

Here is the proof that descending links are connected.

Proof of Proposition 5.3. We first claim that if ι is any descending ideal edge, then there is a subordinate 2-letter ideal edge ι' which is also descending; see Section 3.2 for definitions. It will then follow from Lemma 3.5 and Corollary 6.2 that there is a path in $\mathrm{Lk}_<(\rho)$ between the 1-edge blowup realizing ι to the 1-edge blowup realizing ι' (the graph simultaneously realizing ι and ι' is the midpoint of the path).

To prove the claim, we need some notation. First, recall the notations ρ, a_i, v_i, and M from above. Also, say (without loss of generality) that $\iota = a_{i_1} + a_{i_2} + \cdots + a_{i_m}$, and denote v_{i_j} by w_j. Starting with the matrix with the w_i as rows, we obtain a matrix M' by deleting all columns without a nonzero entry. The ij^{th} entry of M' is denoted w_{ij}.

We proceed in two cases. If the first column of M' is not a coordinate vector, then at least two of the w_{i1} are nonzero, in particular, $w_{11} \neq 0$. Without loss of generality, say $w_{11} > 0$. Since ι is descending, there must be a k so that $w_{k1} < 0$, and since M is a standard representative, we have $|w_{k1}| \leq |w_{11}|$. It follows that $a_{i_1} + a_{i_k}$ is descending, and this completes the proof of the first case.

If the first column of M' is a coordinate vector (i.e. $w_{11} = \pm 1$ and $w_{k1} = 0$ for $k > 1$), then we look at the second column of M'. Without loss of generality, assume $w_{22} > 0$. At this point there are three subcases. If $w_{k2} = 0$ for all $k > 2$, then $a_{i_1} + a_{i_2}$ is descending, since ι is descending. If there is a $k > 2$ so that $w_{k2} < 0$ then $a_{i_2} + a_{i_k}$ is descending (since M is a standard representative). If $w_{k2} \geq 0$ for all $k > 2$ and $w_{k2} \neq 0$ for at least one $k > 2$, then, since ι is descending, it follows that $w_{12} < 0$ and so $a_{i_1} + a_{i_k}$ is descending for any $k > 2$ with $w_k > 0$.

We now claim that given any two descending 2-letter ideal edges, there is a path between the corresponding points in $\mathrm{Lk}_<(\rho)$. This follows, as above, from Lemma 3.5 and Corollary 6.2, in addition to the fact that opposite 2-letter ideal edges cannot both be descending (Lemma 6.3). This completes the proof. □

6.3. **Completely descending link.** We now shift our attention to Proposition 4.2. Let ρ be a rose represented by a marked graph (Γ, g), and say that Γ has edges $a_1, \ldots, a_n$ labelled by $v_1, \ldots, v_n$.

The main argument for the proof (Section 6.4 below) is purely combinatorial, referring only to isomorphism types of labelled graphs. As things stand, however, we cannot describe $\mathrm{Lk}_<(\rho)$ in terms of combinatorial graphs without metrics. Indeed, given a marked graph in $\mathrm{Lk}_<(\rho)$, if we shrink the descending edges to have length less than 1 (while staying in the frontier by enlarging a nondescending edge), then the resulting marked graph is not in $\mathrm{Lk}_<(\rho)$ (Corollary 6.2).

To remedy this problem we perform a deformation retraction of $\mathrm{Lk}_<(\rho)$ onto the *completely descending link*, which we define to be the subset of $\mathrm{Lk}_<(\rho)$ consisting of marked graphs where each edge not labelled $\pm v_i$ is descending. The deformation retraction is achieved by simply shrinking all edges which correspond to nondescending ideal edges. Recall that these edges form a forest (Proposition 3.1(2)), so there is no obstruction. We denote the completely descending link of ρ by $\mathrm{Lk}_\ll(\rho)$.

Lemma 6.4. *For any given rose ρ, the completely descending link $\mathrm{Lk}_\ll(\rho)$ is a strong deformation retract of the descending link $\mathrm{Lk}_<(\rho)$. In particular, the two are homotopy equivalent.*

We now endow $\mathrm{Lk}_{\ll}(\rho)$ with a cell structure where a cell is given by a combinatorial type of labelled graph and the cells are parameterized by the lengths of the edges in the graph (compare with the cell structure on cactus space in Proposition 3.8). To be more precise, let (Γ', g') be a marked graph in $\mathrm{Lk}_{\ll}(\rho)$, and for each i, let k_i be the number of edges of Γ' labelled $\pm v_i$. For each i we thus get a $(k_i - 1)$-simplex by projecting

$$\{(t_1, \ldots, t_{k_i}) \in [0,1]^{k_i} : t_j = 1 \text{ for some } j\}$$

to the simplex $\Delta_i = \{\sum t_i = 1\}$. This projection is a homeomorphism (of abstract cells). For each edge not labelled $\pm v_i$, we allow its length to vary arbitrarily within $[0, 1]$, as long as at least one such edge has length 1. If k_0 is the number of such edges, then, as above, we get a $(k_0 - 1)$-simplex Δ_0. Thus, the cell corresponding to (Γ', g') has a cell structure given by the product

$$\Delta_0 \times \cdots \times \Delta_n$$

We now summarize some of the important features of this cell structure.

