

Rational homological stability for groups of partially symmetric automorphisms of free groups

MATTHEW C B ZAREMSKY

Let F_{n+m} be the free group of rank $n + m$, with generators x_1, \dots, x_{n+m} . An automorphism ϕ of F_{n+m} is called partially symmetric if for each $1 \leq i \leq m$, $\phi(x_i)$ is conjugate to x_j or x_j^{-1} for some $1 \leq j \leq m$. Let ΣAut_n^m be the group of partially symmetric automorphisms. We prove that for any $m \geq 0$ the inclusion $\Sigma\text{Aut}_n^m \rightarrow \Sigma\text{Aut}_{n+1}^m$ induces an isomorphism in rational homology for dimensions i satisfying $n \geq (3(i + 1) + m)/2$, with a similar statement for the groups $P\Sigma\text{Aut}_n^m$ of pure partially symmetric automorphisms. We also prove that for any $n \geq 0$ the inclusion $\Sigma\text{Aut}_n^m \rightarrow \Sigma\text{Aut}_n^{m+1}$ induces an isomorphism in rational homology for dimensions i satisfying $m > (3i - 1)/2$.

20F65; 20F28, 57M07

1 Introduction

Let $\text{Aut}(F_{n+m})$ be the group of automorphisms of the free group F_{n+m} . For a fixed basis $\{x_1, \dots, x_{n+m}\}$ of F_{n+m} , an automorphism ϕ of F_{n+m} is called *partially symmetric* if for each $1 \leq i \leq m$, $\phi(x_i)$ is conjugate to x_j or x_j^{-1} for some $1 \leq j \leq m$. If ϕ is an automorphism such that each $\phi(x_i)$ is even conjugate to x_i we call ϕ *pure partially symmetric*. Call these first m generators *distinguished* and the other n *undistinguished*. Let ΣAut_n^m be the group of partially symmetric automorphisms of F_{n+m} , and $P\Sigma\text{Aut}_n^m$ the group of pure partially symmetric automorphisms.

We prove that the rational homologies of ΣAut_n^m and $P\Sigma\text{Aut}_n^m$ are stable in the parameter n , and the rational homology of ΣAut_n^m is also stable in m . This means that the rational homology is independent of the parameters once they are large enough. This question was posed by McEwen in his thesis [15], where a general strategy was outlined, involving a hypothetical Morse function on a version of a space introduced by Bux, Charney and Vogtmann [4]. As a first step, McEwen [15] and Zaremsky [16] construct a Morse function for the spine of Auter space, which provided a simplified proof of the so-called Degree theorem of Hatcher and Vogtmann [9]. From the Degree theorem, the rational homological stability of $\text{Aut}(F_n) = \Sigma\text{Aut}_n^0$ can be deduced. With this Morse-theoretic approach in hand for the classical case, it was supposed that one

should then be able to generalize the situation to ΣAut_n^m , but this was left in the conjectural stage in [15]. In the present work we complete this project; namely, we exhibit a Morse function that yields a generalized version of the Degree theorem, from which we deduce rational homological stability for ΣAut_n^m .

To keep the notation straight, we mention that in [4] the “outer” version of the group we are calling $P\Sigma\text{Aut}_n^m$ is denoted $P\Sigma(n, k)$, where n is the total rank and k the number of distinguished generators. Jensen and Wahl [14] denote the same group by A_n^k , where n and k are the number of undistinguished and distinguished generators, respectively. They also consider certain other groups denoted $A_{n,k}$, which are central extensions of A_n^k , but these are not the same as the groups ΣAut_n^m considered here. For example, the automorphisms that properly permute the distinguished generators of F_{n+m} appear only in ΣAut_n^m , and not in $P\Sigma\text{Aut}_n^m = A_n^m$ or in $A_{n,m}$.

The relevant existing results are as follows. Hatcher and Vogtmann [9] showed that the homology of $\text{Aut}(F_n) = \Sigma\text{Aut}_n^0$ is stable with respect to n . In [7, Corollary 1.2], Galatius showed that the stable rational homology is even trivial, namely, $H_i(\text{Aut}(F_n); \mathbb{Q}) = 0$ for all $n > 2i + 1$. At the other end of the spectrum, in [12] Hatcher and Wahl showed that the group of *symmetric automorphisms* $\Sigma\text{Aut}(F_m) = \Sigma\text{Aut}_0^m$ is homologically stable in m , and it turns out the rational homology actually vanishes in every dimension by independent results of Griffin [8] and Wilson [19]. In contrast, the pure case is quite different. The rational homology of $P\Sigma\text{Aut}_0^m$ is not stable in m [14], and in fact the cohomology ring has been completely computed by Jensen, McCammond and Meier [13]. To use the notation of [14], so $P\Sigma\text{Aut}_0^m$ is denoted A_n^m , while the A_0^m are not homologically stable, the groups $A_{n,m}$ are in fact stable in n and m , even with coefficients in \mathbb{Z} , by work of Hatcher and Wahl [11]. We remark that the methods used to prove stability for $A_{n,m}$ are very different from how we will prove stability for ΣAut_n^m here.

We actually obtain stability results for a range of families of subgroups of ΣAut_n^m , which includes the groups $P\Sigma\text{Aut}_n^m$. Consider any family of groups G_n^m such that

$$P\Sigma\text{Aut}_n^m \leq G_n^m \leq \Sigma\text{Aut}_n^m$$

for each n and m , and such that the inclusion

$$\Sigma\text{Aut}_n^m \hookrightarrow \Sigma\text{Aut}_{n+1}^m,$$

given by extending $\phi \in \Sigma\text{Aut}_n^m$ to F_{m+n+1} via $\phi(x_{n+m+1}) = x_{n+m+1}$, restricts to an inclusion $G_n^m \hookrightarrow G_{n+1}^m$. Of course $P\Sigma\text{Aut}_n^m$ and ΣAut_n^m are examples of such families of groups. Our main result for these groups is the following theorem.

Theorem 1.1 (Stability in n) For any $m \geq 0$ and $i \geq 0$, and any family of groups G_n^m satisfying the above conditions, the map

$$H_i(G_n^m; \mathbb{Q}) \rightarrow H_i(G_{n+1}^m; \mathbb{Q})$$

induced by inclusion is an isomorphism for $n \geq (3(i + 1) + m)/2$.

Corollary The rational homology of ΣAut_n^m is stable in n , as is the rational homology of $P\Sigma \text{Aut}_n^m$. \square

We also consider stability in the other parameter, m . Renumber the elements of the basis by $\{x_1, \dots, x_n, x_{n+1}, \dots, x_{n+m}\}$, so an automorphism ϕ is partially symmetric if for all $1 \leq i \leq m$, $\phi(x_{n+i})$ is conjugate to x_{n+j} or x_{n+j}^{-1} for some $1 \leq j \leq m$. We now have a natural inclusion map

$$\Sigma \text{Aut}_n^m \hookrightarrow \Sigma \text{Aut}_n^{m+1},$$

given by extending $\phi \in \Sigma \text{Aut}_n^m$ to F_{n+m+1} via $\phi(x_{n+m+1}) = x_{n+m+1}$.

Theorem 1.2 (Stability in m) For any $n \geq 0$ and $i \geq 0$, the map

$$H_i(\Sigma \text{Aut}_n^m; \mathbb{Q}) \rightarrow H_i(\Sigma \text{Aut}_n^{m+1}; \mathbb{Q})$$

induced by inclusion is an isomorphism for $m > (3i - 1)/2$.

In Section 2 we provide some background on the spine of Outer space K_{n+m} , and describe a contractible subcomplex ΔK_n^m that admits a nice ΣAut_n^m action. We also filter ΔK_n^m using the notion of *weighted degree*, a generalization of *degree* from [9]. In Section 3 we define a height function h on ΔK_n^m , which generalizes the height function from the classical case, constructed by McEwen [15] and McEwen and Zaremsky [16]. We then show how the main result of Section 5, Proposition 5.14, about connectivity of descending links with respect to h , implies our so-called Generalized degree theorem, Theorem 3.5. In Section 4 we show how the Generalized degree theorem yields our homological stability results. In Section 5 we prove Proposition 5.14. This is done by separately considering two join factors, the *d-down link*, in Section 5.1, and the *d-up link*, in Section 5.2.

Acknowledgements The author is grateful to Kai-Uwe Bux for his guidance in this project, and to James Griffin, Rob McEwen and Jenny Wilson for their helpful advice along the way. This work was done while the author was supported by the SFB 701 in Bielefeld, Germany, and this support is gratefully acknowledged. Finally, many thanks are due to the anonymous referee, for a careful and thorough reading that resulted in a number of improvements to the paper.

2 Outer space and our space of interest

We will analyze the homology of ΣAut_n^m by considering its action on a certain simplicial complex. Our starting point is the well-studied *spine of Outer space* K_n introduced by Hatcher and Vogtmann in [9]. Let R_n be the rose with n edges, ie the graph with a single vertex p_0 and n edges. Here by a *graph* we mean a connected one-dimensional CW-complex, with the usual notions of vertices and edges. We identify F_n with $\pi_1(R_n)$. If Γ is a graph with basepoint vertex p , a homotopy equivalence $\rho: R_n \rightarrow \Gamma$ is called a *marking* on Γ if ρ takes p_0 to p . We will consider two markings to be equivalent if there is a basepoint-preserving homotopy between them. We will be interested in the set of equivalence classes of triples (Γ, p, ρ) . We only consider graphs such that p is at least bivalent and all other vertices are at least trivalent. Note that we do allow *separating edges*, that is edges whose complement in the graph is disconnected.

For graphs Γ_1 and Γ_2 , a basepoint-preserving homotopy equivalence $d: \Gamma_1 \rightarrow \Gamma_2$ is called a *forest collapse* or a *blow-down* if it amounts to collapsing a subforest F of Γ_1 . We write the blow-down as Γ/F . The reverse of a blow-down is, naturally, called a *blow-up*. This gives us a partial ordering on the set of equivalence classes of triples (Γ, p, ρ) , namely $(\Gamma', p, \rho') \leq (\Gamma, p, \rho)$ if there is a forest collapse $d: \Gamma \rightarrow \Gamma'$ such that ρ' is equivalent to $d \circ \rho$. The spine K_n of Outer space is then the geometric realization of the poset of equivalence classes of triples (Γ, p, ρ) with Γ a rank- n graph, with this partial ordering. In particular the vertices of K_n are equivalence classes of marked basepointed graphs.

Since we are identifying F_n with $\pi_1(R_n)$, we can also identify $\text{Aut}(F_n)$ with the group of basepoint-preserving homotopy equivalences of R_n , up to homotopy. This is the same as the group of markings of R_n , so we can denote markings on R_n by elements of $\text{Aut}(F_n)$. There is a (right) action of $\text{Aut}(F_n)$ on K_n in the following way: given a vertex (Γ, p, ρ) in K_n and $\phi \in \text{Aut}(F_n)$, we have

$$\phi(\Gamma, p, \rho) = (\Gamma, p, \rho \circ \phi).$$

This action only affects markings, and in fact $\text{Aut}(F_n)$ permutes markings arbitrarily.

A space for ΣAut_n^m

We now describe a subcomplex of K_{n+m} on which ΣAut_n^m acts nicely. First we will restrict to only allowing certain markings, using a standard technique, and then we will restrict further to only allowing certain graphs, à la Bux, Charney and Vogtmann [4].

Let (R_{n+m}, p_0, ϕ) be a marked rose in K_{n+m} , so the marking ϕ is really an element of $\text{Aut}(F_{n+m})$. Let W be the set of conjugacy classes in F_{n+m} of the distinguished

generators and their inverses, $x_1^{\pm 1}, \dots, x_m^{\pm 1}$. Note that $\text{Aut}(F_{n+m})$ acts on the set of all conjugacy classes, and ΣAut_n^m is precisely the stabilizer of W . We say that the rose (R_{n+m}, p_0, ϕ) has *minimal norm* if the quantity

$$\sum_{w \in W} |\phi(w)|$$

is minimized, where $|\phi(w)|$ is the length of a cyclically reduced representative of the conjugacy class $\phi(w)$, with respect to the generating set $\{x_i^{\pm 1}\}_{i=1}^{n+m}$. The *star* of this rose is the subcomplex of K_{n+m} spanned by vertices that are obtained by blowing up the rose. Denote by K_{n+m}^W the subcomplex of K_{n+m} that is the union of the stars of all roses with minimal norm.

