New York Journal of Mathematics

New York J. Math. **20** (2014) 217–228.

A combinatorial proof of the Degree Theorem in Auter space

Robert McEwen and Matthew C. B. Zaremsky

ABSTRACT. We use discrete Morse theory to give a new proof of Hatcher and Vogtmann's *Degree Theorem* in Auter space A_n. There is a filtration of A_n into subspaces $A_{n,k}$ using the *degree* of a graph, and the Degree Theorem says that each $A_{n,k}$ is $(k-1)$-connected. This result is useful, for example to calculate stability bounds for the homology of $\mathrm{Aut}(F_n)$. The standard proof of the Degree Theorem is global in nature. Here we give a proof that only uses local considerations, and lends itself more readily to generalization.

CONTENTS

1. Introduction

In this note we provide an alternate proof of Hatcher and Vogtmann's *Degree Theorem* in Auter space [HV98], using discrete Morse theory. The advantage of our proof is that it relies only on local data, and also lends itself more readily to certain generalizations. *Auter space* A_n is the space of rank-n basepointed marked metric graphs. In [HV98], a measurement called the *degree* of a graph was used to filter A_n into highly connected sublevel sets $A_{n,k}$, which were then used to produce stability bounds for the rational and integral homology of $\mathrm{Aut}(F_n)$. The key result was:

Theorem (Degree Theorem)**.** [HV98] *$A_{n,k}$ is $(k-1)$-connected.*

Received August 18, 2013.
2010 *Mathematics Subject Classification.* Primary 20F65; Secondary 57M07, 20F28.
Key words and phrases. Auter space, Degree Theorem, automorphisms of free groups.
The second named author gratefully acknowledges support from the SFB 701 of the DFG.

ISSN 1076-9803/2014

The proof of the Degree Theorem in [HV98] is done by globally deforming disks in A_n via an iterated process. Our proof here uses discrete Morse theory, as in [BB97], to reduce the problem to a purely local one. First we shift focus to the *spine* of Auter space, which we denote L_n. This is a combinatorial model for A_n that is a deformation retract. We construct a *height function* h on L_n that reduces the problem to asking whether the *descending links* with respect to h are highly connected. This is advantageous for being a local rather than global problem, and also lends itself more readily to generalization. For example a similar method has been used in [Zar14] to get stability results for the groups $\Sigma\mathrm{Aut}_n^m$ of *partially symmetric* automorphisms.

In Section 2 we describe the spine of Auter space L_n, and define the notion of the *degree* d_0 of a graph. We use the degree to filter L_n into sublevel sets $L_{n,k}$, as in [HV98]. We then define a height function h on L_n refining d_0, and consider the descending links of vertices in L_n with respect to h. The descending link of a vertex decomposes as a join of two complexes, called the *d-down-link* and *d-up-link*. In Section 3 we analyze the connectivity of the d-down-link, and in Section 4 we do the same for the d-up-link. The upshot of this is Corollary 5.1, that the descending links are all highly connected. From this we quickly obtain that $L_{n,k}$, and hence $A_{n,k}$ is $(k-1)$-connected; see Theorem 5.2.

Acknowledgments. The authors would like to thank Kai-Uwe Bux, who helped with a preliminary version of this paper, and Allen Hatcher for his comments and suggestions. Parts of this paper are based on results in the first named author's Ph.D. thesis [McE10], done at the University of Virginia.

2. Auter space, degree, and a height function

We begin by describing the *spine of Auter space* L_n introduced in [HV98]. Let R_n be the rose with n edges, i.e., the graph with a single vertex p_0 and n edges. Here by a *graph* we always mean a finite connected one-dimensional CW-complex, with the usual notions of vertices and edges. If Γ is a rank n graph with basepoint vertex p, a homotopy equivalence $\rho\colon R_n \to \Gamma$ taking p_0 to p is called a *marking* on Γ. Two markings are equivalent if there is a basepoint-preserving homotopy between them. We only consider graphs such that p is at least bivalent and all other vertices are at least trivalent. The spine L_n of Auter space is then the complex of marked basepointed rank n graphs (Γ, p, ρ), up to equivalence of markings.

To be more precise, L_n is a simplicial complex with a vertex for every equivalence class of triples (Γ, p, ρ). An r-simplex is given by a chain of *forest collapses* $\Gamma_r \xrightarrow{d_r} \Gamma_{r-1} \xrightarrow{d_{r-1}} \cdots \xrightarrow{d_1} \Gamma_0$ and markings $\rho_i\colon R_n \to \Gamma_i$ with the following diagram commuting up to homotopy.