Proposition 6.5. *Consider a cell C of $\mathrm{Lk}_{\ll}(\rho)$ as above.*

(1) *Passing to faces of C corresponds to collapsing forests in Γ'.*
(2) *C is top-dimensional if and only if all vertices of Γ' have valence 3.*
(3) *If Γ' has v vertices, then C has dimension $v - 2$.*

Part (1) implies that the proposed structure is indeed a cell structure. In fact, we have provided a cell structure on all of $\mathrm{Lk}(\rho)$, which we call the *coarse structure* on $\mathrm{Lk}(\rho)$. The key point is that $\mathrm{Lk}_{\ll}(\rho)$ is a full subcomplex in this structure, whereas $\mathrm{Lk}_{<}(\rho)$ is not.

6.4. **Proof of Proposition 4.2.** In this section we show that the completely descending link for any rose is homotopy equivalent to a complex of dimension $2n - 5$ (Proposition 6.6). Since the completely descending link is a deformation retract of the descending link (Lemma 6.4), Proposition 4.2 follows as a corollary.

As usual, let $\rho = (\Gamma, g)$ be a rose in $\mathcal{Y}_n$, with edges $a_1, \ldots, a_n$ labelled by $v_1, \ldots, v_n$. If (Γ', g') is any marked graph in $\mathrm{St}(\rho)$, we define the *v_i-loop* as the union of edges with label containing v_i; this is a topological circle (see Proposition 3.1(3)).

Proposition 6.6. *Let $n \geq 3$. For any rose ρ in $\mathcal{Y}_n$, there is a strong deformation retraction of $\mathrm{Lk}_{\ll}(\rho)$ onto a complex of dimension $2n - 5$.*

Proof. If any top-dimensional cell of $\mathrm{Lk}_{\ll}(\rho)$ has a free face in $\mathrm{Lk}_{\ll}(\rho)$, then there is a homotopy equivalence (deformation retraction) of $\mathrm{Lk}_{\ll}(\rho)$ which collapses away this cell. We perform this process inductively until we arrive at a subcomplex L where no top-dimensional cell has a free face.

We now suppose that L is $(2n - 4)$-dimensional, i.e., it has at least one top-dimensional cell. Among these, choose a cell C where the total number of edges ℓ of a v_1-loop is minimal. Call the loop P and choose one of its edges labelled $\pm v_1$ and call it e; see the leftmost diagram in Figure 6. Say that C is given by a marked graph (Γ', g').

Firstly, note that ℓ is not 1, since there are no graphs with separating edges in $\mathcal{Y}_n$.

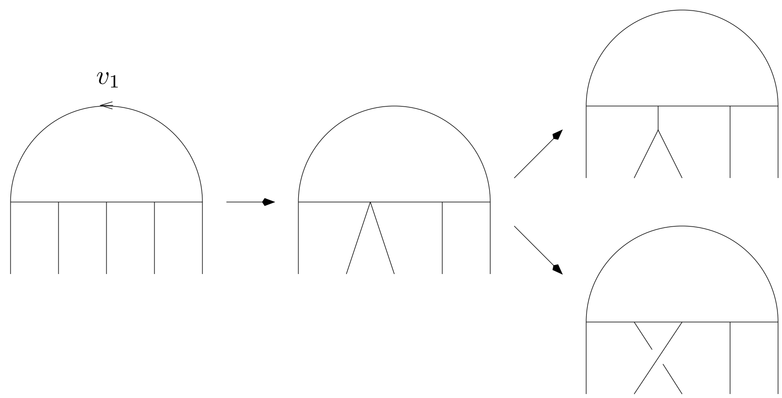

FIGURE 6. The top edge is e. The horizontal path plus e is P.

If we collapse any edge of $P - e$ (middle of Figure 6), we move to a codimension 1 face of C. There are two ways to move to a new top-dimensional cell, since there are two other blowups of the resulting valence 4 vertex. One way reduces the length of P (top right of Figure 6), so by the minimality assumption for C, this is not a cell of L. Since we are assuming C does not have any free faces, the other top-dimensional cell (bottom right of Figure 6), call it C', must be in L. The marked graphs in C and C' have the same labels outside of P; the difference is that the order of the edges leaving P has changed (Proposition 2.2(1) is applied twice).

Continuing in this way, we see that if we permute the edges leaving P in any way, we arrive at cells which are necessarily part of L. In particular, the graph obtained by taking the edge which leaves P at one endpoint of e and moving it to the other endpoint of e gives a descending cell $\bar{C}$.