Similar complexes, for other W , were crucial to the original proof of contractibility of Outer space by Culler and Vogtmann [6]. Our K_{n+m}^W was considered by McEwen in [15], and the ‘‘Outer’’ version was considered in [4, Section 3.1]. The complex K_{n+m}^W is contractible and ΣAut_n^m -invariant. If the marked basepointed graph (Γ, p, ρ) lies in K_{n+m}^W we will call the marking ρ *admissible*.

Having restricted to admissible markings, we next impose restrictions on the graphs. Let (Γ, p, ρ) be a vertex in K_{n+m}^W , so $\rho: R_{n+m} \rightarrow \Gamma$ is an admissible marking. We consider the closed paths, or *cycles*, $\rho(x_i)$ in Γ for $1 \leq i \leq m$. These cycles may not be reduced, so for each $1 \leq i \leq m$ let C_i be the reduced cycle in Γ obtained by reducing $\rho(x_i)$. Here *reducing* a path means inductively removing any subpaths consisting of an edge followed immediately by the same edge in the reverse direction, and a path is *reduced* if no such reduction is possible. Since ρ is admissible, C_i is an embedded simple cycle, that is, it is homeomorphic to a circle; this is for the same reason as in the nonbasepointed case [4, Lemma 15(1)].

Definition 2.1 (Viable graph) With the above notation, we call a graph Γ *viable* if the C_i are all pairwise disjoint.

See Figure 1 for an example. For brevity we will just define a *viable marked graph* to be a viable graph with an admissible marking. Let ΔK_n^m be the subcomplex of K_{n+m}^W consisting of viable marked graphs. Compare this to the nonbasepointed version, denoted $D_{n,k}$, in [4].

The cycles C_i for $1 \leq i \leq m$ are called *distinguished cycles*, and we similarly refer to vertices, edges, half-edges and edge paths as *distinguished* if they are contained in some C_i . A forest F in a viable marked graph Γ is called *admissible* if Γ/F is again viable and the induced marking is again admissible, ie, it is still a point in ΔK_n^m .

Lemma 2.2 (Admissible trees) *Let F be an admissible forest in Γ and T a tree in F . Then T intersects at most one distinguished cycle C , and if $T \cap C$ is nonempty then it must either be a single vertex or a connected edge path in C . Moreover, if F is a forest such that every tree in F satisfies this property, then F is admissible.*

Proof If T meets two distinguished cycles, C and C' , then the images of these cycles in Γ/F meet, violating viability. Now suppose $T \cap C \neq \emptyset$. The image of C in Γ/F must be homeomorphic to S^1 ; this tells us that $T \cap C$ is connected. But then the only options are a single vertex or a connected edge path.

For the converse, we need only observe that blowing down a tree that meets a single distinguished cycle in a connected subspace yields a viable graph. □

An example of an admissible and an inadmissible forest (for some marking ρ) are shown in gray in Figure 1.




Figure 1: From left to right: a viable graph, an admissible forest and an inadmissible forest

The action of ΣAut_n^m on K_{n+m}^W only affects markings, so we can consider the action of ΣAut_n^m on ΔK_n^m . Let

$$\Delta Q_n^m := \Delta K_n^m / \Sigma \text{Aut}_n^m$$

be the orbit space.

Proposition 2.3 *The subcomplex ΔK_n^m is contractible, and ΣAut_n^m acts on ΔK_n^m with finite stabilizers and finite quotient ΔQ_n^m .*

Proof There is an equivariant deformation retraction of K_{n+m}^W onto ΔK_n^m . This follows by a parallel argument to the proofs of Propositions 16 and 17 in [4]. The only difference is that our graphs have basepoints, but all the arguments carry through. That the stabilizers are finite follows because the stabilizers in $\text{Aut}(F_{n+m})$ of vertices in K_{n+m} are already finite. Lastly, ΣAut_n^m is transitive on markings of a given (viable) graph, and there are only finitely many homeomorphism types of graphs with rank $n + m$, so ΔQ_n^m is finite. □

If an element of ΣAut_n^m stabilizes a simplex then it fixes it pointwise, since the vertices of any simplex correspond to pairwise nonisomorphic graphs. This, together with the previous proposition, imply that ΔQ_n^m and ΣAut_n^m have the same rational homology; see for example Brown [2, Exercise 2, page 174].

A useful further reduction

There is a nice subcomplex of ΔK_n^m that will prove useful for our purposes, namely the subcomplex ∇K_n^m spanned by marked basepointed graphs in ΔK_n^m in which the basepoint p is not contained in a distinguished cycle. The action of ΣAut_n^m on ∇K_n^m similarly features finite stabilizers and finite quotient

$$\nabla Q_n^m := \nabla K_n^m / \Sigma\text{Aut}_n^m.$$

To keep straight which is which notationally, note that the symbol ∇ is “top heavy” compared to Δ , indicating that the distinguished cycles cannot be down at the basepoint.

Weighted degree

It is difficult to analyze ΔQ_n^m and ∇Q_n^m directly, and so we will work with a certain filtration. For a vertex (Γ, p, ρ) in ΔK_n^m , define the *weighted valency* $\text{val}_w(v)$ of a vertex v to be the number of undistinguished half-edges at v , plus half the number of distinguished half-edges. Define the *weighted degree* $d_w(\Gamma)$ to be

$$d_w(\Gamma) := 2n + m - \text{val}_w(p).$$

Note that $1 \leq \text{val}_w(p) \leq 2n + m$, and so $0 \leq d_w(\Gamma) \leq N$, where $N := 2n + m - 1$. As an example, the reader can verify that the weighted degree of the graph in Figure 1 is 10. We will also make use of the notion of *degree* from Hatcher and Vogtmann [9], which we define to be

$$d_0(\Gamma) := 2n + 2m - \text{val}(p).$$

If c denotes the number of distinguished cycles not containing p then $d_w = d_0 - c$. The reader curious about the motivation for defining weighted degree this way should glance ahead to the paragraph after Definition 4.2.

For $k \in \mathbb{N}_0$ let $\Delta K_{n,k}^m$ be the full subcomplex of ΔK_n^m spanned by marked basepointed graphs with weighted degree less than or equal to k . In particular for $k \geq N$, $\Delta K_{n,k}^m = \Delta K_n^m$. Also let $\nabla K_{n,k}^m = \Delta K_{n,k}^m \cap \nabla K_n^m$. The sequence of spaces

$$\nabla K_{n,0}^m \subseteq \nabla K_{n,1}^m \subseteq \dots$$

is a filtration of ∇K_n^m , and not of the whole contractible complex ΔK_n^m , but these smaller complexes will prove to be the right ones to inspect for various reasons. Note that when $m = 0$, $\nabla K_{n,k}^0 = \Delta K_{n,k}^0 = K_{n,k}$, the filtration of K_n by degree used in [9]. As a bit of foreshadowing to Section 4, note that the undistinguished loop and/or the distinguished loop on a stick (or “lollipop”) at the basepoint in Figure 1 could be removed without changing the weighted degree. This property is precisely the motivation for filtering ΔK_n^m and ∇K_n^m by weighted degree.

At this point we can state the Generalized degree theorem, which we will prove in the next section.

Theorem 3.5 (Generalized degree theorem) *For each $0 \leq k < N$, $\nabla K_{n,k}^m$ is $(k - 1)$ -connected.*

A lot of notation has been introduced by now, some of which will not appear again for a while, so for reference we collect it here. All the notation in the table will be considered fixed for the rest of the paper.

Symbol	Meaning
m	number of distinguished generators
n	number of undistinguished generators
p	basepoint
c	number of distinguished cycles not at p
K_{n+m}	space of all marked graphs
ΔK_n^m	only allow admissible markings and viable graphs
∇K_n^m	same as above, plus no distinguished cycles at basepoint
C_1, \dots, C_m	distinguished cycles
$\text{val}(v)$ (resp. $\text{val}_w(v)$)	valency (resp. weighted valency) of vertex v
$d_0(\Gamma)$ (resp. $d_w(\Gamma)$)	degree (resp. weighted degree) of graph Γ
$\nabla K_{n,k}^m$	points of ∇K_n^m with weighted degree at most k
$\nabla Q_{n,k}^m$	the quotient of $\nabla K_{n,k}^m$ by ΣAut_n^m

Table 1: Notation

3 A height function

We now define a height function h on the vertices of ΔK_n^m . This height function is related to the one defined by McEwen and Zaremsky in [16] on the space $K_n = \Delta K_n^0$. This will allow us to inspect the connectivity of $\nabla K_{n,k}^m$ using discrete Morse theory; see Bestvina and Brady [1] for background on discrete Morse theory.

Definition 3.1 (Features of graphs) Let (Γ, p, ρ) be a basepointed viable marked graph. For vertices v, v' in Γ , let the *distance* $d(v, v')$ be the number of edges in a minimal length edge path from v to v' . Also, for a subforest F of Γ , define the *level* $D(F)$ of F to be the smallest i such that F has a vertex at distance i from p . Let

$$\Lambda_i(\Gamma) := \{v \in \Gamma \mid d(p, v) = i\}$$

be the i^{th} level of Γ , so for example $\Lambda_0(\Gamma) = \{p\}$. If v is a vertex that is in a distinguished cycle C , and $d(p, v) \leq d(p, v')$ for any other vertex v' in C , then we will say that v is a *base vertex for* C , and call $i_C := d(p, v)$ the *base height* of C . If v is a base vertex for some C , call v a base vertex.

Note that the basepoint p is a base vertex if and only if it is distinguished, if and only if $c = m - 1$, where recall that c is defined as the number of distinguished cycles not containing p . In Figure 2 the distinguished cycle C is indicated by thick edges, the base vertices are the larger dots, and the basepoint is the largest dot at the bottom.




Figure 2: Distinguished cycle C with $i_C = 1$

Measurements contributing to the height function

For each $i \geq 0$, let $m_i(\Gamma)$ denote the number of base vertices in $\Lambda_i(\Gamma)$, define $n_i(\Gamma) := -|\Lambda_i(\Gamma)|$ and let

$$d_i(\Gamma) := \sum_{v \notin \Lambda_i} (\text{val}(v) - 2).$$

Note that $m_0 = m - c$, n_0 is constant -1 and $d_0 = 2n + 2m - \text{val}(p)$ is the degree. In general d_i can be thought of as counting the number of vertices *not* at level i , with higher valence vertices “counting for more.” Now define

$$h_i(\Gamma) := (m_i(\Gamma), n_i(\Gamma), d_i(\Gamma)) \quad \text{and set} \quad h(\Gamma) = (h_0(\Gamma), h_1(\Gamma), h_2(\Gamma), \dots)$$

with the lexicographic order. We remark that the height function used in [16] on the spine of Auter space was $(d_0, n_1, d_1, n_2, d_2, \dots)$, which encodes the same information

as our h when $m = 0$. Extend h to the vertices of ΔK_n^m via $h(\Gamma, p, \rho) = h(\Gamma)$. In general we will just write Γ to denote vertices of ΔK_n^m , with the basepoint and marking understood.

How forests affect the measurements

We need to understand how our height function changes upon blowing down a forest. The first thing to check is that it does indeed change.

Lemma 3.2 (Height always changes) *Let F be an admissible forest in Γ . Then blowing down F either increases or decreases $h_{D(F)}$. More precisely, blowing down F either increases $n_{D(F)}$, or else does not change $n_{D(F)}$ and decreases $d_{D(F)}$. Also, blowing down F does not change any h_i for $i < D(F)$.*

Proof First consider how the d_i are affected by blowing down a single edge ε . The endpoints of the edge become a single vertex, whose valency equals the sum of the valencies of the endpoints, minus two. If neither endpoint of ε lies in Λ_i , then the sum $\sum_{w \notin \Lambda_i} \text{val}(w)$ drops by 2 upon blowing down ε . It also decreases the number of vertices not in Λ_i by one, which implies that the sum $d_i = \sum_{w \notin \Lambda_i} (\text{val}(w) - 2)$ does not change. If one endpoint of ε lies in Λ_i and the other in Λ_{i+1} , then blowing down ε just eliminates one term of d_i , strictly decreasing it. Finally, if both endpoints of ε lie in Λ_i then blowing down ε does not change d_i .