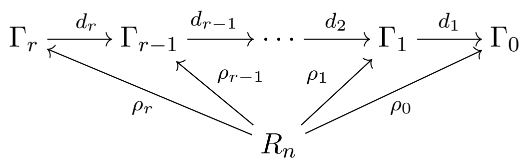

Here a *forest collapse* or *blow-down* $d\colon \Gamma \to \Gamma'$ is a (basepoint-preserving) homotopy equivalence of graphs that is given by collapsing each component of a forest F in Γ to a point. We will write the resulting graph as Γ/F. The reverse of a blow-down is, naturally, called a *blow-up*.

Let Γ be a graph with rank n, basepoint p and vertex set $V(\Gamma)$. The *degree* of Γ can be defined as

$$d_0(\Gamma) := \sum_{p \neq v \in V(\Gamma)} (\mathrm{val}(v) - 2)$$

or equivalently as $d_0(\Gamma) = 2n - \mathrm{val}(p)$ [HV98, Section 3]. Here $\mathrm{val}(v)$ is the valency of v, that is the number of half-edges incident to v. This is sometimes called the "degree" of the vertex, but we have reserved this word for the degree of a graph.

Definition 2.1 (Filtration by degree)**.** For $k \geq 0$, let $L_{n,k}$ be the subcomplex of L_n spanned by vertices represented by triples (Γ, p, ρ) with $d_0(\Gamma) \leq k$.

The Degree Theorem says that $A_{n,k}$ is $(k-1)$-connected, and this is equivalent to $L_{n,k}$ being $(k-1)$-connected [HV98, Section 5.1], which is what we will prove.

We now define some other measurements on Γ. For $v \in V(\Gamma)$ let $d(p,v)$ denote the minimum length of an edge path in Γ from v to p, and call $d(p,v)$ the *level* of v. Here we are treating each edge in the graph as having length 1. Define $\Lambda_i(\Gamma) := \{v \in V(\Gamma) \mid d(p,v) = i\}$, $n_i(\Gamma) := -|\Lambda_i(\Gamma)|$ and

$$d_i(\Gamma) := \sum_{v \in V(\Gamma) \setminus \Lambda_i(\Gamma)} (\mathrm{val}(v) - 2)$$

for $i \geq 0$. Note that $\Lambda_0(\Gamma) = \{p\}$, $n_0(\Gamma) = -1$, and $d_0(\Gamma)$ agrees with the definition of degree, so this is not an abuse of notation. Finally, define

$$h(\Gamma) := (d_0(\Gamma), n_1(\Gamma), d_1(\Gamma), n_2(\Gamma), d_2(\Gamma), \dots)$$

to be the *height* of the graph Γ, considered with the lexicographic ordering. This height function is a refinement of the degree function. Extend the definition of h to the vertices of L_n via $h(\Gamma, p, \rho) = h(\Gamma)$. For brevity, in the future we will often just refer to vertices in L_n as being graphs, rather than equivalence classes of triples (Γ, p, ρ).

Observation 2.2. *$L_{n,k}$ is the sublevel set of L_n defined by the inequality*

$$h(\Gamma) \leq (k, 1, 0, 0, \dots).$$

Proof. If $h(\Gamma) \leq (k, 1, 0, 0, \dots)$ then $d_0(\Gamma) \leq k$. Now suppose $d_0(\Gamma) \leq k$. If $d_0(\Gamma) < k$ then $h(\Gamma) < (k, 1, 0, 0, \dots)$. If $d_0(\Gamma) = k$ then since $n_1(\Gamma) \leq 0$ we have $h(\Gamma) < (k, 1, 0, 0, \dots)$. □

Any blow-down necessarily increases some n_i (that is, decreases some $|\Lambda_i|$), and so adjacent vertices in L_n have different heights. Hence h is a "true" height function, in the sense of [BB97]. This, together with Observation 2.2, means that the connectivity of $L_{n,k}$ can be deduced by inspecting the *descending links* with respect to h of vertices in $L_n \setminus L_{n,k}$. For a vertex Γ in L_n, the *descending star* $\mathrm{st}{\downarrow}(\Gamma)$ with respect to h is the set of simplices in the star of Γ whose vertices other than Γ all have strictly lower height than Γ. The *descending link* $\mathrm{lk}{\downarrow}(\Gamma)$ is the set of faces of simplices in $\mathrm{st}{\downarrow}(\Gamma)$ that do not themselves contain Γ.

There are two types of vertices in $\mathrm{lk}{\downarrow}(\Gamma)$: 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 full subcomplex of $\mathrm{lk}{\downarrow}(\Gamma)$ 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 $\mathrm{lk}{\downarrow}(\Gamma)$ is the simplicial join of the d-up- and d-down-links.