Our goal now is to apply Lemma 6.3 to argue that C and $\bar{C}$ cannot both be cells of $\mathrm{Lk}_{\ll}(\rho)$. Consider either endpoint of e in Γ'. This is a valence 3 vertex, as shown in Figure 7. By Proposition 2.2(3) and Proposition 3.1(4), the labels must be as in the left hand side of the figure. When we move the edge labelled $\sum k_i v_i$ to the other end of e (as above), the labels must be as shown in the right hand side of the figure; the key point is that the labels and orientations do not change for e and the edge being moved. It is then possible for us to determine the label for the third edge leaving the vertex where these edges meet. By Lemma 6.3 and Corollary 6.2, we have a contradiction. $\square$

7. The toy model and infinite generation of top homology

In this section we will prove the second part of the main theorem, that $H_{2n-4}(\mathcal{T}_n, \mathbb{Z})$ is not finitely generated when $n \geq 3$. In order to do this, we define a subcomplex $\mathcal{M}_n$ of $\mathcal{Y}_n$, called the "toy model", we find an explicit list of infinitely many linearly independent elements of $H_{2n-4}(\mathcal{M}_n, \mathbb{Z})$, and then we show that these elements remain linearly independent in $H_{2n-4}(\mathcal{Y}_n, \mathbb{Z})$ (Theorem 7.8).

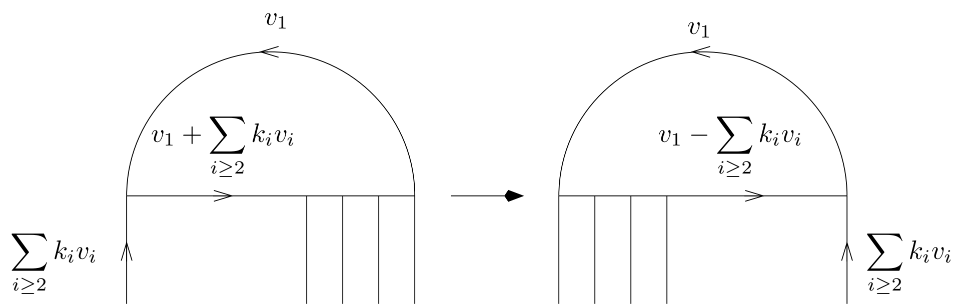

FIGURE 7. Labels at endpoints of e.

7.1. **Description of the toy model.** Let $\rho = (\Gamma, g)$ be the rose in $\mathcal{Y}_n$ with the identity marking, let x_i denote the edges of the standard rose R_n, and let a_i denote the corresponding edges of Γ.

Consider the set of points $\mathcal{M}_n^0 = \{(\Gamma', g')\}$ in $\mathrm{St}(\rho)$ where $g'(x_1) \cup g'(x_2)$ is a rank 2 rose. We define the *toy model* to be the subset $\mathcal{M}_n$ of $\mathcal{Y}_n$ given by

$$\mathcal{M}_n = \bigcup_{p_i, q_i \in \mathbb{Z}} \mathcal{M}_n^0 \cdot \begin{bmatrix} 1 & 0 & p_3 & p_4 & \cdots & p_n \\ 0 & 1 & q_3 & q_4 & \cdots & q_n \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & & & & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \end{bmatrix}$$

Another point of view. We now give a different description of $\mathcal{M}_n$, which will make it easier to find its homotopy type.

For each marked graph (Γ', g') of $\mathcal{M}_n$, the union $g'(x_1) \cup g'(x_2)$ is a rank 2 rose in Γ', and Γ' has $n-2$ edges $a_3, \ldots, a_n$ labelled $v_3, \ldots, v_n$, where v_i is a coordinate vector with $+1$ in the i^{th} entry. By considering the starting and ending points of $a_3, \ldots, a_n$ as points in $g'(x_1) \cup g'(x_2)$, a path in $\mathcal{M}_n$ can be thought of as a path in the configuration space of $n-2$ pairs of points in the universal abelian cover of $g'(x_1) \cup g'(x_2)$, which is $U = (\mathbb{R} \times \mathbb{Z}) \cup (\mathbb{Z} \times \mathbb{R})$. To make this precise, for each metric graph (Γ', g') in $\mathcal{M}_n$, we rescale the metric so that $g'(x_1)$ and $g'(x_2)$ both have length 1. After doing this, the endpoints of the a_i give a well-defined subset of the metric cover U.

If, in the configuration space, we move the two points corresponding to the endpoints of some a_i by the same integral vector, then the corresponding point in $\mathcal{M}_n$ does not change.

Proposition 7.1. *The above construction defines a homeomorphism:*

$$(U^2)^{n-2}/(\mathbb{Z}^2)^{n-2} \to \mathcal{M}_n$$

At this point, the proof is straightforward and is left to the reader.

A typical graph in $\mathcal{M}_7$ is shown in Figure 8. That graph is "maximally blown up" in the sense that it has the greatest number of valence 3 vertices possible in $\mathcal{M}_7$.

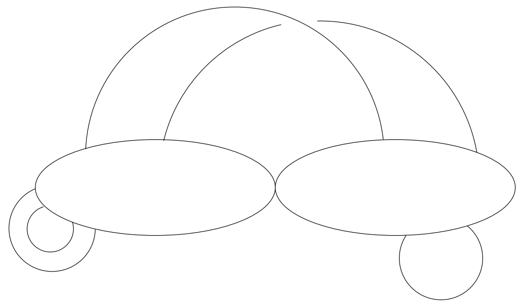

FIGURE 8. A maximally blown up graph in $\mathcal{M}_7$.