Now consider the forest F . If F connects vertices in $\Lambda_{D(F)}$ then blowing down F decreases the number of such vertices and hence increases $n_{D(F)}$. Now suppose F does not connect any vertices in $\Lambda_{D(F)}$, so blowing down F does not change $n_{D(F)}$. Then F must connect a vertex that is not in $\Lambda_{D(F)}$ to a vertex that is in $\Lambda_{D(F)}$. Upon blowing down F one edge at a time (in any order), we will eventually blow down an edge connecting a vertex at level $D(F)$ to one at level $D(F) + 1$, and so we conclude that blowing down F strictly decreases $d_{D(F)}$.

Now let $i < D(F)$. Blowing down F induces a bijection $f: \Lambda_i(\Gamma) \rightarrow \Lambda_i(\Gamma/F)$, and moreover $f(v)$ is a base vertex if and only if v is. This tells us that m_i and n_i are unchanged upon blowing down F . That d_i is unchanged follows from the first paragraph. \square

In general, of all the terms in h changed by blowing down F , there is one that is lexicographically first, which we will call the *essential* term of F . Similarly, any blow-up has an essential term.

It will be important in the future to know some restrictions on the ways in which the various measurements can possibly change.

Observation 3.3 (Some restrictions) A blow-down at level i cannot decrease n_i , and a blow-up at level i cannot decrease d_i , though both blow-downs and blow-ups can either increase or decrease m_i .

Proof The first and last statements need no explanation (though perhaps we should reiterate that n_i is the *negative* of $|\Lambda_i|$). The second statement follows from the proof of Lemma 3.2, which says that a blow-down at level i cannot increase d_i . \square

Lemma 3.4 (Sublevel sets) *The complex ∇K_n^m is the sublevel set of ΔK_n^m defined by the inequality*

$$h \leq (0, 0, 0, \dots).$$

Also, $\nabla K_{n,k}^m$ is the sublevel set of ΔK_n^m defined by

$$h \leq (0, -1, k + m + 1, -1, 0, \dots).$$

Proof If $h(\Gamma) \leq (0, 0, 0, \dots)$ then in particular $m_0(\Gamma) = 0$ and so $\Gamma \in \nabla K_n^m$. Conversely, if $\Gamma \in \nabla K_n^m$ then $m_0(\Gamma) = 0$, and so $h(\Gamma) = (0, -1, d_0(\Gamma), \dots) \leq (0, 0, 0, \dots)$. This proves the first claim.

Now suppose $h(\Gamma) \leq (0, -1, k + m + 1, -1, 0, \dots)$, so $m_0(\Gamma) = 0$ and $d_0(\Gamma) \leq k + m + 1$. Since $m_0 = 0$ we have $d_0 = d_w + m$, so $d_w(\Gamma) \leq k + 1$. If this is even an equality, that is if $d_w(\Gamma) = k + 1$, then $h(\Gamma) = (0, -1, k + m + 1, m_1(\Gamma), \dots)$ and so $m_1(\Gamma) \leq -1$, which is absurd. Thus in fact $d_w(\Gamma) \leq k$ and so $\Gamma \in \nabla K_{n,k}^m$. Finally suppose $\Gamma \in \nabla K_{n,k}^m$. Then $m_0(\Gamma) = 0$ and $d_w(\Gamma) \leq k$. Hence $d_0(\Gamma) \leq k + m$, and so we conclude that $h(\Gamma) \leq (0, -1, k + m + 1, -1, 0, \dots)$. \square

The upshot of Lemmas 3.2 and 3.4 is that connectivity of $\nabla K_{n,k}^m$ can be determined by looking at descending links of vertices with respect to h . For a vertex Γ in ΔK_n^m , the *descending star* $st\downarrow(\Gamma)$ with respect to h is the set of simplices in the star of Γ whose other vertices all have strictly lower height than Γ . The *descending link* $lk\downarrow(\Gamma)$ consists of the faces of simplices in $st\downarrow(\Gamma)$ that do not themselves contain Γ .

The proof of the following main result will make up all of Section 5.

Proposition 5.14 (Connectivity of descending links) *The descending link $lk\downarrow(\Gamma)$ is either contractible or a wedge of spheres of dimension $d_w(\Gamma) - 1$.*

From this, our so-called Generalized degree theorem stated below follows quickly, as we now show. Recall that the weighted degree d_w of a graph can never exceed $N = 2n + m - 1$. Moreover, $d_w = N$ if and only if the basepoint p has valency 2 and is a base vertex. In addition to Proposition 5.14, we will need to use Lemma 5.12, which will also be proved later, but we state a relevant version of it here for reference.

Lemma 5.12 (Contractible case) *If p is a base vertex and $\text{val}(v) > 2$, then $\text{lk}\downarrow(\Gamma)$ is contractible.*

Theorem 3.5 (Generalized degree theorem) *For each $0 \leq k < N$, $\nabla K_{n,k}^m$ is $(k-1)$ -connected.*

Proof Since ΔK_n^m is contractible, it suffices by [1, Corollary 2.6] to show that for any vertex Γ in $\Delta K_n^m \setminus \nabla K_{n,k}^m$, the descending link $\text{lk}\downarrow(\Gamma)$ is at least $(k-1)$ -connected. Let Γ be such a vertex, so either $d_w(\Gamma) > k$, or else $d_w(\Gamma) \leq k$ and $m_0(\Gamma) = 1$. In the former case, $\text{lk}\downarrow(\Gamma)$ is $(d_w - 2)$ -connected by Proposition 5.14, and hence $(k-1)$ -connected. In the latter case, the basepoint p is a base vertex, and since $d_w(\Gamma) < N$, $\text{val}(p) > 2$, so by Lemma 5.12, $\text{lk}\downarrow(\Gamma)$ is contractible. \square

In the rest of this section we do some preliminary work with $\text{lk}\downarrow(\Gamma)$. Then in Section 4 we show how the Generalized degree theorem gives us homological stability results. Finally in Section 5 we prove Proposition 5.14, and, along the way, Lemma 5.12.

The d-up link and d-down link

There are two types of vertices in $\text{lk}\downarrow(\Gamma)$: those obtained from Γ by a descending blow-up, and those obtained by a descending blow-down. Here we say that a blow-up or blow-down is *descending* if the resulting graph has a lower height than the starting graph. Call the subcomplex of $\text{lk}\downarrow(\Gamma)$ spanned by vertices of the first type the *d-up link*, and the subcomplex spanned by vertices of the second type the *d-down link*. Any vertex in the d-up link is related to every vertex in the d-down link by a blow-down, so $\text{lk}\downarrow(\Gamma)$ is the simplicial join of the d-up- and d-down links. We remark that we only consider admissible blow-downs, and on the other hand observe that any blow-up of a viable graph is again viable. If a forest blow-down is descending we call the forest itself *descending*, and similarly we refer to *ascending* forests. By Lemma 3.2, every forest is either descending or ascending.

It will be important to have a somewhat explicit description of which forests are descending.

Lemma 3.6 (Interpreting the height function h) *Let F be an admissible forest in Γ with $i := D(F)$.*

- (i) *If $m_i(\Gamma/F) < m_i(\Gamma)$, then F is descending.*
- (ii) *If $m_i(\Gamma/F) > m_i(\Gamma)$, then F is ascending.*
- (iii) *If $m_i(\Gamma/F) = m_i(\Gamma)$ and F connects vertices in Λ_i , then F is ascending.*
- (iv) *If $m_i(\Gamma/F) = m_i(\Gamma)$ and F does not connect vertices in Λ_i , then F is descending.*

Proof The essential term of F occurs in h_i (Lemma 3.2), so the first two claims are immediate. Suppose $m_i(\Gamma/F) = m_i(\Gamma)$. If F connects vertices in Λ_i , then blowing down F increases n_i (Lemma 3.2) and so is ascending. If F does not connect vertices in Λ_i , then blowing down F does not change n_i , but decreases d_i (Lemma 3.2), so F is descending. \square

Again we have accumulated a lot of notation, which we collect in a table for easy reference.

Symbol	Meaning
$D(F)$	level of F
$\Lambda_i(\Gamma)$	i^{th} level of Γ
i_C	base height of distinguished cycle C , ie level of any base vertex of C
m_i	number of base vertices in Λ_i
n_i	negative number of vertices in Λ_i
d_i	sum of valencies minus two of vertices not in Λ_i
h_i	i^{th} term of height function h
$\text{lk}\downarrow(\Gamma)$	descending link of Γ in ΔK_n^m with respect to h

Table 2: More notation

4 Homological stability

As in Section 1, let G_n^m be any family of groups such that $P\Sigma\text{Aut}_n^m \leq G_n^m \leq \Sigma\text{Aut}_n^m$ for each n and m , and such that the inclusion $\Sigma\text{Aut}_n^m \hookrightarrow \Sigma\text{Aut}_{n+1}^m$ restricts to an inclusion $G_n^m \hookrightarrow G_{n+1}^m$. For any $0 \leq k < N$, the action of G_n^m on $\nabla K_{n,k}^m$ has finite stabilizers and finite quotient $\nabla K_{n,k}^m / G_n^m$. Hence by the Generalized degree theorem, $\nabla K_{n,k}^m / G_n^m$ has the same rational homology as G_n^m in dimensions i with $i < k$. To be precise, we have the following.

Lemma 4.1 (From groups to orbit spaces) *For any $0 \leq k < N$, we have that $H_i(\nabla K_{n,k}^m / G_n^m; \mathbb{Q})$ is isomorphic to $H_i(G_n^m; \mathbb{Q})$ for $i < k$, and $H_k(\nabla K_{n,k}^m / G_n^m; \mathbb{Q})$ surjects onto $H_k(G_n^m; \mathbb{Q})$.* \square

To get homological stability in n for G_n^m we can now look for homological stability in n for $\nabla K_{n,k}^m / G_n^m$. We will do this in a similar way as was done in the classical $m = 0$ case by Hatcher and Vogtmann [9, Section 5]. The vertices of $\nabla K_n^m / P\Sigma\text{Aut}_n^m$ are the homeomorphism types of basepointed graphs with m distinguished oriented cycles, disjoint and distinguishable from each other and disjoint from the basepoint. In

$\nabla K_{n,k}^m / \Sigma \text{Aut}_n^m$ the cycles become nonoriented and indistinguishable from each other, and in general $\nabla K_{n,k}^m / G_n^m$ interpolates between these two extremes. Exactly as in [9], we have a map

$$\nu: \nabla K_{n,k}^m / G_n^m \hookrightarrow \nabla K_{n+1,k}^m / G_{n+1}^m$$

induced by sending a graph Γ to $\Gamma \vee S^1$, that is the graph with an extra (undistinguished) loop wedged to its basepoint.

To get stability in n , we want to be able to “detect” loops and theta subgraphs at the basepoint. If Γ has a loop at the basepoint p then Γ is in the image of ν , which is why we want to be able to detect loops. We will see in Proposition 4.5 why theta subgraphs at the basepoint are also useful.

First we set up the situation for stability in m . Instead of loops and theta subgraphs we will use certain subgraphs defined as follows.

Definition 4.2 (Lollipops and double lollipops) A *lollipop* in Γ is a subgraph ℓ consisting of an undistinguished nonloop edge ε (the *stick*) and a distinguished loop δ sharing a vertex $v \neq p$, such that ε and δ are the only edges incident to v .

To define *double lollipops* we first make precise the notion of an r -iterated lollipop wedge. Take a collection of lollipops ℓ_1, \dots, ℓ_r , with their univalent vertices specified as basepoints. Set $W_1 = \ell_1$. Assuming W_i has been constructed, we let W_{i+1} be the result of wedging ℓ_{i+1} , at its basepoint, onto W_i at *any* point of W_i , even perhaps a nonvertex point. A *double lollipop* then is a subgraph of Γ that is a 2-iterated lollipop wedge.