If blowing down the forest F is a descending blow-down, we will call the forest itself *descending*, and similarly a forest can be ascending. It will be a good idea to describe precisely which forests in a graph are ascending and descending. For a forest F in Γ define $D(F) := \min\{i \mid F \text{ has a vertex in } \Lambda_i\}$ to be the *level* of F. If there is an edge path in F from a vertex in $\Lambda_{D(F)}$ to another, distinct vertex in $\Lambda_{D(F)}$, we say that F *connects vertices in* $\Lambda_{D(F)}$.

Lemma 2.3. *If F connects vertices in $\Lambda_{D(F)}$ then F is ascending. Otherwise F is descending.*

Proof. Let $i := D(F)$. Blowing down F does not change any n_j or d_j for $j < i$. If F connects vertices in Λ_i, then blowing down F increases n_i, so F is ascending. If F does not connect any vertices in Λ_i, then blowing down F will not change n_i, but since each non-basepoint vertex of Γ is at least trivalent, d_i will be smaller in Γ/F than in Γ, and so F is descending. □

As a corollary to the proof we obtain:

Corollary 2.4. *A blow-up at a vertex $v \in \Lambda_i$ is descending if and only if it decreases n_i, that is increases $|\Lambda_i|$.* □

An example of a descending blow-up is given in Figure 1. Here d_0 stays constant 4, and n_1 decreases from -1 to -2.

We close this section with some definitions regarding edges in graphs.

Definition 2.5. Let ε be an edge in Γ, with vertices v_1 and v_2. We call ε *horizontal* if $d(p, v_1) = d(p, v_2)$, and *vertical* if $d(p, v_1) \neq d(p, v_2)$. Let ε

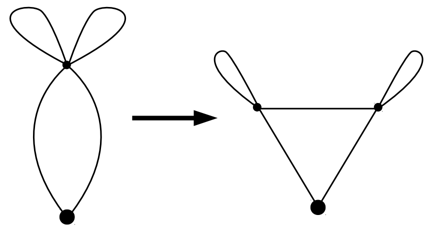

FIGURE 1. A descending blow-up.

be a vertical edge with vertices v_1 and v_2 such that $d(p, v_1) > d(p, v_2)$. We call v_1 the *top* of ε and v_2 the *bottom*. A half-edge can also have a top or a bottom (or neither, if it comes from a horizontal edge). We say that ε is *decisive* if it is the unique vertical edge having v_1 as its top, that is if any minimal length edge path from v_1 to p must begin with ε.

3. Connectivity of the d-down-link

In this section we analyze the d-down-link of Γ. In order for a certain induction to run, it will become necessary to consider (connected) graphs with vertices of valency 1 and 2. It turns out that h does not "work correctly" on such graphs, for instance Lemma 2.3 no longer holds. Therefore in this section we will use Lemma 2.3 as a guide for which forests we want to consider.

Recall that we say F connects vertices in $\Lambda_{D(F)}$ provided that there is an edge path in F between distinct vertices of $\Lambda_{D(F)}$.

Definition 3.1. Let Γ be a connected graph with basepoint p, and with no restriction on the valency of vertices. Let F be a subforest of Γ, with level $D(F)$. We will call F *bad* if it connects vertices in $\Lambda_{D(F)}$, and *good* if it does not.

Thanks to Lemma 2.3, if Γ actually comes from L_n then a forest in Γ is good if and only if it is descending. Let $P(\Gamma)$ be the poset of good forests in Γ, ordered by inclusion, so if Γ comes from L_n then the geometric realization $|P(\Gamma)|$ of $P(\Gamma)$ is the d-down-link of Γ. Let V be the number of vertices in Γ and E the number of edges. In what follows we will suppress the bars indicating geometric realization, so posets themselves will be said to have a homotopy type. Recall that an empty wedge of spheres is a single point.

Proposition 3.2 (Homotopy type of the d-down-link). *$P(\Gamma)$ is homotopy equivalent to a (possibly empty) wedge of spheres of dimension $V - 2$.*

Proof. Our proof is similar to the proof of Proposition 2.2 in [Vog90]. We induct on the number of edges E. We can assume that Γ has no single-edge loops, since they do not affect V or $P(\Gamma)$. We remark that already after this reduction the vertices may have arbitrary valency, so it is important that we are considering "good" forests instead of "descending" forests. Also, if Γ has a separating edge ε then $P(\Gamma)$ is a cone with cone point ε, so without loss of generality Γ has no separating edges.

The base case is $E = 0$, for which $V = 1$ and $P(\Gamma) = \emptyset = S^{V-2}$ as desired.