7.2. **Homotopy type of the toy model.** We start by focusing our attention on the rank 3 toy model $\mathcal{M}_3$. In general, we have $\mathcal{M}_n \cong (\mathcal{M}_3)^{n-2}$, and so we will be able to deduce the finiteness properties of $\mathcal{M}_n$ from those of $\mathcal{M}_3$.

Via Proposition 7.1, we can think of $\mathcal{M}_3$ as pairs of points in U. However, it will simplify our analysis if we thicken U to a space V, which we now define. First, let D be the closed region inside the ellipse in $\mathbb{R}^2$ given by $x^2/r^2 + y^2/(r+\epsilon)^2 = 1$ for some small r and some ϵ small with respect to r. Then, for any integers p and q, denote by $D_{p,q}$ the translate of D by the vector $(p+1/2, q+1/2)$. Finally, define V to be the complement in $\mathbb{R}^2$ of the interiors of the $D_{p,q}$. The straight line retraction of V onto U gives a homotopy equivalence from $V^2/\mathbb{Z}^2$ to $U^2/\mathbb{Z}^2 \cong \mathcal{M}_3$.

We will understand the homotopy type of $V^2/\mathbb{Z}^2$ via Morse theory for manifolds with corners, the details of which are given in Appendix B. Our "Morse function" on $V^2/\mathbb{Z}^2$, denoted d, is the Euclidean distance between the corresponding points in $\mathbb{R}^2$. The minset of d, which is the diagonal of $V^2/\mathbb{Z}^2$, is homeomorphic to a torus with one boundary component; this is the only place where d fails to be a true Morse function (it is not smooth). However, since there are no critical points with d-value very close to 0, we can apply the proof of Theorem B.1 to see that small sublevel sets are homotopy equivalent to the minset.

The critical points for d (points where the derivative does not surject onto $\mathbb{R}$) are classified as follows; c.f. Figure 9. The indices are exactly the indices of the critical points in the classical sense, when d is restricted to the boundary tori $\partial D_{0,0} \times \partial D_{p,q}$; we denote these tori by $Z_{p,q}$.

(1) There are two critical points of index 1 corresponding to the two orderings of the two points of $\partial D_{0,0}$ lying on the minor (shorter) axis of $D_{0,0}$.
(2) There are two critical points of index 2 corresponding to the two orderings of the two points of $\partial D_{0,0}$ lying on the major axis of $D_{0,0}$.
(3) For every $(p,q) \neq (0,0)$, there is an index 2 critical point corresponding to the two points of tangency of $\partial D_{0,0}$ and $\partial D_{p,q}$ with the smallest circle which is tangent to both and contains both in its interior.

(4) For every $(p,q) \neq (0,0)$, there is a critical point corresponding to the two points of tangency of $\partial D_{0,0}$ and $\partial D_{p,q}$ with the smallest circle which is tangent to both and contains both in its exterior.

The fourth type of critical point is not a true critical point, in the sense that it does not contribute to the homotopy type of $V^2/\mathbb{Z}^2$; the reason is that the derivative of d surjects onto $\mathbb{R}_{\leq 0}$ at these points.

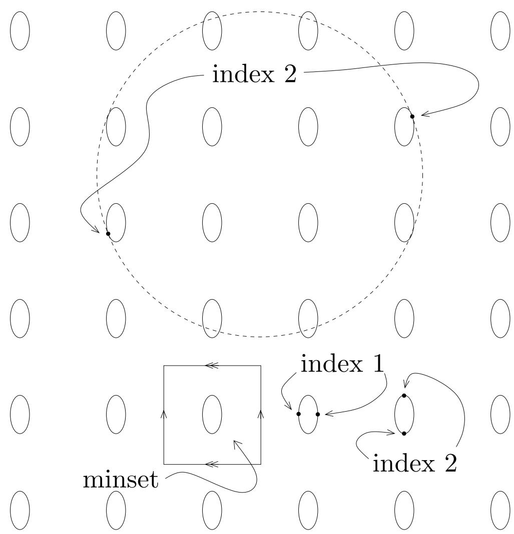

FIGURE 9. Critical points for the toy model in rank 3. A pair of points in the complement of the disks in $\mathbb{R}^2$ corresponds to a point in an abelian cover of the toy model.

In order to apply Morse theory, one needs to check that the critical points are nondegenerate, that is, the Hessian of the restriction of d to $Z_{p,q}$ is nonsingular. We illustrate the argument for the third type of critical point, which corresponds to the maximum of a $Z_{p,q}$. Let C be the smallest circle in $\mathbb{R}^2$ which contains $D_{0,0}$ and $D_{p,q}$; the critical point corresponds to the two points of tangency of C with the ellipses. We want to show that if we move along any tangent vector along $Z_{p,q}$, then d decreases to the second order. By comparing the curvatures of C and the ellipses, we see that the sum of the distances to the center of C decreases to the second order, and hence so does d.

