THE TOPOLOGY AT INFINITY OF $Out(F_n)$

MLADEN BESTVINA AND MARK FEIGHN

ABSTRACT. We construct a bordification of Outer Space analogous to the Borel-Serre bordification of symmetric spaces. As an application, we show that $Out(F_n)$ is $(2n-5)$-connected at infinity and that it is a virtual duality group of dimension $2n-3$.

CONTENTS

Date: October 1999.
Both authors gratefully acknowledge the support by the National Science Foundation.

1. Introduction

A locally compact space X is said to be *r-connected at infinity* provided for every compact set $K \subseteq X$ there exists a compact set $L \subseteq X$ such that $L \supseteq K$ and any map $\alpha : S^d \to X \setminus L$ with $d \leq r$ extends to a map $\tilde{\alpha} : B^{d+1} \to X \setminus K$. For example, X is (-1)-connected at infinity iff it is non-compact, and it is 0-connected at infinity iff it has 1 end.

A discrete group Γ is said to be *r-connected at infinity* if it acts cocompactly, freely, and simplicially on an r-connected simplicial complex X which is r-connected at infinity. If Γ acts cocompactly, freely, and simplicially on another r-connected simplicial complex Y, then Y is necessarily also r-connected at infinity. For example, a group is (-1)-connected at infinity iff it is infinite. The aim of this paper is to prove:

Theorem 1.1. *For $n \geq 2$, $Out(F_n)$ is $(2n-5)$-connected at infinity.*

There are several natural contractible spaces on which $Out(F_n)$ acts, namely Culler-Vogtmann's Outer Space [CV86] and its variants. Outer Space is to be regarded as the analog of a homogeneous space with the action of an arithmetic group or the Teichmüller space with the action of the mapping class group of a compact surface: the action is not quite free (point stabilizers are finite) and it is not quite cocompact. The first problem is not serious at all and is dealt with using the following observation:

Lemma 1.2. *If Γ acts cocompactly, simplicially, and with finite stabilizers of points on an r-connected simplicial complex X which is r-connected at infinity, and if there is a homomorphism $\phi : \Gamma \to \Delta$ to a finite group with torsion-free kernel, then Γ is r-connected at infinity.*

Proof. Δ acts freely on an r-connected finite simplicial complex Y. The diagonal action of Γ on $X \times Y$ (induced by ϕ in the second coordinate) is cocompact and free, and $X \times Y$ is r-connected and r-connected at infinity. $\square$

The group $Out(F_n)$ satisfies the assumptions of Lemma 1.2. Indeed, the natural homomorphism $Out(F_n) \to GL_n(\mathbb{Z}/3)$ has torsion-free kernel [1] [BT68].

There are several ways to resolve the second problem. There is a cocompact equivariant deformation retract of Outer Space, called the Spine of Outer Space [CV86]. Alternatively, we follow the constructions of Siegel [Sie59] for $GL_n(\mathbb{Z})$ and Borel-Serre [BS73] for general arithmetic groups to produce the *Bordification* of Outer Space. The procedure consists of adding ideal points at infinity rather than "cutting off cusps". Roughly speaking, a point of Outer Space is a marked metric graph (of rank n) of total volume 1 in which essential loops

[1] More precisely, they show that $Ker[Out(F_n) \to GL_n(\mathbb{Z})]$ is torsion-free. The desired statement follows then from the standard fact (see [Bro82, p.40]) that $Ker[GL_n(\mathbb{Z}) \to GL_n(\mathbb{Z}/3)]$ is torsion-free.

have positive length. At an ideal point, the metric is allowed to vanish on a subgraph which is not a forest, but a part of the structure at this ideal point is the "secondary" metric of volume 1 defined only on the core of the subgraph on which the primary metric vanishes. Similarly, the secondary metric is allowed to vanish non-trivially and there may be a "tertiary" metric, etc. Thus an ideal point is represented by a hierarchy of metrics defined on a decreasing sequence of subgraphs of a rank n marked metric graph. For example, imagine a sequence of metrics ℓ_i, $i = 2, 3, \cdots$ on the rose with loops a, b, c such that $\ell_i(a) = 1 - 1/i$ and $\ell_i(b) = \ell_i(c) = 1/2i$. The sequence converges to an ideal point with primary metric ℓ given by $\ell(a) = 1$ and $\ell(b) = \ell(c) = 0$ and the secondary metric ℓ' defined only on b and c by $\ell'(b) = \ell'(c) = 1/2$. The secondary metric is obtained by rescaling the restriction of ℓ_i to $b \cup c$ (the subgraph on which ℓ vanishes) and taking the limit. For another example, the sequence of metrics ℓ_i defined by $\ell_i(a) = 1 - 1/i - 1/i^2$, $\ell_i(b) = 1/i$, $\ell_i(c) = 1/i^2$ converges to the ideal point determined by a primary metric that assigns 1 to a and 0 to b and c, a secondary metric defined on $b \cup c$ that assigns 1 to b and 0 to c, and a tertiary metric defined on c and assigning 1 to it.