Define a map

$$\mu: \nabla K_{n,k}^m / \Sigma \text{Aut}_n^m \hookrightarrow \nabla K_{n,k}^{m+1} / \Sigma \text{Aut}_n^{m+1}$$

by sending Γ to $\Gamma \vee \ell$, where ℓ is a lollipop wedged to the basepoint. Unlike attaching an undistinguished loop, attaching a lollipop in this way changes the degree, but it does not change the weighted degree, so this is still fine. (Indeed this was precisely the impetus for defining weighted degree as we did.) We now describe how to detect the presence of these various subgraphs at the basepoint, as in [9, Lemma 5.2]. Following that, we will see why this gives us stability.

Lemma 4.3 (Detecting features at the basepoint) *Let Γ be a graph with basepoint p , rank $n + m$, weighted degree d_w , and m pairwise disjoint distinguished cycles, disjoint from p . The following hold:*

- (i) If $n > 2d_w + m$ then Γ has a loop at the basepoint.
- (ii) If $n > (3d_w + m)/2$ then Γ has either a loop at the basepoint or a theta graph wedge summand.
- (iii) If $m > 2d_w$ then Γ has a lollipop at the basepoint.
- (iv) If $m > 3d_w/2$ then Γ has a lollipop or a double lollipop at the basepoint.

Proof Since p is not contained in a distinguished cycle, we have that the degree d_0 is $d_0 = d_w + m$. The first two parts of the lemma then follow from [9, Lemma 5.2]. Next suppose that there are no lollipops at p , and we want to show that $m \leq 2d_w + 1$. We will induct on n . If $n = 0$ then every undistinguished edge in Γ is a separating edge. Let Γ' be the graph obtained by blowing down every undistinguished edge. Now Γ' is a *cactus graph* as in Collins [5], ie every edge is contained in a unique reduced cycle. Note that Γ' is no longer in ΔK_0^m , since the distinguished cycles are not disjoint, but Γ' has the same weighted degree d_w as Γ . Let b' be the number of cycles in Γ' at p and $c' = m - b'$ the number of cycles not at p . Since Γ had no lollipops (or loops) at p , Γ' has no loops at p . This tells us that $b' \leq c'$, and since $m = b' + c'$ we see that $m \leq 2c'$. Also, in Γ' , $c' = m - \text{val}(p)/2 = d_w$, so indeed $m \leq 2d_w$. This finishes the base case, and we also note that if additionally Γ has no double lollipops then $b' \leq c'/2$, so $m \leq 3c'/2 = 3d_w/2$.

Now assume $n > 0$. Then there exists an undistinguished edge ε that is not a separating edge. Let Γ_1 be the graph obtained from Γ by removing ε , and then if any bivalent vertices $v \neq p$ arise (or univalent vertices v), blowing down one of the edges containing v . Then Γ_1 is a connected graph with undistinguished rank $n - 1$ and m distinguished cycles. Let $a \in \{0, 1, 2\}$ be such that the weighted degree $d_w(\Gamma_1)$ of Γ_1 is $d_w - a$. In particular $a = 0$ if and only if ε is a loop at p , and $a = 1$ if and only if p is an endpoint of ε and ε is not a loop. The graph Γ_1 has at most two lollipops at the basepoint, say there are b of them, so $b \in \{0, 1, 2\}$. Let Γ_2 be the graph obtained by removing all lollipops at p in Γ_1 . Then the weighted degree $d_w(\Gamma_2)$ of Γ_2 is the same as Γ_1 , the undistinguished rank is $n - 1$, and there are $m - b$ distinguished cycles. By induction, $m - b \leq 2(d_w - a)$, so $m \leq 2d_w - (2a - b)$. It now suffices to show that $2a \geq b$. If $a = 0$ then $b = 0$, so suppose $a > 0$. Then the only case to check is when $b = 2$. But then p cannot be an endpoint of ε , so $a = 2$ and the result follows. We remark that the stronger statement $a \geq b$ even holds.

Lastly suppose that Γ has no lollipops or double lollipops at p . Let $b \in \{0, 1, 2\}$ be the number of lollipops in Γ_1 and $c \in \{0, 1, 2\}$ the number of double lollipops in Γ_1 , so $b + c \in \{0, 1, 2\}$. Let Γ_3 be the graph obtained by removing all lollipops and double lollipops at p in Γ_1 . Let $a \in \{0, 1, 2, 3, 4\}$ be such that Γ_3 has weighted

degree $d_w - a$. Again, $a = 0$ if and only if ε is a loop at p . Also, if ε is not a loop but p is an endpoint of ε then $a = 1 + c$, and otherwise $a = 2 + c$. See Figure 3 for some examples. By the induction hypothesis $m - (b + 2c) \leq 3(d_w - a)/2$, so $m \leq 3d_w/2 - (3a/2 - (b + 2c))$. It now suffices to show that $3a \geq 2b + 4c$. If $a = 0$ then $b = c = 0$, so suppose $a > 0$. If p is an endpoint of ε then $b + c \leq 1$ and $a = 1 + c$, so $2b + 4c \leq 2 + 2c = 2a < 3a$. Now suppose p is not an endpoint of ε , so $b + c \leq 2$ and $c = a - 2$. Then $2b + 4c \leq 4 + 2c = 2a < 3a$ and we are done. Again, we find that a stronger statement holds, namely $a \geq b + 2c$. \square




Figure 3

Remark 4.4 In the last two paragraphs of the proof, it is interesting that the induction would have run even with sharper bounds. In fact, whatever the best possible bound is for the $n = 0$ case automatically extends to all cases, as long as the slope is not less than 1. In particular, we can detect wedge summands that are r -iterated lollipop wedges, with increasingly better bounds as r grows. Ultimately, we find that whenever $m > d_w$, there is always some nontrivial wedge summand that is an r -iterated lollipop wedge for some r . However, since we currently do not have a way to make use of this fact to get better bounds for homological stability, we will content ourselves with just detecting lollipops and double lollipops.

Proposition 4.5 (Stability in n) *The map*

$$\nu: \nabla K_{n,k}^m / G_n^m \hookrightarrow \nabla K_{n+1,k}^m / G_{n+1}^m$$

is a homeomorphism for $2k + m < n + 1$ and a homotopy equivalence for $\frac{3k+m}{2} < n + 1$.

Proof The proof is very similar to the proof of Proposition 5.4 in [9]. If $2k + m < n + 1$ then every Γ in $\nabla K_{n+1,k}^m / G_{n+1}^m$ has a loop at p , so ν is a homeomorphism. Now suppose $(3k + m)/2 < n + 1$, and let Γ be a vertex not in the image of ν . Then Γ has no loops at p but does have at least one theta graph wedge summand. Let Θ be the subgraph of Γ consisting of all such theta graphs at p , say there are $r \geq 1$ of

them. Then $\Gamma = \Theta \vee \Gamma'$, for some Γ' with rank $n + m + 1 - 2r$. Now, the open star of Γ in $\nabla K_{n+1,k}^m / G_{n+1}^m$ is the product of open stars of Θ in $\nabla K_{2r,r}^0 / G_{2r}^0$ and Γ' in $\nabla K_{n+1-2r,k-r}^m / G_{n+1-2r}^m$. The former consists of a single simplex, since all nonloop edges in Θ are equivalent under automorphisms of Θ ; moreover, every other vertex of this star has lower weighted degree since blowing down any edge reduces d_w by 1. So, collapsing any nonloop edge of Θ gives a deformation retraction of the star of Γ into the image of ν . \square

As a remark, in [10] Hatcher and Vogtmann give some bounds to detect wedge summands of higher degree, and the possibility of collapsing these in a similar way to the theta wedge summands is examined. In the present situation though, this collapse could cause p to become distinguished, which is a problem. Hence we cannot immediately improve the bound to $(5k + m)/4 < n + 1$, as was done for the $m = 0$ case in [10]. It seems likely that we could nonetheless improve this bound by directly inspecting examples with low (weighted) degree, in the spirit of [10], but we leave this for future work.

Proposition 4.6 (Stability in m) *Let $\nabla Q_{n,k}^m := \nabla K_{n,k}^m / \Sigma \text{Aut}_n^m$. The map*

$$\mu: \nabla Q_{n,k}^m \hookrightarrow \nabla Q_{n,k}^{m+1}$$

is a homeomorphism for $2k < m + 1$ and a homotopy equivalence for $3k/2 < m + 1$.

Proof If $2k < m + 1$ then every Γ in $\nabla Q_{n,k}^{m+1}$ has a lollipop at p , so μ is a homeomorphism. Now suppose $3k/2 < m + 1$, and let Γ be a vertex not in the image of μ . Then Γ has no lollipops at p but does have at least one double lollipop. Let $\Lambda\Lambda$ be the subgraph of Γ consisting of all double lollipops at p , say there are $r \geq 1$ of them. Then $\Gamma = \Lambda\Lambda \vee \Gamma'$, for some Γ' with rank $n + m + 1 - 2r$. The open star of Γ in $\nabla Q_{n,k}^{m+1}$ is the product of open stars of $\Lambda\Lambda$ in $\nabla Q_{0,r}^{2r}$ and Γ' in $\nabla Q_{n-2r,k-r}^{m+1}$. We claim that there is a retraction of the former that yields a retraction of the star of Γ into the image of μ , similar to the previous proof. Consider the height function h from Section 3, thought of on $\nabla K_{0,r}^{2r}$, and note that since h only depends on ρ inasmuch as ρ determines which cycles are distinguished, h descends to a function \bar{h} on $\nabla Q_{0,r}^{2r}$. Since $\nabla Q_{0,r}^{2r}$ is not simplicial we think of \bar{h} as a height function in the sense of Bux [3]. It now suffices to show that the descending link $\overline{\text{lk}}_{\downarrow}(\Gamma)$ is contractible.

There are three homeomorphism types of double lollipops, depending on where the first lollipop is wedged to the second. If it is wedged to a point in the interior of the stick, call this Type 1. If it is wedged to a point on the distinguished cycle not in the stick, call this Type 2. If it is wedged to the vertex shared by the loop and the

stick call this Type 3. See Figure 4. If $\Lambda\Lambda$ has a double lollipop of Type 1 then blowing down the edge connecting the wedge point to p is descending (with essential term d_0). Moreover, every simplex in $\overline{\text{lk}}\downarrow(\Gamma)$ is compatible with this move since descending blow-ups cannot affect double lollipops of Type 1, so it is a cone point of $\overline{\text{lk}}\downarrow(\Gamma)$. Next, if $\Lambda\Lambda$ has a double lollipop of Type 2, then blowing down either edge connecting the wedge point to the top of the stick is descending (with essential term d_0). These edges differ by a homeomorphism of Γ , so they actually correspond to the same blow-down. Again, every simplex in $\overline{\text{lk}}\downarrow(\Gamma)$ is compatible with this move since descending blow-ups cannot affect double lollipops of Type 2, so it is a cone point of $\overline{\text{lk}}\downarrow(\Gamma)$. Finally suppose $\Lambda\Lambda$ has a double lollipop of Type 3. Consider the blow-up that pushes the base of the first cycle away from the wedge point, creating a double lollipop of Type 1. This is descending, with essential term m_1 , and since descending (admissible) blow-downs cannot affect double lollipops of Type 3, it is a cone point for $\overline{\text{lk}}\downarrow(\Gamma)$. We conclude that attaching Γ does not change the homotopy type, by [3, Lemma 4], so the result follows. \square




Figure 4: Types of double lollipops

There is evidence to suggest that the descending links $\overline{\text{lk}}\downarrow(\Gamma)$ are always contractible whenever there is a nontrivial wedge summand that is an iterated lollipop wedge. As indicated by Remark 4.4, this would imply that μ is a homotopy equivalence whenever $k \leq m$. From this we would also recover the fact that ΣAut_0^m has trivial rational homology, shown independently by Griffin and Wilson [8; 19]. For now though, we will content ourselves with the double lollipop situation.

Before proving our stability theorems, the reader may want a hint as to why we get n -stability for both ΣAut_n^m and $P\Sigma\text{Aut}_n^m$, but m -stability only for ΣAut_n^m . The key is, when looking at $P\Sigma\text{Aut}_n^m$, the distinguished cycles in the orbit space are distinguishable from each other; they each have a different “color,” so to speak. Thus when we look for features at the basepoint, we would be hunting for, say, a lollipop of one specific color, and this is too much to ask for. In the nonpure case, all the distinguished cycles have the same color so all we care about is finding a lollipop at the basepoint, period.