As in classical Morse theory, Appendix B now tells us that $\mathcal{M}_3$ is homotopy equivalent to a cell complex with one 0-cell and four 1-cells (for the minset and first type of critical point), and infinitely many 2-cells (the second and third types of critical points). Since there are no 3-cells in this decomposition, $H_2(\mathcal{M}_3, \mathbb{Z})$ is infinitely generated.

Proposition 7.2. *The $Z_{p,q}$, for $p, q \in \mathbb{Z}$, freely generate a subgroup of $H_2(\mathcal{M}_3, \mathbb{Z})$.*

As mentioned, we have $\mathcal{M}_n \cong (\mathcal{M}_3)^{n-2}$, and so we obtain the desired corollary for $\mathcal{M}_n$. In $\mathcal{M}_n$, we define $Z_{p,q}$ to be the $(n-2)$-fold product of the $Z_{p,q}$ in $\mathcal{M}_3$.

Corollary 7.3. *The $Z_{p,q}$, for $p, q \in \mathbb{Z}$, freely generate a subgroup of $H_{2n-4}(\mathcal{M}_n, \mathbb{Z})$.*

We remark that the image of $\pi_1(Z_{p,q}) \cong \mathbb{Z}^{2n-4}$ in $\pi_1(\mathcal{Y}_n) \cong \mathcal{T}_n$ is a conjugate of the subgroup G of $\mathcal{T}_n$ described in the introduction. To see this, one simply needs to understand the effect of blowups and blowdowns on the *homotopy* classes of marked graphs; see [6].

7.3. **Independence in homology.** We now set out to prove that the $Z_{p,q}$ represent independent classes in $\mathcal{Y}_n$ (Theorem 7.8). In particular, this will prove the second part of the main theorem.

By understanding the homotopy equivalences

$$(V^2)^{n-2}/(\mathbb{Z}^2)^{n-2} \to (U^2)^{n-2}/(\mathbb{Z}^2)^{n-2} \to \mathcal{M}_n$$

we can give a concrete description of the $Z_{p,q}$ in terms of marked graphs. First of all, we have the following.

Lemma 7.4. *Each $Z_{p,q}$ is contained in the union of stars of roses with marking matrices of the form*

$$\begin{bmatrix} 1 & 0 & p_3 & p_4 & \cdots & p_n \\ 0 & 1 & q_3 & q_4 & \cdots & q_n \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & & & & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \end{bmatrix}$$

where each $p_i \in [p-1, p+1]$ and $q_i \in [q-1, q+1]$.

The main fact we will need about the stars of roses listed in Lemma 7.4 is that any ideal edge of the form

$$\pm v_1 \pm v_2 + \sum_{i \geq 3} k_i v_i$$

is *ascending*. We will also need the following observation about the $Z_{p,q}$.

Lemma 7.5. *For any $Z_{p,q}$ and any rose ρ, the intersection $Z_{p,q} \cap \mathrm{Lk}(\rho)$ is a full subcomplex of $\mathrm{Lk}(\rho)$ with the coarse structure (c.f. Proposition 6.5). In particular, if a marked graph in $Z_{p,q} \cap \mathrm{Lk}(\rho)$ has an edge of length less than 1 corresponding to an ideal edge ι, then there is a point in that $Z_{p,q}$ which realizes ι.*

For the remainder, fix a $Z_{p,q}$, and let ρ be a rose of greatest norm whose star intersects $Z_{p,q}$. By Lemma 7.4, if p and q are both nonzero, then ρ is unique; if one of them is zero, then there are 2^{n-2} choices for ρ; and if $p = q = 0$, then there are 2^{2n-4} choices. From Lemma 7.4, we deduce the following key fact.

Lemma 7.6. *Any map from $\{Z_{p,q}\}$ to roses, sending $Z_{p,q}$ to any rose ρ of maximal norm with $Z_{p,q} \cap \mathrm{St}(\rho) \neq \emptyset$, is injective. In particular, given any finite subset of $\{Z_{p,q}\}$, a rose of maximal norm whose star intersects this set has nonempty intersection with exactly one torus in this set.*

Let $\mathrm{Lk}_<(\rho)$ denote the descending link for ρ as defined in Section 4.2.

Lemma 7.7. *The intersection $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ is homeomorphic to S^{2n-5}.*

Proof. We assume $p, q > 0$, with the other cases handled similarly.

Under the homotopy equivalence $(V^2)^{n-2}/(\mathbb{Z}^2)^{n-2} \to (U^2)^{n-2}/(\mathbb{Z}^2)^{n-2}$, we can identify $Z_{p,q}$ with the configuration space of $n-2$ pairs of points in $U \subset \mathbb{R}^2$ where the first point z_i in each pair lies on the coordinate square with vertices at $(0,0)$ and $(1,1)$ and the second point z_i' in each pair lies on the square with vertices (p,q) and $(p+1, q+1)$.

The rose ρ is realized when each z_i is at $(0,0)$ and each z_i' is at the point $(p+1, q+1)$. The points of $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ are exactly the set of points where each z_i is within a distance of $1/2$ from the origin, each z_i' is within $1/2$ of $(p+1, q+1)$, and at least one z_i or z_i' has distance exactly $1/2$.