The reason for working with bordification instead with the spine is that we analyze the topology of the underlying space using a Morse function, and a natural Morse function would have too many (canceling) critical points at the frontier of the cusps that were cut off. We succeed in understanding the topology at infinity because of the lucky event that a certain naturally defined Morse-like function on the bordified Outer Space has all critical points of the same index, so there is no need to worry about possible cancellations.

A model Morse theory argument for investigating the topology at infinity is:

Lemma 1.3. *Suppose M^n is an open manifold and $f : M \to \mathbb{R}_+ = [0, \infty)$ is a proper Morse function such that at all critical points the ascending sphere has dimension $\geq r + 1$ (i.e. the index is $\leq n - r - 2$). If M is r-connected, then it is r-connected at infinity.*

Outer Space is not a manifold and we develop a Morse theory that applies to (bordified) Outer Space in the style of [BB97]. Roughly speaking, for any conjugacy class α in F_n the map "length of α" is a good candidate for a Morse function on the bordified Outer Space. Unfortunately, such a function does not have isolated critical points.

To remedy this, we work with an infinite collection (an "ample sequence") of conjugacy classes $\alpha_1, \alpha_2, \cdots$ and consider a function that to a marked graph assigns the tuple of lengths of the α_i's. Worse, the topology near the ideal points is not captured by ordinary length, and we are forced to consider higher order lengths as well. The target of our "Morse function" is not $\mathbb{R}_+$ but an infinite product of $\mathbb{R}_+$'s, ordered lexicographically, so it makes sense to talk about sub- and super-level sets.

The analysis of the topology of "ascending links" is inductive, and we are forced to consider more general bordified Outer Spaces, allowing the underlying graphs to be disconnected and to have distinguished points.

As a corollary of our analysis of the topology of bordified Outer Space we obtain:

Theorem 1.4. *For $n \geq 2$, $Out(F_n)$ is a virtual duality group of dimension $2n-3$.*

It was shown by Culler and Vogtmann [CV86] that the virtual cohomological dimension (see [Bro82]) of $Out(F_n)$ is $2n-3$. The notion of a duality group was introduced by Bieri and Eckmann [BE73] and it generalizes the concept of Poincaré duality groups. One definition (see [Bro82, p.220]) is that a group Γ that acts freely and cocompactly on a contractible cell complex X (by permuting cells) is a duality group of dimension n if $H^i(\Gamma;\mathbb{Z}\Gamma)$ vanishes when $i \neq n$ and is torsion-free (as an abelian group) when $i = n$. The cohomology group $H^i(\Gamma;\mathbb{Z}\Gamma)$ is isomorphic (see [Bro82, p.209]) to the compactly supported cohomology $H^i_c(X;\mathbb{Z})$. From the Culler-Vogtmann result we have $H^i(\Gamma;\mathbb{Z}\Gamma) = 0$ for $i > 2n-3$, and it is a theorem of Geoghegan-Mihalik [GM85] that if X as above is k-connected at infinity, then $H^i_c(X;\mathbb{Z})$ vanishes for $i \leq k+1$ and is torsion-free for $i = k+2$. Applying this result with $k = 2n-5$ gives a proof of Theorem 1.4. In fact, our analysis gives a concrete description of $H^{2n-3}_c(X;\mathbb{Z})$ from which it is clear that it is torsion-free (and moreover a free abelian group).

When $n = 2$ we don't say anything new. Vogtmann [Vog95] showed that $Out(F_n)$ has one end for $n \geq 3$ and that it is 1-connected at infinity for $n \geq 5$, and she also showed [Vog90] that Outer Space and its variants are Cohen-Macauley complexes, thus providing evidence for Theorem 1.4. Rickert [Ric] proved that $Out(F_4)$ is 1-connected at infinity. The statement that $Out(F_3)$ is 1-connected at infinity is new. For arithmetic groups, the equivalents of Theorems 1.1 and 1.4 were proved by Borel and Serre [BS73], and for mapping class groups of compact surfaces by Harer [Har86].

The layout of the paper is as follows. In Section 2 we give the definition of Outer Space, its bordification, and some variants, and analyze the local topology of these spaces. In particular, the bordified Outer Space can be thought of as a cell-complex with each cell a convex cell in Euclidean space, much like Siegel's bordification of $GL_n(\mathbb{R})/O_n(\mathbb{R})$. As with many other things in this paper, this is not technically how things go, and we only have large charts with convex polyhedral image. In fact, the "cells" are only that modulo the Poincaré Conjecture.

In Section 3 we define our Morse function and show that its restriction to each "cell" is standard (empty or contractible super-level sets). An immediate consequence of this is that each critical point is a vertex (so the critical points are isolated).