Now suppose $E > 0$. Choose an edge ε with endpoints v_1, v_2 maximizing the quantity $d(p, v_1) + d(p, v_2)$. In other words, ε is as far as possible from the basepoint; note that $D(\varepsilon)$ is also maximized. Let $P_1(\Gamma) \subseteq P(\Gamma)$ be the poset of all good forests in Γ except the forest just consisting of the edge ε. Also let $P_0(\Gamma) \subseteq P_1(\Gamma)$ be the poset of good forests that do not contain ε.

Claim 1. $P_1(\Gamma) \simeq P_0(\Gamma)$.

Proof of Claim 1. For any $F \in P_1(\Gamma)$, $F - \varepsilon$ is again a good forest by definition, so the poset map $g \colon P_1(\Gamma) \to P_1(\Gamma)$ given by $F \mapsto F - \varepsilon$ is well defined. Here $F - \varepsilon$ is just the forest obtained by removing ε from F. By construction, g is the identity on its image $P_0(\Gamma)$, and $g(F) \leq F$ for all $F \in P_1(\Gamma)$, so g induces a homotopy equivalence between $P_1(\Gamma)$ and $P_0(\Gamma)$ [Qui78, Section 1.3]. □

Now consider the graph $\Gamma - \varepsilon$ obtained by removing ε from Γ. Since ε is not a separating edge, $\Gamma - \varepsilon$ is connected.

Claim 2. $P_0(\Gamma) \cong P(\Gamma - \varepsilon)$.

Proof of Claim 2. Consider the map $\iota \colon P(\Gamma - \varepsilon) \to P_0(\Gamma)$ induced by $\Gamma - \varepsilon \hookrightarrow \Gamma$. Since $D(\varepsilon)$ is maximized and ε is not a separating edge, 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 $\Gamma - \varepsilon$ is good or bad, so ι is an isomorphism. □

Since $\Gamma - \varepsilon$ has $E-1$ edges and V vertices, by induction $P(\Gamma - \varepsilon) \simeq \bigvee S^{V-2}$. Then Claims 1 and 2 tell us that $P_1(\Gamma) \simeq \bigvee S^{V-2}$.

With $P_1(\Gamma)$ in hand, we now ask about $P(\Gamma)$ itself. If ε is horizontal then it is bad, so $P_1(\Gamma) = P(\Gamma)$ and we are done. Assume instead that ε is vertical, hence good, which means $P(\Gamma) = P_1(\Gamma) \cup \mathrm{st}(\varepsilon)$ with $P_1(\Gamma) \cap \mathrm{st}(\varepsilon) = \mathrm{lk}(\varepsilon)$, where link and star are taken in $P(\Gamma)$.

Consider the graph Γ/ε. This has $E - 1$ edges and $V - 1$ vertices, so by induction, $P(\Gamma/\varepsilon) \simeq \bigvee S^{V-3}$. Hence it suffices now to prove the following:

Claim 3. $\mathrm{lk}(\varepsilon) \cong P(\Gamma/\varepsilon)$.

Proof of Claim 3. First note that for a forest $F \neq \varepsilon$ in Γ, F is good if and only if F/ε is, where F/ε is the image of F in Γ/ε. Indeed, if $D(F) < D(\varepsilon)$ then this is trivial; if $D(F) \geq D(\varepsilon)$ then by our choice of ε, $D(F) = D(\varepsilon)$, and it is then evident that F is good if and only if F/ε is. Now consider the map $c \colon \mathrm{lk}(\varepsilon) \to P(\Gamma/\varepsilon)$ sending F to F/ε. This is well-defined by the previous observation. We claim that c is bijective. Let $\Phi \in P(\Gamma/\varepsilon)$. There are precisely two forests in Γ that map to Φ under blowing down ε, one that contains ε and one that does not (this shows that c is injective). Let Φ' be the one that does. If Φ was good then so is Φ', again by the previous observation, so $\Phi' \in \text{lk}(\varepsilon)$. Hence c is an isomorphism. □

This finishes the proof of the Proposition 3.2. □

It will also be convenient to establish one specific case when $P(\Gamma)$ is contractible.

Lemma 3.3. *If Γ has a decisive edge then $P(\Gamma)$ is contractible.*

Proof. The proof is almost the same as the proof of the previous proposition. We again induct on E. If $E = 0$ then Γ does not have any edges, much less any decisive edges, and so the claim is vacuously true. Now assume $E > 0$ and Γ has a decisive edge η. If η has maximum distance to the base point among edges in Γ then it is separating and $P(\Gamma)$ is contractible with η serving as a cone point. Otherwise, let $\varepsilon \neq \eta$ be an edge in Γ that has maximum distance to the basepoint, and define $P_1(\Gamma)$ and $P_0(\Gamma)$ as in the previous proof.