In other words, each z_i and z_i' is allowed to move within a closed interval, and such a configuration is in $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ if at least one of the points is on the boundary of its interval. Thus, $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ is homeomorphic to $\partial I^{2n-4} \cong S^{2n-5}$. □

The following completes the proof of the main theorem.

Theorem 7.8. *Let $n \geq 3$. The subgroup of $H_{2n-4}(\mathcal{M}_n, \mathbb{Z})$ generated by the $Z_{p,q}$ injects into $H_{2n-4}(\mathcal{Y}_n, \mathbb{Z})$. In particular, the $Z_{p,q}$, for $p, q \in \mathbb{Z}$, form an infinite set of independent classes in $H_{2n-4}(\mathcal{Y}_n, \mathbb{Z}) \cong H_{2n-4}(\mathcal{T}_n, \mathbb{Z})$.*

Proof. Given a finite subset A of $\{Z_{p,q}\}$, let $Z_{p,q}$ be an element which intersects the star of a rose ρ of highest norm (the marking matrix for ρ has the form described in Lemma 7.4; let v_i denote the i^{th} row). We know that there is a strong deformation retraction of $\mathrm{Lk}_<(\rho)$ onto a complex of dimension $2n-5$ (Lemma 6.4 plus Proposition 6.6). The goal is to show that the image of the sphere $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ is embedded in this complex. Even better, we will show that the deformation retractions of Lemma 6.4 and Proposition 6.6 do not move the points of $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$.

To see why this proves the theorem, we consider the long exact sequence associated to the pair $(\mathrm{St}(\rho), \mathrm{Lk}_<(\rho))$:

$$\begin{aligned} \cdots \to H_{2n-4}(\mathrm{St}(\rho)) &\to H_{2n-4}(\mathrm{St}(\rho), \mathrm{Lk}_<(\rho)) \\ &\to H_{2n-5}(\mathrm{Lk}_<(\rho)) \to H_{2n-5}(\mathrm{St}(\rho)) \to \cdots \end{aligned}$$

By excision, $Z_{p,q}$ corresponds to a class in $H_{2n-4}(\mathrm{St}(\rho), \mathrm{Lk}_<(\rho))$. The image in $H_{2n-5}(\mathrm{Lk}_<(\rho))$ is the class $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$, which is nontrivial once we show $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ is embedded in the $(2n-5)$-dimensional deformation retract of $\mathrm{Lk}_<(\rho)$ (Lemma 7.7). It follows that $Z_{p,q}$ is nontrivial in $H_{2n-4}(\mathrm{St}(\rho), \mathrm{Lk}_<(\rho))$ and hence, via excision, in $H_{2n-4}(\mathcal{Y}_n^\alpha, \mathcal{Y}_n^\beta)$ where $\mathcal{Y}_n^\beta$ is the largest initial segment (see Section 4.2) not containing ρ and $\mathcal{Y}_n^\alpha = \mathcal{Y}_n^\beta \cup \mathrm{St}(\rho)$. By Lemma 7.6, each element of A other than $Z_{p,q}$ is trivial in $H_{2n-4}(\mathcal{Y}_n^\alpha, \mathcal{Y}_n^\beta)$, and so $Z_{p,q}$ is linearly independent from these, which is what we wanted to show.

Thus, we are reduced to showing that the two deformation retractions do not move the sphere $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$. We handle each in turn.

For the deformation retraction of $\mathrm{Lk}_<(\rho)$ onto $\mathrm{Lk}_\ll(\rho)$, we need to show that $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ is already contained in $\mathrm{Lk}_\ll(\rho)$. Suppose that there were a point of $Z_{p,q} \cap \mathrm{Lk}_<(\rho)$ which were not contained in $\mathrm{Lk}_\ll(\rho)$. By Lemma 7.5, there is a point which realizes an *ascending* ideal edge (Lemma 4.1), and this implies that $Z_{p,q}$ intersects the star of some rose of higher norm, contradicting the choice of ρ.

We now focus on the deformation retraction of Proposition 6.6. A marked graph representing a maximal cell of $\mathrm{Lk}_\ll(\rho)$ must have disjoint v_1- and v_2-loops. Indeed, there are no valence 4 vertices, so the overlap would have to contain an ascending edge by Proposition 2.2(3), the statement after Lemma 7.4, and Lemma 6.1. The codimension 1 cells which get collapsed are obtained from these maximal cells by collapsing an edge of the v_1-loop. Thus, the corresponding graphs still have disjoint v_1- and v_2-loops. On the other hand, in any graph of $Z_{p,q}$, the v_1-loop and the v_2-loop intersect in exactly 1 point. Thus, no points of $Z_{p,q}$ are moved during this retraction, so we are done. □

Appendix A. Proof of Proposition 5.5

This appendix contains Magnus's proof of Proposition 5.5. In this section, we freely use the notation of Section 5.

Let K be the subgroup of $\mathrm{Out}(F_n)$ generated by the K_{ik} and K_{ikl} for distinct i, k, and l. We now prove Proposition 5.5, that K is normal in $\mathrm{Out}(F_n)$.