Since ν is natural with respect to $G_n^m \hookrightarrow G_{n+1}^m$ and μ is natural with respect to $\Sigma\text{Aut}_n^m \hookrightarrow \Sigma\text{Aut}_n^{m+1}$, we can now prove our main stability results.

Proof of Theorem 1.1 We know that for $0 \leq k < N$, if $(3k + m)/2 < n + 1$ then

$$H_i(G_n^m; \mathbb{Q}) \rightarrow H_i(G_{n+1}^m; \mathbb{Q})$$

is an isomorphism for all $i < k$, by Lemma 4.1 and Proposition 4.5. Assume that $n \geq (3(i + 1) + m)/2$, so in particular $n \geq 2$, and set $k = i + 1$. Then $(3k + m)/2 < n + 1$ and $k \leq (2n - m)/3$, which is less than N since $n \geq 2$. The result now follows. \square

Note that when $m = 0$, so $G_n^0 = \text{Aut}(F_n)$, we recover the stability bound for $\text{Aut}(F_n)$ given in [9], though not the improved one given in [10].

Proof of Theorem 1.2 We know that for $0 \leq k < N$, if $3k/2 < m + 1$ then

$$H_i(\Sigma\text{Aut}_n^m; \mathbb{Q}) \rightarrow H_i(\Sigma\text{Aut}_n^{m+1}; \mathbb{Q})$$

is an isomorphism for all $i < k$, by Lemma 4.1 and Proposition 4.6. If $n = 0$ then the homology groups are all 0 by [8; 19], so we can assume $n \geq 1$. Suppose $m > (3i + 1)/2$, so in particular $m \geq 1$, and set $k = i + 1$. Then $3k/2 = 3(i + 1)/2 < m + 1$, and also since $n, m \geq 1$ we get $k < (2m + 2)/3 \leq 2n + m - 1 = N$, so $k < N$. The result now follows. \square

5 Connectivity

The rest of this paper is devoted to proving Proposition 5.14, that $\text{lk}\downarrow(\Gamma)$ is $(d_w(\Gamma) - 1)$ -spherical. In reading these subsections, the reader may find it helpful to refer to the corresponding sections in McEwen and Zaremsky [16], which cover what amounts here to the classical $m = 0$ case.

We first collect some natural definitions that will be used in these subsections, including the important notion of a *decisive edge* in a graph.

Definition 5.1 (Edges in graphs) For an edge ε in a basepointed graph Γ with vertices v and v' , we call ε *horizontal* if $d(p, v) = d(p, v')$. Otherwise we call ε *vertical*. If ε is vertical, by comparing $d(v, p)$ and $d(v', p)$ we get a natural notion of the *top* vertex and *bottom* vertex of ε . A half-edge may also have either a top or a bottom. If a vertex v has only one incident vertical edge ε with v as its top, we call that edge *decisive at v* . In other words, if every minimal length path from v to p begins with ε , then ε is decisive at v . If an edge ε in Γ is decisive at its top vertex we call it a *decisive edge*. For example any separating edge is decisive.

Outline of the section

There are many technical arguments in this section, so we record here the important steps of the proof. The descending link $\text{lk}\downarrow(\Gamma)$ is the join of the d-down link and d-up link. We analyze the d-down link first, modeled as the geometric realization of the *poset of good admissible forests* in Γ , denoted $P(\Gamma)$. The notion of a forest being *good* is defined in Definition 5.2, along with the important related notion of it being *arcfree* or *arced*. The main result of the subsection on the d-down link is the following.

Proposition 5.8 (Homotopy type of the d-down link) *$P(\Gamma)$ is homotopy equivalent to a (possibly empty) wedge of spheres of dimension $V - c - 2$, where V is the number of vertices in Γ .*

To prove this, we attempt to relate the homotopy of $P(\Gamma)$ to that of certain subposets $P_0(\Gamma, \varepsilon) \subseteq P_1(\Gamma, \varepsilon)$, where ε is a cleverly chosen edge (called an *optimal edge*). For now we denote these by P_0 and P_1 for brevity. The easier of these to analyze is P_0 . If ε is distinguished, P_0 is contractible (Lemma 5.4). The case when ε is not distinguished is dealt with during the induction argument in the proof of Proposition 5.8, by realizing that $P_0(\Gamma, \varepsilon)$ is isomorphic to $P(\Gamma \setminus \varepsilon)$; the latter is well defined since ε is undistinguished. The connection between $P_1(\Gamma)$ and $P(\Gamma)$ is established in Lemma 5.7. The remaining step, and the hardest, is to establish the connection between P_0 and P_1 .

Proposition 5.6 (From P_0 to P_1) *Let ε be an optimal edge. Then $P_1(\Gamma, \varepsilon)$ is homotopy equivalent to $P_0(\Gamma, \varepsilon)$.*

First we enlarge P_0 to an intermediate poset called $P_{1/2}$. If ε is undistinguished then this actually equals P_1 and we are done. Now assume ε is distinguished. We define a new height function e on P_1 with $P_{1/2}$ as a level set, and analyze descending links in $P_1 \setminus P_{1/2}$ with respect to e . We claim they are contractible, and prove this by looking at a join factor called the *d-in link*. The key is to retract the d-in link to the star of a cleverly chosen forest (which is a path) γ .

This finishes the analysis of $P(\Gamma)$, and hence of the d-down link. After showing in Lemma 5.9 that the d-down link of Γ is contractible in the special case where Γ has a nonbase vertex with an admissible decisive edge, we turn our attention to the d-up link of Γ . This part does not deviate much from the corresponding part of [16].

Lemma 5.13 (Homotopy type of the d-up link) *Suppose Γ has no nonbase vertices with an admissible decisive edge. Moreover suppose every base vertex has minimal valency. Then the d-up link of Γ is homotopy equivalent to $\bigvee S^{d_0 - V}$.*

In Lemma 5.10 we decompose the d-up link into a join of complexes $\text{BU}\downarrow(v)$ of descending blow-ups at vertices v . These can be modeled as complexes $\Sigma\downarrow(v)$ of certain partitions of $\{1, \dots, n\}$. If v is not a base vertex, $\Sigma\downarrow(v)$ is a wedge of spheres (Lemma 5.11). If v is a base vertex, the complex $\Sigma\downarrow(v)$ is “usually” contractible (Lemma 5.12). The remaining cases are dealt with in the proof of Lemma 5.13.

With the homotopy types of the d-down link and d-up link in hand, we conclude that the descending link $\text{lk}\downarrow(\Gamma)$, being their join, is highly spherical.

This ends the outline. Now we begin the details.

5.1 Connectivity of the d-down link

In this section we analyze the d-down link of Γ . In order to get an induction to run in the proof of Proposition 5.8, we will need to lift the restriction on the valency of vertices. Our height function h does not work well with such graphs though, for instance the fact that nonbasepoint vertices have valency at least 3 was crucial to the proof in Lemma 3.2 that blowing down a forest F either increases $n_{D(F)}$ or decreases $d_{D(F)}$. Thanks to Lemma 3.6 though, we have a condition on forests that is equivalent to being descending for graphs $\Gamma \in \Delta K_n^m$, and does not refer to the functions n_i or d_i . For lack of a more clever name, we will call such forests *good* (defined below). For the rest of this subsection, Γ is a connected graph with basepoint p and m disjoint distinguished cycles, with no restriction on the valency of vertices. The definitions of Λ_i and m_i remain valid, and are as given previously. A reduced, non-self-intersecting edge path γ in Γ will be called an *arc* if both of its endpoints lie in $\Lambda_{D(\gamma)}$.

Definition 5.2 (Good forests) Let F be an admissible forest in Γ . Define

$$\Delta m_i(\Gamma, F) := m_i(\Gamma/F) - m_i(\Gamma)$$

for any i . Now let $i := D(F)$. If $\Delta m_i(\Gamma, F) < 0$ call F *base-decreasing*, if $\Delta m_i(\Gamma, F) > 0$ call F *base-increasing* and if $\Delta m_i(\Gamma, F) = 0$ call F *base-preserving*. If F connects vertices in Λ_i , or equivalently if F contains an arc γ with $D(\gamma) = D(F)$, call F *arc-ed*. If F does not connect vertices in Λ_i , call F *arc-free*. Finally, if F is base-decreasing, or if it is base-preserving and arc-free, call F *good*. A forest is *bad* if it is not good.

Lemma 3.6 says that for any $\Gamma \in \Delta K_n^m$, a forest F in Γ is descending if and only if it is good.

There are a few important technical observations about some “basic” admissible forests that we collect here.

Observation 5.3 (Good/bad edges and distinguished paths) Let ε be an admissible forest that consists of one edge. Let γ be an admissible forest that consists of a reduced, non-self-intersecting edge path contained in a distinguished cycle C .

- (i) Suppose ε is vertical. Then it is arcfree and cannot be base-decreasing. If it is distinguished, then it is base-preserving and hence good.
- (ii) Suppose ε is horizontal. Then it is arced and cannot be base-increasing. If it is base-decreasing then it must be distinguished. Hence ε is good if and only if it is distinguished and connects two base vertices.
- (iii) The distinguished edge path γ is bad if it is arced and $D(\gamma) > i_C$, and otherwise is good. In any case γ is not base-increasing.

Proof Item (i) follows by definition, and by the observation that a vertical distinguished edge cannot be base-increasing. For item (ii), note that if ε is base-decreasing then it connects base vertices, and then since it is admissible, it must be distinguished. The other points follow by definition. For item (iii), if $D(\gamma) > i_C$ then since $\gamma \subseteq C$, blowing down γ does not change $m_{D(\gamma)}$; hence if γ is arced then it is bad by definition. If γ is arcfree then it is base-preserving, and hence good by definition. For the final case, if γ is arced and $D(\gamma) = i_C$, then γ is base-decreasing and hence good. \square

See Figure 5 for some examples of item (iii) in the observation.




Figure 5: Good and bad distinguished edge paths

Posets of forests Let $P(\Gamma)$ be the poset of good admissible forests in Γ , ordered by inclusion. For $\Gamma \in \Delta K_n^m$, the d -down link of Γ is the geometric realization $|P(\Gamma)|$ of $P(\Gamma)$, so the goal of this section is to calculate the homotopy type of $|P(\Gamma)|$. For the rest of this section we will omit the vertical bars, and just refer to $P(\Gamma)$ as having a homotopy type. For each edge ε of Γ , let $P_1(\Gamma, \varepsilon)$ be the poset of all good admissible forests except the forest just consisting of ε , and let $P_0(\Gamma, \varepsilon) \subseteq P_1(\Gamma, \varepsilon)$ be the poset of good admissible forests that do not contain ε . Whenever Γ and ε are understood from context we will just write P , P_1 and P_0 . We call $P_1(\Gamma, \varepsilon)$ the *deletion* of ε , and $P_0(\Gamma, \varepsilon)$ the *strong deletion* of ε .

Lemma 5.4 (Strong deletion of distinguished edge) *For an admissible distinguished edge ε , $P_0(\Gamma, \varepsilon)$ is contractible.*

Proof Let C be the distinguished cycle containing ε , and let ϕ be the forest consisting of all edges of C other than ε . Since $D(\phi) = i_C$, ϕ is good by Observation 5.3(iii). Let $f: P_0 \rightarrow P_0$ be given by

$$F \mapsto F \cup \phi.$$

We claim that for $F \in P_0$, $F \cup \phi$ is an admissible good forest, so f is well defined. Since $\varepsilon \notin F$ and F is admissible, we see that $F \cup \phi$ is an admissible forest (Lemma 2.2). Let ϕ' be the image of ϕ in Γ/F , so

$$\Gamma/F \cup \phi = (\Gamma/F)/\phi'.$$

By the same argument as for ϕ , ϕ' is good. Now, F , being good, is by definition either base-decreasing, or else base-preserving and arcfree. Since ϕ' is good, if F is base-decreasing then so is $F \cup \phi$. In this case the claim follows. Now suppose F is base-preserving and arcfree. If ϕ' is base-decreasing then so is $F \cup \phi$, and we are done. The other option is that ϕ' is base-preserving and arcfree (like F). Then $F \cup \phi$ is base-preserving. We need to show it is arcfree, and then the claim will follow. Suppose not, and let $\gamma \subseteq F \cup \phi$ be an arc. Since F is arcfree, $\gamma \not\subseteq F$, and so the image γ' of γ in Γ/F is again an edge path. In fact γ' is an arc. But $\gamma' \subseteq \phi'$, so this is impossible. This finishes the proof of the claim, that $F \cup \phi \in P_0$.