By Claims 1 and 2 in the previous proof, $P_1(\Gamma) \simeq P_0(\Gamma) \cong P(\Gamma - \varepsilon)$. This is contractible by induction since $\Gamma - \varepsilon$ has fewer edges and still contains the decisive edge η. If ε is horizontal, $P(\Gamma) = P_1(\Gamma)$ and we are done, so assume ε is vertical. As in the previous proof, it then suffices to show that $\text{lk}(\varepsilon)$ has the appropriate homotopy type, i.e., is contractible. By Claim 3 in the previous proof, $\text{lk}(\varepsilon) \simeq P(\Gamma/\varepsilon)$. Let η' be the image of η in Γ/ε. Since η is decisive, ε and η have different tops. Since ε is at maximal distance from p, η' is a decisive edge in Γ/ε. Hence $P(\Gamma/\varepsilon)$ is contractible by induction, and we are done. □

4. Connectivity of the d-up-link

We now inspect the d-up-link. We first focus on one vertex at a time. Let $\text{BU}(v)$ be the poset of all blow-ups at the vertex v. We can describe $\text{BU}(v)$ using the combinatorial framework for graph blow-ups described in [CV86] and [Vog90], namely $\text{BU}(v)$ is the poset of *compatible partitions* of the set of incident half-edges, which we now recall.

Compatible partitions. Let $[m] := \{1, 2, \ldots, m\}$, and consider partitions of $[m]$ into two blocks. Denote such a partition by $\alpha = \{a, \bar{a}\}$, where $1 \in a$. Define the *size* of α be

$$s(\alpha) := |\bar{a}|.$$

Recall that distinct partitions $\{a, \bar{a}\}$ and $\{b, \bar{b}\}$ are said to be *compatible* if either $a \subset b$ or $b \subset a$. For $m \geq 3$ let $\Sigma(m)$ denote the simplicial complex of partitions $\alpha = \{a, \bar{a}\}$ of $[m]$ into blocks a and $\bar{a}$ such that a and $\bar{a}$ each have at least two elements, so $2 \leq s(\alpha) \leq m - 2$. That is, the vertices of $\Sigma(m)$ are such partitions, and a j-simplex is given by a collection of $j + 1$ distinct, pairwise compatible partitions. Note that $\Sigma(3) = \emptyset$. Also define a similar complex $\Sigma'(m)$ for $m \geq 2$, identical to $\Sigma(m)$ except that we allow partitions $\alpha = \{a, \bar{a}\}$ with $|\bar{a}| = 1$. We do not allow $|a| = 1$ though, so for example $\Sigma'(2) = \emptyset$.

For $v \neq p$ with $m := \mathrm{val}(v)$, fix a labeling $1, \dots, m$ of the half-edges at v. Then the geometric realization of $\mathrm{BU}(v)$ is isomorphic to the barycentric subdivision of $\Sigma(m)$. In other words, a blow-up at v is encoded by a chain of compatible partitions. A single partition describes an *ideal edge*, i.e., 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 [CV86] and [Vog90] for more details.

Separating blow-ups. Thanks to Corollary 2.4 we know precisely when a blow-up at $v \in \Lambda_i$ is descending, namely when it increases the number of vertices in Λ_i. Hence a blow-up at v is descending if and only if it separates the set of half-edges at v whose top is equal to v. We say that such a blow-up *separates at* v. Let $\mathrm{SBU}(v)$ be the poset of blow-ups at v that separate at v. Note that blow-ups at the basepoint p are never separating, so $\mathrm{SBU}(p) = \emptyset$.

Splitting partitions. We will say that a partition $\alpha = \{a, \bar{a}\}$ of $[m]$ *splits* a subset $S \subseteq [m]$ if $S \not\subseteq a$ and $a \not\subseteq S$. Define the *splitting level* $\ell(\alpha)$ to be the minimum element of $\bar{a}$, i.e., the smallest ℓ such that α splits $[\ell]$. Note that $2 \leq \ell(\alpha) \leq m-1$ for $\alpha \in \Sigma(m)$ and $2 \leq \ell(\alpha) \leq m$ for $\alpha \in \Sigma'(m)$. Let $\Sigma(m,r)$ be the sublevel set of $\Sigma(m)$ spanned by partitions α with $\ell(\alpha) \leq r$, and similarly define $\Sigma'(m,r)$.

The next lemma gives a reformulation of $\Sigma(m,r)$ in terms of graph blow-ups. We assume now that in our fixed labeling of the half-edges of v, those half-edges whose top is v, say there are r of them, are labeled precisely by $1, \dots, r$.