Proof of Proposition 5.5. We choose the following generating set for $\mathrm{Out}(F_n)$.

$$\begin{aligned} \delta_{12} &= [x_1x_2, x_2, \ldots, x_n] \\ \Omega_1 &= [x_1^{-1}, x_2, \ldots, x_n] \\ \Pi_{i-1} &= [x_1, \ldots, x_{i-2}, x_i, x_{i-1}, x_{i+1}, \ldots, x_n] \end{aligned}$$

It suffices to show that the conjugates of the K_{ik} and K_{ikl} by the chosen generators of $\mathrm{Out}(F_n)$ (and their inverses) are elements of K.

We have the following simplifications.

(1) Operations on disjoint sets of elements commute.
(2) Since Ω_1 and Π_{i-1} have order 2, we do not need to conjugate by their inverses.
(3) We do not need to conjugate by δ_{12}^{-1} since
$$(\Pi_1\Omega_1\Pi_1)\delta_{12}(\Pi_1\Omega_1\Pi_1)^{-1} = \delta_{12}^{-1}$$
(4) Since $K_{ikl} = K_{ilk}^{-1}$, we may assume $k < l$.
(5) Any outer automorphism ψ of the form
$$[x_1, \ldots, g'x_ig, \ldots, x_n]$$
where gg' is an element of the commutator subgroup of the subgroup H of F_n generated by $\{x_k : k \neq i\}$ is an element of K.

To see that $\psi \in K$, first note that, by postcomposing with a product of $K_{i\star}^{\pm 1}$, we may assume that $g' = 1$. Now, we know that the commutator subgroup of H is normally generated by the $[x_k, x_l]$, where k and l are both different from i. Therefore, it suffices to handle the case of

$$g = h[x_k, x_l]h^{-1} = [hx_kh^{-1}, hx_lh^{-1}]$$

where $h = x_{i_1} \cdots x_{i_p}$ is an arbitrary element of H. It is elementary to check that

$$\psi = P^{-1}K_{ikl}P$$

where

$$P = \prod_{j \neq i} K_{ji_p} \cdots \prod_{j \neq i} K_{ji_2} \prod_{j \neq i} K_{ji_1}$$

Given these simplifications, it is straightforward to check (case by case) that the conjugates by Π_{i-1}, Ω_1, and δ_{12} of each K_{ik} and K_{ikl} are elements of K. There is one exception; we give Magnus's computation for this difficult case here:

$$\delta_{12}K_{2l1}\delta_{12}^{-1} = K_{l2}K_{l1}^{-1}K_{1l}K_{l1}K_{2l1}K_{12l}K_{l2}^{-1}K_{2l}^{-1}$$

□

Appendix B. Morse theory for manifolds with corners

Let M be a smooth manifold with corners (see the appendix of [4] for background). For each point $p \in M$, let $T_p(M)$ be the tangent space of M at p. The subset $C_p(M) \subseteq T_p(M)$ of those tangent vectors at p not leaving M (within a chart, and this can be seen to be invariant) is called the *inward tangent cone*. The name is appropriate: the inward tangent cone is a convex cone with cone point $0 \in T_p(M)$.

We can classify points in M as corner points, edge points, face points, etc.; the stratification is most conveniently defined by a number $d(p)$, which is defined to be the maximum dimension of a linear subspace in $T_p(M)$ that lies completely within $C_p(M)$. Let M^k be $\{p \in M : d(p) = k\}$. Each of these strata is a manifold without boundary of dimension k.

Let $f : M \to \mathbb{R}$ be a smooth function. A point $p \in M$ is *critical* if $D_pf(C_p(M)) \neq \mathbb{R}$. Since the derivative D_pf is a linear form on $T_p(M)$, it will map $C_p(M)$ to a convex cone in $\mathbb{R}$. Thus, there are three types of critical points, corresponding to whether $D_pf(C_p(M))$ is $\{0\}$, $\mathbb{R}_{\geq 0}$, or $\mathbb{R}_{\leq 0}$. In either of the first two cases, we say that the critical point is *positive*.

For any real number s, we denote by M_s the sublevel set $\{p \in M : f(p) \leq s\}$. We have the following Morse lemma due to Braess [5, Satz 4.1].

Theorem B.1. *For $s < t$, if $f^{-1}([s, t])$ is compact and does not contain any positive critical points, then M_s is a deformation retract of M_t.*

Before we discuss the effects of the positive critical points, we need a little preparation. Recall that each point in an n-manifold with corners has a coordinate chart U (containing p) with coordinate functions $x_1, \ldots, x_n : U \to \mathbb{R}$ which together give a homeomorphism of the open set U onto an open subset of $\mathbb{R}^n_{\geq 0}$. We can always arrange these functions so that $x_i(p) > 0$ for $i \leq d(p)$ and $x_i(p) = 0$ for $i > d(p)$. In this case, we say that the chart is *admissible*.