We conclude that f is well defined, and so it follows from work of Quillen [17, Section 1.5] that P_0 is contractible. □

Optimal edges For an admissible edge ε with endpoints v_1 and v_2 , call ε *maximally distant* if among all admissible edges, ε maximizes the quantity $d(p, v_1) + d(p, v_2)$. This quantity is even (resp. odd) if ε is horizontal (resp. vertical). Hence all maximally distant edges have the same orientation, ie, horizontal or vertical. If a maximally distant edge ε maximizes the quantity $\Delta m_{D(\varepsilon)}(\Gamma, \varepsilon)$ among all maximally distant edges, call ε *optimal*.

Observation 5.5 (Good optimal edges) If there exists a *good* optimal edge, then either every maximally distant edge is vertical and good, or else every maximally distant edge is horizontal and connects base vertices (and so is good).

Proof Let ε be a good optimal edge and let δ be another maximally distant edge, so δ has the same orientation as ε . If ε is vertical then $\Delta m_{D(\varepsilon)}(\Gamma, \varepsilon) \geq 0$ (Observation 5.3(i)), but ε is also good, so this quantity must equal 0. Then since ε is optimal, it maximizes

this quantity, whence $\Delta m_{D(\varepsilon)}(\Gamma, \delta) = 0$ and δ is good. The horizontal case follows immediately from Observation 5.3(ii). \square

Next we want to relate P_0 and P_1 . This is the most involved step in the analysis of P .

Proposition 5.6 (From P_0 to P_1) *Let ε be an optimal maximally distant edge. Then $P_1(\Gamma, \varepsilon)$ is homotopy equivalent to $P_0(\Gamma, \varepsilon)$.*

Proof We begin by finding an intermediate poset that is more apparently homotopy equivalent to P_0 . Let $P_{1/2} = P_{1/2}(\Gamma, \varepsilon)$ be the subcomplex of P_1 spanned by good admissible forests F for which $F \setminus \{\varepsilon\}$ is again a (nonempty) good admissible forest. Call $P_{1/2}$ the *sufficiently strong deletion* of ε . By definition,

$$P_0 \subseteq P_{1/2} \subseteq P_1.$$

Let $f: P_{1/2} \rightarrow P_{1/2}$ be given by $F \mapsto F \setminus \{\varepsilon\}$. This is a well defined poset map that is the identity on its image P_0 , and so induces a homotopy equivalence between $P_{1/2}$ and P_0 by [17, Section 1.3]. It remains to relate $P_{1/2}$ to P_1 .

Case 1: Undistinguished optimal edge First suppose that ε is undistinguished, and we claim that $P_{1/2} = P_1$. Let $F \in P_1$ and let $i := D(F)$. We want to show that $F \setminus \{\varepsilon\}$ is good. We may assume ε is (properly) contained in F , which since ε is maximally distant tells us that $D(F \setminus \{\varepsilon\}) = i$. If ε' is the image of ε in $\Gamma/(F \setminus \{\varepsilon\})$ then ε' is undistinguished, and so cannot be base-decreasing (Observation 5.3(i) and (ii)). Hence

$$\Delta m_i(\Gamma, F) = \Delta m_i(\Gamma/(F \setminus \{\varepsilon\}), \varepsilon') + \Delta m_i(\Gamma, F \setminus \{\varepsilon\}) \geq \Delta m_i(\Gamma, F \setminus \{\varepsilon\}).$$

If F is arcfree then $F \setminus \{\varepsilon\}$ is too. From this fact and the above equation, we conclude (from the definition of good) that if F is good then so is $F \setminus \{\varepsilon\}$. We remark that so far we have not used the hypothesis that ε is optimal, just that it is maximally distant.

Case 2: Distinguished optimal edge Now assume ε is distinguished, so we know $\Delta m_{D(\varepsilon)}(F, \varepsilon) \leq 0$ (Observation 5.3(iii)). We have to do a bit more work in this case. Define a height function e on P_1 as follows. For $F \in P_1$, if $F \in P_{1/2}$ set $e(F) = 0$ and otherwise let $e(F)$ be the number of edges in F . Since adjacent vertices (forests) in $P_1 \setminus P_{1/2}$ have different e values, we can build up from $P_{1/2}$ to P_1 by gluing in vertices along their descending links. We claim these descending links are contractible, so by [1, Corollary 2.6] the homotopy type does not change, and the result follows. The descending link of $F \in P_1 \setminus P_{1/2}$ is the join of two subcomplexes, which we will call the *d-out link* and the *d-in link*. The d-out link is spanned by forests in $P_{1/2}$ containing F , and the d-in link by forests in P_1 properly contained in F . In words, the d-in link is the complex of good admissible forests properly contained in F . It suffices to show that the d-in link is contractible.

Calculating Δm_i A forest F in P_1 but not in $P_{1/2}$ is characterized by F being good and $F \setminus \{\varepsilon\}$ being bad. This is a relatively specific situation, and we will be able to restrict the possibilities quite a bit. First of all, $\varepsilon \subseteq F$, and ε is maximally distant so $D(F \setminus \{\varepsilon\}) = i := D(F)$. Now we claim the following.

Claim The forests F and $F \setminus \{\varepsilon\}$ are arced, and

$$\begin{aligned} \Delta m_i(\Gamma, F) &= -1, \\ \Delta m_i(\Gamma/(F \setminus \{\varepsilon\}), \varepsilon') &= -1, \\ \Delta m_i(\Gamma, F \setminus \{\varepsilon\}) &= 0. \end{aligned}$$

Proof of the claim Consider again the equation

$$\Delta m_i(\Gamma, F) = \Delta m_i(\Gamma/(F \setminus \{\varepsilon\}), \varepsilon') + \Delta m_i(\Gamma, F \setminus \{\varepsilon\}),$$

where ε' is the image of ε in $\Gamma/(F \setminus \{\varepsilon\})$. Since F is good and $F \setminus \{\varepsilon\}$ is bad, $\Delta m_i(\Gamma, F) \leq 0$ and $\Delta m_i(\Gamma, F \setminus \{\varepsilon\}) \geq 0$ by definition. Also, $\Delta m_i(\Gamma/(F \setminus \{\varepsilon\}), \varepsilon')$ can only be $-1, 0$ or 1 since ε' is a single edge. In fact it cannot be 1 , since ε' is distinguished (Observation 5.3(i) and (ii)). If it equals 0 then all the terms in the equation are 0 , but then the definitions of good and bad necessitate that F is arcfree and $F \setminus \{\varepsilon\}$ is arced, which is absurd. Thus $\Delta m_i(\Gamma/(F \setminus \{\varepsilon\}), \varepsilon') = -1$. In particular $D(\varepsilon') = i$ in $\Gamma/(F \setminus \{\varepsilon\})$ (Lemma 3.2).

Now there are two possibilities for $\Delta m_i(\Gamma, F)$ and $\Delta m_i(\Gamma, F \setminus \{\varepsilon\})$, namely they either equal -1 and 0 (which we want), or else 0 and 1 . We know that ε' , being base-decreasing, must connect base vertices (Observation 5.3(i) and (ii)), and so in particular F must be arced, with an arc containing ε and connecting base vertices. Since F is good it therefore, by definition, must be base-decreasing, and so $\Delta m_i(\Gamma, F) = -1$. We now know that the equalities in the claim all hold, and F is arced. That $F \setminus \{\varepsilon\}$ is even arced follows by definition, now that we know it is base-preserving and bad. \square

Having understood the situation sufficiently, we now hunt for a way to retract the d-in link to a point.

A crucial arc in F Let C be the distinguished cycle containing ε . Since $\varepsilon \subseteq F$ and F is admissible, we know by Lemma 2.2 that $F \cap C$ is a forest (ie, it does not have isolated vertices). Let γ' be the connected edge path in $F \cap C$ containing ε . By the proof of the claim, γ' must contain an arc at level $D(F)$ that in turn contains ε . Let γ be the shortest arc in γ' containing ε with $D(\gamma) = D(F)$. If $\gamma = \varepsilon$ then $D(F) = D(\varepsilon)$, and ε being both an arc and an optimal edge implies that it, and so every edge of F , is horizontal and connects base vertices (Observation 5.5). Hence

$F \setminus \{\varepsilon\}$ is base-decreasing, which we know is not the case. We can therefore assume γ properly contains ε . According to Observation 5.3(iii), γ is base-decreasing, hence good, and $\gamma \setminus \{\varepsilon\}$ is nonbase-increasing. By minimality of γ , $\gamma \setminus \{\varepsilon\}$ is also arcfree, hence good. Since $F \setminus \{\varepsilon\}$ is bad, this means γ does not equal F . Hence γ is a good, admissible proper subforest of F , so γ is in the d-in link of F . See Figure 6 for an idea of γ' and γ .




Figure 6: The forest $F \cap C$, and edge paths γ' and γ

Contractibility of the d-in link The idea now is to retract the d-in link to the relative star of γ . We claim that for any F' in the d-in link, $F' \cup \gamma$ is also in the d-in link. The things to show are that it is admissible, good, and does not equal all of F . First note that $F' \cup \gamma$ is admissible, since it is contained in F . Next we want to see that it is good. If $\gamma \subseteq F'$ there is nothing to show, so we can assume rather that the image of γ in Γ/F' is an arc, which necessarily connects base vertices and so is base-decreasing. Since F' is good we conclude that $F' \cup \gamma$ is base-decreasing, and so is also good.

It remains only to show that $F' \cup \gamma \neq F$.

Claim For any $\emptyset \neq \delta \subseteq \gamma$, $F \setminus \delta$ is bad.

This can be phrased colloquially as: if removing ε from F turns it bad, then removing any part of γ from F turns it bad. Since F' is good, this will then imply that it is not of the form $F \setminus \delta$, and so $F' \cup \gamma \neq F$.

Proof of claim Note that if $\varepsilon \notin \delta$ and $F \setminus \delta$ is good, then the connected component of $(F \setminus \delta) \cap C$ containing ε does not connect base vertices (by minimality of γ). As seen in the previous paragraphs, this was a necessary requirement for a forest to be in $P_1 \setminus P_{1/2}$, so we conclude that, instead, $F \setminus \delta$ is in $P_{1/2}$, whence by definition we have that $F \setminus (\delta \cup \{\varepsilon\})$ is good. All of this is to say that, if $F \setminus \delta$ is good then so is $F \setminus (\delta \cup \{\varepsilon\})$. Taking the contrapositive, if $F \setminus (\delta \cup \{\varepsilon\})$ is bad then so is $F \setminus \delta$, so we can assume without loss of generality that $\varepsilon \subseteq \delta$.

Now, since $F \setminus \{\varepsilon\}$ is arced and γ is minimal we have $D(F \setminus \gamma) = i$, and so also $D(F \setminus \delta) = i$. Since the edges of γ are distinguished, we know from Observation 5.3(iii) that

$$\Delta m_i(\Gamma, F \setminus \delta) \geq \Delta m_i(\Gamma, F \setminus \{\varepsilon\}) = 0,$$

so to show $F \setminus \delta$ is bad, it suffices to assume it is arcfree and prove it is base-increasing. We will use the equation

$$\Delta m_i(\Gamma, F \setminus \delta) + \Delta m_i(\Gamma/(F \setminus \delta), (\delta \setminus \{\varepsilon\})') = \Delta m_i(\Gamma, F \setminus \{\varepsilon\}) = 0,$$

where $(\delta \setminus \{\varepsilon\})'$ is the image of $\delta \setminus \{\varepsilon\}$ in $\Gamma/(F \setminus \delta)$. For $F \setminus \{\varepsilon\}$ to be arced and $F \setminus \delta$ to be arcfree, there must exist an arc in $F \setminus \{\varepsilon\}$ containing an edge of $\delta \setminus \{\varepsilon\}$. In particular, $(\delta \setminus \{\varepsilon\})'$ is an arced forest consisting of distinguished edges, with an arc connecting base vertices. This must be base-decreasing, again by Observation 5.3(iii), so $\Delta m_i(\Gamma/(F \setminus \delta), (\delta \setminus \{\varepsilon\})') < 0$ and the above equation becomes

$$\Delta m_i(\Gamma, F \setminus \delta) > 0,$$

so $F \setminus \delta$ is bad and we are done. □

This finishes the proof of the claim, so now we know $F' \cup \gamma$ is in the d-in link of F . In particular the d-in link is contractible by [17, Section 1.5]. Then this finishes the proof that $P_1(\Gamma, \varepsilon) \simeq P_0(\Gamma, \varepsilon)$. □

Next we want to relate P_1 to P .