Lemma 4.1 (Separating blow-ups and splitting partitions)**.** *Let $v \neq p$ be a vertex in Γ with m incident half-edges. Let r be the number of half-edges with top v. Then $|\mathrm{SBU}(v)| \simeq \Sigma(m,r)$.*

Proof. The geometric realization $|\mathrm{SBU}(v)|$ contains the barycentric subdivision of $\Sigma(m,r)$ as a subcomplex. Also, any simplex in $|\mathrm{SBU}(v)|$ has at least one vertex in $\Sigma(m,r)$. Hence there is a map $|\mathrm{SBU}(v)| \to |\mathrm{SBU}(v)|$ sending each simplex to its face spanned by vertices in $\Sigma(m,r)$. This induces a deformation retraction from $|\mathrm{SBU}(v)|$ to $\Sigma(m,r)$. □

We now want to calculate the homotopy type of $\Sigma(m,r)$, and perhaps unsurprisingly we will use Morse theory. Consider the height function

$$z(\alpha) := (\ell(\alpha), s(\alpha))$$

on $\Sigma(m)$, with the lexicographic ordering. Since compatible partitions have different sizes, they also have different z-values. Note that $\Sigma(m,r)$ is a sublevel set with respect to z, namely $\Sigma(m,r) = \Sigma(m)^{z \leq (r, m-2)}$. Hence we can analyze the homotopy type of $\Sigma(m,r)$ by looking at descending links in $\Sigma(m)$ with respect to z. We can also think of z as a height function on $\Sigma'(m)$, and before handling $\Sigma(m,r)$ it will be convenient to first calculate the homotopy type of $\Sigma'(m,r)$.

Lemma 4.2. *For any $m \geq 2$ and $2 \leq r \leq m$, $\Sigma'(m,r) \simeq \bigvee S^{m-3}$.*

Proof. We induct on m. Since $\Sigma'(2) = \emptyset$, we already know that $\Sigma'(2,r) = \emptyset = S^{2-3}$ for any r, which handles the base case. Now let $m > 2$ and consider the complex $\Sigma'(m,2)$. This is spanned by partitions $\{a, \bar{a}\}$ in which the set $\{1,2\}$ is split, and so any such a will be $a = \{1\} \cup T$ for T a non-empty subset of $\{3,4,\dots,m\}$. Thus $\Sigma'(m,2)$ is isomorphic to the barycentric subdivision of an $(m-3)$-simplex, and so is contractible.

We now analyze the descending links of partitions with respect to z. Let $\alpha = \{a, \bar{a}\}$ be a partition in $\Sigma'(m,r) \setminus \Sigma'(m,2)$ and set $\ell := \ell(\alpha) > 2$ and $s := s(\alpha)$. A partition $\beta = \{b, \bar{b}\}$ compatible with α is in the z-descending link $\mathrm{lk}{\downarrow_z}(\alpha)$ of α precisely when either $\ell(\beta) < \ell$, or $\ell(\beta) = \ell$ and $a \subsetneq b$. Note that in the first case $b \subseteq a$, so any partition of the first type is compatible with every partition of the second type. Hence the z-descending link of α is a join, of a *d-in-link* and a *d-out-link*. The d-in-link is the full subcomplex of $\mathrm{lk}{\downarrow_z}(\alpha)$ spanned by partitions of the first type, and the d-out-link is spanned by partitions of the second type. See Figure 2 for an example.

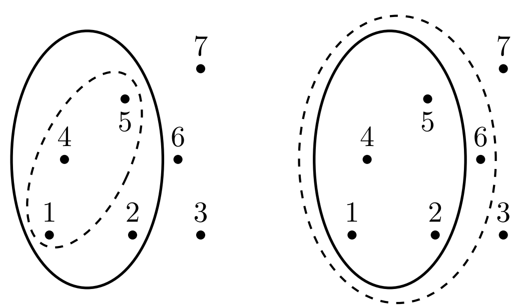

FIGURE 2. A partition in the d-in-link, and one in the d-out-link, of a partition with size $s = 3$ and splitting level $\ell = 3$.

First consider the d-out-link. Partitions $\beta = \{b, \bar{b}\}$ in the d-out-link are characterized by the property that $a \subsetneq b$ and $\ell \in \bar{b}$. Treating a as a single point, this amounts to saying that $a \subsetneq b$ and β splits $\{a, \ell\}$. Hence the d-out-link is isomorphic to $\Sigma'(s+1,2)$. If $s = 1$ this is empty, and if $s > 1$ this is contractible as explained above. In particular if $s > 1$ then $\mathrm{lk}{\downarrow_z}(\alpha)$ is already contractible. Now assume $s = 1$, so the d-out-link is empty and $\mathrm{lk}{\downarrow_z}(\alpha)$ just equals the d-in-link. Then the d-in-link is isomorphic to the complex of partitions of $[m-1]$ that split $[\ell-1]$, and so is given by $\Sigma'(m-1,\ell-1)$. This is $(m-1-3)$-spherical by induction, so we conclude that all descending links are either contractible or $(m-4)$-spherical. Since $\Sigma'(m,2)$ is $(m-3)$-spherical this implies that $\Sigma'(m,r)$ is also $(m-3)$-spherical [BB97, Corollary 2.6]. $\square$