Let $p \in M$ be a positive critical point and fix an admissible chart $x_1, \ldots, x_n : U \to \mathbb{R}$ around p. Since $M^{d(p)}$ is a submanifold of dimension $d(p)$ containing p, we have

$$\frac{\partial f}{\partial x_i}(p) \begin{cases} = 0 & i \leq d(p) \\ \geq 0 & i > d(p) \end{cases}$$

We say that the critical point p satisfies *condition B* if

$$\ker(D_p f) \cap C_p(M) = T_p(M^{d(p)})$$

Equivalently, $\partial f / \partial x_i > 0$ for $i > d(p)$. A critical point which satisfies condition B is *nondegenerate* if the Hessian

$$\left(\frac{\partial^2 f}{\partial x_i \partial x_j}(p) \right)_{i,j \leq d(p)}$$

is nonsingular. The *index* of a nonsingular critical point p which satisfies condition B is the index in the sense of classical Morse theory, when we restrict the Morse function f to the submanifold $M^{d(p)}$.

Finally, we have the second Morse lemma, also due to Braess [5, Satz 7.1].

Theorem B.2. *For $s < t$, assume that $f^{-1}([s,t])$ is compact and contains k positive nondegenerate critical points $p_1, \ldots, p_k$ which satisfy condition B, satisfies $s < f(p_i) < t$, and have index λ_i for $1 \leq i \leq k$. Then M_t is homotopy equivalent to M_s with one λ_i-cell attached for each i.*

Braess only states this theorem in the case $k = 1$, but his proof holds in this generality: small perturbations of the Morse function are used to push critical points away from the corners into the manifold; these perturbations are chosen individually for each critical point.

References

[1] Problems in low-dimensional topology. In Rob Kirby, editor, *Geometric topology (Athens, GA, 1993)*, volume 2 of *AMS/IP Stud. Adv. Math.*, pages 35–473. Amer. Math. Soc., Providence, RI, 1997.

[2] Toshiyuki Akita. Homological infiniteness of Torelli groups. *Topology*, 40(2):213–221, 2001.

[3] Gilbert Baumslag and Tekla Taylor. The centre of groups with one defining relator. *Math. Ann.*, 175:315–319, 1968.

[4] A. Borel and J.-P. Serre. Corners and arithmetic groups. *Comment. Math. Helv.*, 48:436–491, 1973. Avec un appendice: Arrondissement des variétés à coins, par A. Douady et L. Hérault.

[5] Dietrich Braess. Morse-Theorie für berandete Mannigfaltigkeiten. *Math. Ann.*, 208:133–148, 1974.

[6] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.

[7] John L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.*, 84(1):157–176, 1986.

[8] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[9] Dennis Johnson. The structure of the Torelli group. I. A finite set of generators for $\mathcal{I}$. *Ann. of Math. (2)*, 118(3):423–442, 1983.

[10] Sava Krstić and James McCool. The non-finite presentability of $\mathrm{IA}(F_3)$ and $\mathrm{GL}_2(\mathbf{Z}[t, t^{-1}])$. *Invent. Math.*, 129(3):595–606, 1997.

[11] W. Magnus. Über n-dimensional Gittertransformationen. *Acta Math.*, 64:353–367, 1934.
[12] Darryl McCullough and Andy Miller. The genus 2 Torelli group is not finitely generated. *Topology Appl.*, 22(1):43–49, 1986.
[13] Geoffrey Mess. Unit tangent bundle subgroups of the mapping class groups. *Preprint IHES/M/90/30*, 1990.
[14] Geoffrey Mess. The Torelli groups for genus 2 and 3 surfaces. *Topology*, 31(4):775–790, 1992.
[15] J. Nielsen. Die isomorphismengruppe der freien Gruppen. *Math. Ann.*, 91:169–209, 1924.
[16] Jean-Pierre Serre. Cohomologie des groupes discrets. In *Prospects in mathematics (Proc. Sympos., Princeton Univ., Princeton, N.J., 1970)*, pages 77–169. Ann. of Math. Studies, No. 70. Princeton Univ. Press, Princeton, N.J., 1971.
[17] John Smillie and Karen Vogtmann. Automorphisms of graphs, p-subgroups of $\mathrm{Out}(F_n)$ and the Euler characteristic of $\mathrm{Out}(F_n)$. *J. Pure Appl. Algebra*, 49(1-2):187–200, 1987.
[18] John Smillie and Karen Vogtmann. A generating function for the Euler characteristic of $\mathrm{Out}(F_n)$. In *Proceedings of the Northwestern conference on cohomology of groups (Evanston, Ill., 1985)*, volume 44, pages 329–348, 1987.
[19] Karen Vogtmann. Automorphisms of free groups and outer space. In *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000)*, volume 94, pages 1–31, 2002.

Mladen Bestvina: Department of Mathematics, University of Utah, 155 S 1440 East, Salt Lake City, UT 84112-0090

E-mail address: bestvina@math.utah.edu

Kai-Uwe Bux: Department of Mathematics, University of Virginia, Kerchof Hall 229, Charlottesville, VA 22903-4137

E-mail address: kb2ue@virginia.edu

Dan Margalit: Department of Mathematics, University of Utah, 155 S 1440 East, Salt Lake City, UT 84112-0090

E-mail address: margalit@math.utah.edu