Decomposing P using ε In general if ε is any admissible good edge, then we have

$$P(\Gamma) = P_1(\Gamma, \varepsilon) \cup \text{st}(\varepsilon),$$

$$P_1(\Gamma, \varepsilon) \cap \text{st}(\varepsilon) = \text{lk}(\varepsilon),$$

where star and link here are taken in $P(\Gamma)$. The results up to this point provide tools to analyze $P_1(\Gamma, \varepsilon)$, and the next lemma tells us something about $\text{lk}(\varepsilon)$.

Lemma 5.7 (Links in the d-down link) *Let ε be an optimal edge in Γ such that $\varepsilon \in P(\Gamma)$, ie, ε is good. Let F be an admissible forest properly containing ε . Then $F \in P(\Gamma)$ if and only if $F/\varepsilon \in P(\Gamma/\varepsilon)$. Moreover, $\text{lk}(\varepsilon) \cong P(\Gamma/\varepsilon)$.*

Proof Let $i := D(F) = D(F/\varepsilon)$. Since ε is good, $\Delta m_i(\Gamma, \varepsilon) \in \{-1, 0\}$. First suppose that $\Delta m_i(\Gamma, \varepsilon) = 0$, for example if $D(\varepsilon) > i$. Since ε is optimal, F is arced if and only if F/ε is arced. Also,

$$\Delta m_i(\Gamma, F) = \Delta m_i(\Gamma/\varepsilon, F/\varepsilon) + \Delta m_i(\Gamma, \varepsilon),$$

so $\Delta m_i(\Gamma, F) = \Delta m_i(\Gamma/\varepsilon, F/\varepsilon)$. Hence, F is base-decreasing if and only if F/ε is, and F is base-preserving and arcfree if and only if F/ε is, which implies that $F \in P(\Gamma)$ if and only if $F/\varepsilon \in P(\Gamma/\varepsilon)$.

Next suppose $\Delta m_i(\Gamma, \varepsilon) = -1$, so $D(\varepsilon) = i$. We claim that in fact F and F/ε must both be base-decreasing, and hence good. First observe that every maximally distant edge is optimal, since the maximized quantity $\Delta m_i(\Gamma, \varepsilon) = -1$ is also minimized. We know that ε , and indeed every maximally distant edge, is horizontal and connects base vertices (Observations 5.3(i), (ii) and 5.5). In particular since $D(\varepsilon) = i$, every edge of F must be maximally distant, and so connects base vertices. Since F has more than one edge, $\Delta m_i(\Gamma, F) \leq -2$, so F is base-decreasing. Also,

$$\Delta m_i(\Gamma/\varepsilon, F/\varepsilon) = \Delta m_i(\Gamma, F) - \Delta m_i(\Gamma, \varepsilon) \leq -2 + 1 = -1,$$

so F/ε is base-decreasing.

Now consider the map

$$f: \text{lk}(\varepsilon) \rightarrow P(\Gamma/\varepsilon)$$

sending F to F/ε . This is well-defined by the previous paragraphs. We claim that f is bijective. Let $\Phi \in P(\Gamma/\varepsilon)$. There are two forests in Γ that map to Φ under blowing down ε , one that contains ε and one that does not (already this shows that f is injective). Let Φ' be the one that does, so $\Phi' \in \text{lk}(\varepsilon)$ and $f(\Phi') = \Phi$. If Φ was admissible then Φ' is too. Also, if Φ was good then so is Φ' , again by the previous paragraphs. So f is an isomorphism. \square

Let V be the number of vertices in Γ and E_{ad} the number of admissible edges. The next two results are generalizations of Proposition 3.2 and Lemma 3.3 from [16]. Recall that c is the number of distinguished cycles not at p .

Proposition 5.8 (Homotopy type of the d-down link) *$P(\Gamma)$ is homotopy equivalent to a (possibly empty) wedge of spheres of dimension $V - c - 2$.*

Proof The proof is similar to the proof of Proposition 2.2 in Vogtmann [18] and of Proposition 3.2 in [16]. We induct on the number of admissible edges E_{ad} . Since undistinguished loops do not affect $P(\Gamma)$, V or c , we may assume there are none. The base case is $E_{ad} = 0$, for which $P(\Gamma)$ is empty, ie S^{-1} . When $m > 0$, if there are no admissible edges then $V = m$ and $c = m - 1$. If $m = 0$ and there are no admissible edges then $V = 1$ and $c = 0$. In both cases, $-1 = V - c - 2$, which finishes the base case.

Now assume $E_{ad} > 0$, so in particular there exists a maximally distant edge. Let ε be an optimal (maximally distant) edge. First suppose that ε is distinguished. By Lemma 5.4

and Proposition 5.6, $P_1(\Gamma, \varepsilon)$ is contractible. If ε is bad then $P(\Gamma) = P_1(\Gamma, \varepsilon)$ and we are done, so assume ε is good. Then $\text{lk}(\varepsilon) \cong P(\Gamma/\varepsilon)$ by Lemma 5.7, and $P(\Gamma/\varepsilon)$ has fewer than E_{ad} admissible edges, so by induction $\text{lk}(\varepsilon)$ is $(V - c - 3)$ -spherical. Since

$$P(\Gamma) = P_1(\Gamma, \varepsilon) \cup \text{st}(\varepsilon),$$

$$P_1(\Gamma, \varepsilon) \cap \text{st}(\varepsilon) = \text{lk}(\varepsilon),$$

we conclude that $P(\Gamma)$ is $(V - c - 2)$ -spherical.

Next suppose that ε is not distinguished, and is not a separating edge. By the same argument as above, if ε is good then $\text{lk}(\varepsilon)$ is $(V - c - 3)$ -spherical (and if ε is bad then $P = P_1$ anyway), so we just need to inspect P_1 , which by Proposition 5.6 is homotopy equivalent to P_0 . Since ε is not a separating edge, we can remove it from Γ and we still have a connected graph with m distinguished cycles and V vertices, and strictly fewer admissible edges. By induction then, $P(\Gamma \setminus \varepsilon)$ is $(V - c - 2)$ -spherical (since c did not change either). Consider the map

$$g: P(\Gamma \setminus \varepsilon) \rightarrow P_0(\Gamma, \varepsilon)$$

induced by $\Gamma \setminus \varepsilon \hookrightarrow \Gamma$. Adding ε to the graph cannot affect whether a forest F in $\Gamma \setminus \varepsilon$ is admissible or not. Also, since ε is maximally distant, ε cannot be decisive, so adding ε to the graph does not change the levels Λ_i . In particular adding ε cannot affect whether a forest F in $\Gamma \setminus \varepsilon$ is good or bad, so g is an isomorphism. We conclude that $P_0(\Gamma, \varepsilon)$ is $(V - c - 2)$ -spherical, and hence so is $P(\Gamma)$.

Lastly suppose ε is not distinguished, but is an (admissible) separating edge. If ε is good then for any $F \in P(\Gamma)$, $F \cup \varepsilon$ is again an admissible good forest. In this case $P(\Gamma)$ is contractible by [17, Section 1.5]. Incidentally, this completely finishes the $m = 0$ case. If ε is bad then it is base-increasing (Observation 5.3(i)) and so its top must be a base vertex. Since ε is maximally distant, and Γ has no undistinguished loops, ε is the stick of a lollipop ℓ . The graph $\Gamma \setminus \ell$ has $V - 1$ vertices and $c - 1$ distinguished cycles not at p , and has fewer admissible edges than Γ . Also, $P_0(\Gamma, \varepsilon) \cong P(\Gamma \setminus \ell)$, so by induction, $P(\Gamma) = P_1(\Gamma, \varepsilon) \simeq P_0(\Gamma, \varepsilon)$ is $(V - 1) - (c - 1) - 2 = (V - c - 2)$ -spherical. \square

Lemma 5.9 (Decisive edges) *If Γ has a nonbase vertex with an admissible decisive edge then $P(\Gamma)$ is contractible.*

Proof The proof has essentially the same structure as the previous one. Induct on E_{ad} . In the base case, there are no admissible edges, much less admissible decisive edges, so the claim is vacuously true. Now assume $E_{ad} > 0$. Let ε be an optimal maximally

distant edge, so $P_1(\Gamma, \varepsilon)$ and $P_0(\Gamma, \varepsilon)$ are homotopy equivalent. We want to find a decisive edge $\eta \neq \varepsilon$ with top a nonbase vertex. Being maximally distant, the only way ε could be decisive is if it is separating. If ε is a good separating edge, then $P(\Gamma)$ is already contractible with cone point ε . If ε is a bad separating edge, then its top is a base vertex. Hence our hypothesis allows us to assume there is a decisive edge $\eta \neq \varepsilon$ with top a nonbase vertex.

Now we want to prove that $P(\Gamma)$ is contractible. First suppose that ε is distinguished. Then $P_1(\Gamma, \varepsilon)$ is contractible, so if ε is bad we are done. If ε is good, we still have that $\text{lk}(\varepsilon) \cong P(\Gamma/\varepsilon)$ as in the previous proof. By Observation 5.3(i) and (ii), ε is either vertical, or is horizontal and connects base vertices. In either case, η maps to a decisive edge in Γ/ε , with a nonbase vertex for a top, and so $\text{lk}(\varepsilon)$ is contractible by induction. Therefore $P(\Gamma)$ is contractible. Now suppose ε is not distinguished. Again, $\text{lk}(\varepsilon)$ is contractible if ε is good, so we just need to inspect $P_0(\Gamma, \varepsilon)$. If ε is not a separating edge we may remove it as in the previous proof and get that $P_0(\Gamma, \varepsilon) \cong P(\Gamma \setminus \varepsilon)$ is contractible by induction. The only case remaining is when ε is a separating edge whose top is a distinguished vertex, so it is the stick of a lollipop ℓ . Then η is still a decisive edge in $\Gamma \setminus \ell$, so $P(\Gamma) = P_0(\Gamma, \varepsilon) \cong P(\Gamma \setminus \ell)$ is contractible by induction. \square

5.2 Connectivity of the d-up link

Now consider the d-up link of Γ . We return to only considering graphs coming from ΔK_n^m , so all vertices $v \neq p$ are at least trivalent and p is at least bivalent. Let $\text{BU}(v)$ be the poset of all blow-ups at the vertex v , and let $\text{BU}\downarrow(v)$ be the poset of descending blow-ups at v . We will use the combinatorial framework for graph blow-ups described Culler and Vogtmann in [6] and Vogtmann in [18], so we think of $\text{BU}(v)$ as the poset of *compatible partitions* of the set of incident half-edges. Let $[n] := \{1, \dots, n\}$, and consider partitions of $[n]$ into two blocks. Denote such a partition by $\alpha = \{a, \bar{a}\}$, where $1 \in a$. Distinct partitions $\{a, \bar{a}\}$ and $\{b, \bar{b}\}$ are called *compatible* if either $a \subset b$ or $b \subset a$. Let $\Sigma(v)$ be the simplicial complex of partitions $\alpha = \{a, \bar{a}\}$ of $[\text{val}(v)]$ into blocks a and \bar{a} such that a and \bar{a} each have at least two elements. (If v is the basepoint p , then one block may have size one, since p is allowed to be bivalent.) That is, the vertices of $\Sigma(v)$ are partitions, and a j -simplex is given by a collection of $j + 1$ distinct, pairwise compatible partitions. Also let $\Sigma\downarrow(v)$ be the subcomplex of $\Sigma(v)$ spanned by descending partitions, ie, partitions that correspond to descending single-edge blow-ups.