Proposition 4.3. *For any $m \geq 3$ and $2 \leq r \leq m-1$, $\Sigma(m,r) \simeq \bigvee S^{m-4}$.*

Proof. As in the previous proof we induct on m. When $m = 3$ we only consider $r = 2$, and $\Sigma(3,2)$ is empty. Now let $m > 3$ and consider $\Sigma(m,2)$. As with $\Sigma'(m,2)$, $\Sigma(m,2)$ is spanned by partitions $\{a,\bar{a}\}$ in which the set $\{1,2\}$ is split, and so any such a will be $a = \{1\} \cup T$, for T now a *proper* non-empty subset of $\{3,4,\dots,m\}$. (Now we cannot have $T = \{3,4,\dots,m\}$ since the resulting partition would have size 1.) Thus $\Sigma(m,2)$ is the surface of a barycentrically subdivided $(m-3)$-simplex, and so is homeomorphic to S^{m-4}.

Now consider the descending link $\mathrm{lk}{\downarrow_z}(\alpha)$ of $\alpha = \{a,\bar{a}\}$ with $\ell := \ell(\alpha) > 2$ and $s := s(\alpha)$. The descending link decomposes as before as the join of a d-in-link and d-out-link. By the same argument as in the previous proof, the d-out-link is isomorphic to $\Sigma(s+1,2)$, which is homeomorphic to S^{s-3}. The d-in-link is isomorphic to the complex of partitions of $[m-s]$ that split $[\ell-1]$ and have size at least 1. (Since $\bar{a}$ has elements in it, we do have to consider partitions of $[m-s]$ that have size 1 as a partition of $[m-s]$.) So, the d-in-link is isomorphic to $\Sigma'(m-s,\ell-1)$, and hence is homotopy equivalent to $\bigvee S^{m-s-3}$ by the previous lemma. Then $\mathrm{lk}{\downarrow_z}(\alpha)$ is the join of the d-in- and d-out-links, and so is homotopy equivalent to $(\bigvee S^{m-s-3}) * S^{s-3} = \bigvee S^{m-5}$. Since $\Sigma(m,2)$ is $(m-4)$-spherical and the descending links of partitions in $\Sigma(m,r) \setminus \Sigma(m,2)$ are all $(m-5)$-spherical, we conclude that $\Sigma(m,r)$ is $(m-4)$-spherical [BB97, Corollary 2.6]. $\square$

We remark that since $\Sigma(m,m-1) = \Sigma(m)$, we recover the fact that $\Sigma(m)$ is $(m-4)$-spherical, as shown in [Vog90, Theorem 2.4]. Coupling Proposition 4.3 with Lemma 4.1 we see that if there are least two half-edges with top v, then

$$|\mathrm{SBU}(v)| \simeq \bigvee S^{\mathrm{val}(v)-4}.$$

Now let $A := *_{v \neq p} \mathrm{SBU}(v)$, where the join is taken over all vertices $v \neq p$ in Γ. Recall that V is the number of vertices in Γ.

Corollary 4.4. *If Γ has no decisive edges then $|A| \simeq \bigvee S^{d_0(\Gamma)-V}$.*

Proof. Since there are no decisive edges, for any $v \neq p$ we know that there are at least two half-edges at v with top v. Hence $|\mathrm{SBU}(v)| \simeq \bigvee S^{\mathrm{val}(v)-4}$, and so

$$|A| \simeq *_{v \neq p} \bigvee S^{(\mathrm{val}(v)-2)-2} = \bigvee S^{(d_0(\Gamma)-2(V-1))+(V-2)} = \bigvee S^{d_0(\Gamma)-V}. \quad \square$$

Proposition 4.5 (Homotopy type of the d-up-link)**.** *If Γ has no decisive edges then the d-up-link is homotopy equivalent to $|A|$, and hence to $\bigvee S^{d_0(\Gamma)-V}$.*

Proof. For a poset P, define $\underline{P}$ to be $P \sqcup \{\bot\}$, with $\bot$ a formal minimum element. Then $P * Q \cong \underline{P} \times \underline{Q} \setminus \{(\bot,\bot)\}$ for posets P and Q. The relevant example is that

$$A = *_{v\neq p}\,\mathrm{SBU}(v) \cong \prod_{v\neq p} \underline{\mathrm{SBU}}(v) - \{(\bot)_v\} =: Y.$$

Define

$$X := \left\{ f \in \prod_{v\neq p} \underline{\mathrm{BU}}(v) \;\middle|\; \exists v \in \Lambda_{D(f)} \text{ with } f_v \in \mathrm{SBU}(v) \right\}.$$

Here f_v is the blow-up at vertex v in the tuple f, and $D(f)$ is the minimal level such that $f_v \neq \bot$ for some $v \in \Lambda_{D(f)}$. Note that $Y \subseteq X$. Define a map $r\colon X \to X$ by

$$(f_v)_v \mapsto \left(\begin{cases} f_v & \text{for } f_v \in \mathrm{SBU}(v) \\ \bot & \text{for } f_v \notin \mathrm{SBU}(v) \end{cases} \right)_v .$$

Note that r is a poset map that is the identity on its image Y. Also, $r(f) \leq f$ for all $f \in X$, so r induces a homotopy equivalence between $|X|$ and $|Y|$ [Qui78, Section 1.3]. But $|X|$ is precisely the d-up-link of Γ, so the d-up-link is homotopy equivalent to $\bigvee S^{d_0(\Gamma)-V}$ by Corollary 4.4. □

5. Proof of the main results

Corollary 5.1 (Homotopy type of descending links)**.** *For any vertex Γ in L_n, $\mathrm{lk}{\downarrow}(\Gamma)$ is either contractible or homotopy equivalent to $\bigvee S^{d_0(\Gamma)-1}$.*

Proof. If the d-down-link of Γ is contractible, then so is $\mathrm{lk}{\downarrow}(\Gamma)$. If the d-down-link is not contractible, then Γ has no decisive edges (Lemma 3.3). Hence joining the d-up-link and d-down-link yields

$$\left(\bigvee S^{d_0(\Gamma)-V}\right) * \left(\bigvee S^{V-2}\right) \simeq \bigvee S^{d_0(\Gamma)-1}$$

(Propositions 3.2 and 4.5). □

Theorem 5.2 (Degree Theorem)**.** *$L_{n,k}$ is $(k-1)$-connected.*

Proof. For any vertex Γ in $L_n \setminus L_{n,k}$ we have $d_0(\Gamma) > k$, so by the previous corollary, $\mathrm{lk}{\downarrow}(\Gamma)$ is $(k-1)$-connected. Since L_n is contractible and $L_{n,k}$ is a sublevel set of L_n with respect to h (Observation 2.2), $L_{n,k}$ is $(k-1)$-connected by [BB97, Corollary 2.6]. □

References

[BB97] Bestvina, Mladen; Brady, Noel. Morse theory and finiteness properties of groups. *Invent. Math.* **129** (1997), no. 3, 445–470. MR1465330 (98i:20039), Zbl 0888.20021, doi: 10.1007/s002220050168.

[CV86] Culler, Marc; Vogtmann, Karen. Moduli of graphs and automorphisms of free groups. *Invent. Math.* **84** (1986), no. 1, 91–119. MR0830040 (87f:20048), Zbl 0589.20022, doi: 10.1007/BF01388734.

[HV98] Hatcher, Allen; Vogtmann, Karen. Cerf theory for graphs. *J. London Math. Soc.* (2) **58** (1998), no. 3, 633–655. MR1678155 (2000e:20041), Zbl 0922.57001, doi: 10.1112/S0024610798006644.

[McE10] McEwen, R.A. Homological stability for the groups OutP(n,t+1). PhD thesis. University of Virginia, 2010.

[Qui78] Quillen, Daniel. Homotopy properties of the poset of nontrivial p-subgroups of a group. *Adv. in Math.* **28** (1978), no. 2, 101–128. MR0493916 (80k:20049), Zbl 0388.55007, doi: 10.1016/0001-8708(78)90058-0.

[Vog90] Vogtmann, Karen. Local structure of some Out(F_n)-complexes. *Proc. Edinburgh Math. Soc.* (2) **33** (1990), no. 3, 367–379. MR1077791 (92d:57002), Zbl 0694.20021, doi: 10.1017/S0013091500004818.

[Zar14] Zaremsky, Matthew C.B. Rational homological stability for groups of partially symmetric automorphisms of free groups. To appear. *Algebr. Geom. Topol.* (2014). arXiv:1203.4845.

Ruckersville, VA 22968
mcewen.rob@gmail.com

Department of Mathematical Sciences, Binghamton University, Binghamton, NY 13902
zaremsky@math.binghamton.edu

This paper is available via http://nyjm.albany.edu/j/2014/20-13.html.