For $v \neq p$, the geometric realization $|\text{BU}(v)|$ of $\text{BU}(v)$ is isomorphic to the barycentric subdivision of $\Sigma(v)$. The idea is that a partition describes an *ideal edge*, ie, an edge blow-up at a vertex, and the blocks a and \bar{a} indicate which half-edges attach to

which endpoints of the new edge. See [6] and [18] for more details. The geometric realization $|\text{BU}\downarrow(v)|$ contains the barycentric subdivision of $\Sigma\downarrow(v)$ as a subcomplex, and any simplex in $|\text{BU}\downarrow(v)|$ has at least one vertex in $\Sigma\downarrow(v)$. Hence there is a map $|\text{BU}\downarrow(v)| \rightarrow |\Sigma\downarrow(v)|$ sending each simplex to its face spanned by vertices in $\Sigma\downarrow(v)$, which induces a deformation retraction from $|\text{BU}\downarrow(v)|$ to $|\Sigma\downarrow(v)|$.

The next lemma relates the d-up link of Γ to these complexes $\Sigma\downarrow(v)$. The proof is very similar to the proof of [16, Proposition 4.5].

Lemma 5.10 (Local to global) *Let $\text{BU}\downarrow(\Gamma) := \ast_{v \in \Gamma} \text{BU}\downarrow(v)$. Then $|\text{BU}\downarrow(\Gamma)|$ is homotopy equivalent to the d-up link of Γ .*

Proof For a poset P , let \underline{P} be $P \sqcup \{\perp\}$, where \perp is a formal minimal element. Then we have that $P \ast Q \simeq \underline{P} \times \underline{Q} - \{(\perp, \perp)\}$. Let

$$U := \left\{ f \in \prod_v \underline{\text{BU}}(v) - \{(\perp)_v\} \mid f \text{ is descending} \right\},$$

so the geometric realization $|U|$ is the d-up link. Define a poset map $r: U \rightarrow \text{BU}\downarrow(\Gamma)$ via

$$(f_v)_v \mapsto \left(\begin{cases} f_v & \text{for } f_v \in \text{BU}\downarrow(v), \\ \perp & \text{for } f_v \notin \text{BU}\downarrow(v), \end{cases} \right)_v$$

where f_v is a blow-up at v in the tuple f . This map is well defined since if f is descending then f_v must be descending for some v . By construction, r is the identity when restricted to $\text{BU}\downarrow(\Gamma)$. Also, $r(f) \leq f$ for all $f \in U$, and so by [17, 1.3] this induces a homotopy equivalence between $|U|$ and $|\text{BU}\downarrow(\Gamma)|$. □

In particular the d-up link is homotopy equivalent to $\ast_{v \in \Gamma} \Sigma\downarrow(v)$, so we can analyze the d-up link by looking at the complexes $\Sigma\downarrow(v)$. In light of Lemma 5.9, one important situation is when v is a nonbase vertex with no decisive edges.

Lemma 5.11 (No decisive edges, locally) *Suppose $v \neq p$ is a nonbase vertex with no decisive edge. Then $\Sigma\downarrow(v) \simeq \bigvee S^{\text{val}(v)-4}$.*

Proof We know that among the half-edges at v , at least two correspond to vertical edges with top v . Since v is a nonbase vertex, a blow-up at v is descending if and only if it separates some of these half-edges with top v . (This is due to Observation 3.3. The essential term will be $n_{d(p,v)}$.) Thus $\Sigma\downarrow(v)$ is isomorphic to the complex denoted $\text{SBU}(v)$ in [16], and the result is immediate from [16, Lemma 4.1, Proposition 4.3]. □

Next we describe one very important case for which the d-up link, and hence $\text{lk}\downarrow(\Gamma)$ is contractible. If a vertex $v \neq p$ has valency 3, or if $v = p$ and $\text{val}(v) = 2$, we say v has *minimal valency*. Otherwise we naturally say it has *nonminimal valency*.

Lemma 5.12 (Contractible case) *If Γ has a base vertex with nonminimal valency, then the d-up link is contractible, and so $\text{lk}\downarrow(\Gamma)$ is contractible.*

Proof Let v be a base vertex with nonminimal valency. By Lemma 5.10 it suffices to show that $\Sigma\downarrow(v)$ is contractible. Label the distinguished half-edges at v by c_1 and c_2 , and label the undistinguished half-edges by b_1, \dots, b_q . By hypothesis, $q > 1$, unless $v = p$ in which case $q > 0$. Let α_0 be the ideal edge at v that separates c_1, c_2 from all the other half-edges. See Figure 7 for an example. This is a descending blow-up, with essential term $m_{d(p,v)}$. Also, any partition of $\{c_1, c_2, b_1, \dots, b_q\}$ that separates c_1 and c_2 is ascending, so indeed $\Sigma\downarrow(v)$ is contractible with cone point α_0 . \square




Figure 7: The blow-up at v given by α_0 ; here m_1 goes from 2 to 1

We may now assume every base vertex has minimal valency, and so $\Sigma\downarrow(v)$ is empty for all base vertices v . Let V be the number of vertices of Γ , and recall that here $d_0 = d_0(\Gamma)$ is the degree of Γ , ie, $d_0 = 2n + 2m - \text{val}(p)$.

Lemma 5.13 (Homotopy type of the d-up link) *Suppose Γ has no nonbase vertices with an admissible decisive edge. Moreover suppose every base vertex has minimal valency. Then the d-up link of Γ is homotopy equivalent to $\bigvee S^{d_0-V}$.*

Proof By Lemma 5.10, the d-up link is homotopy equivalent to $\ast_{v \in \Gamma} \Sigma\downarrow(v)$. Since $\Sigma\downarrow(p) = \emptyset$, this is the same as $\ast_{v \neq p} \Sigma\downarrow(v)$. Also, we are assuming that each base vertex $u \neq p$ has valency 3, so $\Sigma\downarrow(u) = \emptyset = S^{\text{val}(u)-4}$. Therefore by Lemma 5.11 the d-up link is homotopy equivalent to

$$\ast_{v \neq p} \left(\bigvee S^{\text{val}(v)-4} \right),$$

which is a wedge of spheres of dimension $(V - 2) + \sum_{v \neq p} (\text{val}(v) - 4)$. Observe that

$$\sum_{v \neq p} (\text{val}(v) - 2) = d_0,$$

so this dimension equals $(V - 2) + d_0 - 2(V - 1) = d_0 - V$. □

We can now prove our main result of this section. Here d_w is the weighted degree, which recall equals $d_0 - c$.

Proposition 5.14 (Connectivity of descending links) *The descending link $\text{lk}\downarrow(\Gamma)$ is either contractible or a wedge of spheres of dimension $d_w - 1$.*

Proof Assume that neither the d-up link nor d-down link is contractible. Then every base vertex has minimal valency (Lemma 5.12), and no nonbase vertex of Γ has a decisive edge (Lemma 5.9). By Proposition 5.8, $P(\Gamma) \simeq \bigvee S^{V-c-2}$, and by Lemma 5.13 the d-up link is homotopy equivalent to $\bigvee S^{d_0-V}$. Hence $\text{lk}\downarrow(\Gamma)$ is homotopy equivalent to

$$\left(\bigvee S^{V-c-2}\right) * \left(\bigvee S^{d_0-V}\right) \simeq \bigvee S^{V-c-2+d_0-V+1} = \bigvee S^{d_0-c-1} = \bigvee S^{d_w-1}.$$

This completes the proof. □

Remark 5.15 (Concluding remarks) We conclude with some questions that now naturally arise. First, the stable rational homology of ΣAut_n^0 in n is trivial, and the rational homology of ΣAut_0^m is trivial in every dimension, so it seems likely that the stable homology in m and n is always trivial; is this indeed the case? Some additional evidence for this is Theorem 7.4 in Jensen and Wahl [14], which implies that $H_1(P\Sigma\text{Aut}_n^m; \mathbb{Q}) = 0$ for any $n > 2$ and any $m \geq 0$. Second, there exist examples where $H_i(\Sigma\text{Aut}_n^0; \mathbb{Q}) = \mathbb{Q}$, but when can nontrivial rational homology occur in general, eg, if $m > 0$? This is an interesting question for outer automorphisms as well. Third, when $n = 0$ or $m = 0$, we have stable integral homology, so an obvious question is whether this holds in general. Fourth, we know that $H_i(P\Sigma\text{Aut}_0^m; \mathbb{Q})$ is not stable in m , but as stated above, $H_1(P\Sigma\text{Aut}_n^m; \mathbb{Q})$ is stably constant 0 for $n > 2$. What can we expect in general for $H_i(P\Sigma\text{Aut}_n^m; \mathbb{Q})$ in terms of stability in m ?

References

- [1] **M Bestvina, N Brady**, *Morse theory and finiteness properties of groups*, Invent. Math. 129 (1997) 445–470 MR1465330

- [2] **K S Brown**, *Cohomology of groups*, Graduate Texts in Mathematics 87, Springer, New York (1982) MR672956
- [3] **K-U Bux**, *Orbit spaces of subgroup complexes, Morse theory, and a new proof of a conjecture of Webb*, from: “Proceedings of the 1999 Topology and Dynamics Conference”, Topology Proc. 24 (1999) 39–51 MR1802676
- [4] **K-U Bux, R Charney, K Vogtmann**, *Automorphisms of two-dimensional RAAGS and partially symmetric automorphisms of free groups*, Groups Geom. Dyn. 3 (2009) 541–554 MR2529947
- [5] **D J Collins**, *Cohomological dimension and symmetric automorphisms of a free group*, Comment. Math. Helv. 64 (1989) 44–61 MR982561
- [6] **M Culler, K Vogtmann**, *Moduli of graphs and automorphisms of free groups*, Invent. Math. 84 (1986) 91–119 MR830040
- [7] **S Galatius**, *Stable homology of automorphism groups of free groups*, Ann. of Math. 173 (2011) 705–768 MR2784914
- [8] **J T Griffin**, *Diagonal complexes and the integral homology of the automorphism group of a free product*, Proc. Lond. Math. Soc. 106 (2013) 1087–1120 MR3066749
- [9] **A Hatcher, K Vogtmann**, *Cerf theory for graphs*, J. London Math. Soc. 58 (1998) 633–655 MR1678155
- [10] **A Hatcher, K Vogtmann**, *Rational homology of $\text{Aut}(F_n)$* , Math. Res. Lett. 5 (1998) 759–780 MR1671188
- [11] **A Hatcher, N Wahl**, *Stabilization for the automorphisms of free groups with boundaries*, Geom. Topol. 9 (2005) 1295–1336 MR2174267
- [12] **A Hatcher, N Wahl**, *Stabilization for mapping class groups of 3-manifolds*, Duke Math. J. 155 (2010) 205–269 MR2736166
- [13] **C A Jensen, J McCammond, J Meier**, *The integral cohomology of the group of loops*, Geom. Topol. 10 (2006) 759–784 MR2240905
- [14] **C A Jensen, N Wahl**, *Automorphisms of free groups with boundaries*, Algebr. Geom. Topol. 4 (2004) 543–569 MR2077676
- [15] **R A McEwen**, *Homological stability for the groups $\text{Out}P(n, t + 1)$* , PhD thesis, University of Virginia (2010)
- [16] **R McEwen, M C B Zaremsky**, *A combinatorial proof of the Degree theorem in Outer space*, to appear in New York J. of Math.
- [17] **D Quillen**, *Homotopy properties of the poset of nontrivial p -subgroups of a group*, Adv. in Math. 28 (1978) 101–128 MR493916
- [18] **K Vogtmann**, *Local structure of some $\text{Out}(F_n)$ -complexes*, Proc. Edinburgh Math. Soc. 33 (1990) 367–379 MR1077791

- [19] **J C H Wilson**, *Representation stability for the cohomology of the pure string motion groups*, *Algebr. Geom. Topol.* 12 (2012) 909–931 MR2928898

*Department of Mathematical Sciences, Binghamton University
Binghamton, NY 13902, USA*

`zaremsky@math.binghamton.edu`

<http://www.math.binghamton.edu/zaremsky/>

Received: 17 January 2013 Revised: 21 October 2013