In Section 4 we study the change in the topology of super-level sets that occurs after passing through a vertex. In the language of our Morse theory, we study the topology of ascending links, and show that they are homotopy equivalent to the wedge of spheres of dimension $2n-4$.

Section 5 deals with the last remaining problem – that our Morse function didn't go to $\mathbb{R}_+$ and wasn't continuous (much less proper). Theorem 5.3 is the model Morse theory statement that replaces Lemma 1.3 in this setting. Theorem 1.1 is restated (and proved) as Theorem 5.6. The trick consists of replacing Bordified Outer Space by a simplicial model, and replacing the Morse function by another one which is proper and has ascending links homotopy equivalent to the ascending links of the original function.

We wish to thank Noel Brady, Vincent Guirardel, Gilbert Levitt, and the referee for pointing out some mistakes in earlier versions of the manuscript.

We end this introductory section by reviewing some concepts we need in this paper.

1.1. **Absolute Neighborhood Retracts.** All spaces will be separable and will admit a complete metric. (The targets of our Morse functions will be totally ordered sets, but we will not topologize them.) A space X is said to be an ANR if whenever embedded in a locally convex topological vector space as a closed subset, it admits a neighborhood retraction. The standard reference is [Bor67]. A finite-dimensional space is an ANR iff it is locally contractible. If $A, B \subseteq X$ are closed subsets such that $A, B, A \cap B$ are ANR's, then $A \cup B$ is an ANR. A map between connected ANR's that induces an isomorphism between homotopy groups is a homotopy equivalence. A contractible ANR is called an *absolute retract* (AR).

1.2. **Semi-algebraic sets.** A subset V of $\mathbb{R}^n$ is *semi-algebraic* if it is the union of finitely many subsets each of which is cut out from $\mathbb{R}^n$ by a finite number of polynomial equations $f(x) = 0$ and inequalities $f(x) > 0$. Every semi-algebraic set is an ANR (i.e. locally contractible). For semi-analytic sets this was first proved by Lojasiewicz in [Loj64]. The semi-algebraic case is easier, and a simple proof is presented by Hironaka in [Hir75]. Moreover, for any $x \in V$ there is an analytic change of coordinates h from a neighborhood U of x in $\mathbb{R}^n$ to $\mathbb{R}^n$ such that $h(x) = 0$ and $h(U \cap V)$ is a cone. In particular, points of V have cone neighborhoods in V. In fact, semi-algebraic sets are triangulable.

1.3. Z**-sets.** A closed subset $A \subset X$ of an ANR X is a *Z-set* if for any open set $U \subseteq X$ inclusion $U \setminus A \hookrightarrow U$ is a homotopy equivalence. Equivalently, there is a homotopy $H : X \times [0,1] \to X$ with $H(x,0) = x$ and $H(x,t) \in X \setminus A$ for $t > 0$. For example, any closed subset of the boundary of a manifold is a Z-set in the manifold. The standard reference is [Hen75]. We find it useful to slightly extend the language of Z-sets. We say that a closed subset A of an ANR X

is *locally* Z *at* $a \in A$ if there is a neighborhood V of a in X and a homotopy $H : V \times [0,1] \to X$ such that $H(x,0) = x$ and $H(x,t) \in X \setminus A$ for $t > 0$. It can be shown with the methods of [Hen75] that if a closed set is locally Z at each point then it is a Z-set.

A *Z-sandwich* is a triple $X \supseteq Y \supseteq Z$ such that X is an ANR and $X \setminus Z$ is a Z-set in X. It follows that inclusions $Z \hookrightarrow Y \hookrightarrow X$ are homotopy equivalences.

2. The Bordification of Outer Space

Definitions 2.1. A finite (possibly disconnected) graph is said to be a *core graph* if it does not have vertices of valence 0 or 1, and none of its components has a separating edge. Equivalently, a core graph is one that can be represented as the union of finitely many circles. Note that every finite graph H which is not a forest contains a unique nonempty maximal subgraph which is a core graph. We call this subgraph the *core* of H, denoted $core(H)$. By $\mathcal{C}(G)$ we denote the collection of all nonempty subgraphs of a graph G which are core graphs.

An $\mathbb{R}$-*metric* on a graph G is a function ℓ that assigns a nonnegative number $\ell(e)$ to each edge e of G. We refer to $\ell(e)$ as the *length* of e (with respect to ℓ). The *volume* of G with respect to ℓ is the sum of the lengths of all edges of G. More generally, the volume of a subgraph is the sum of the lengths of edges in the subgraph. When G is a finite graph, we denote by Σ_G the simplex of all $\mathbb{R}$-metrics on G of volume 1. An $\mathbb{R}$-metric on G is *degenerate* if a nonempty core subgraph has volume 0; otherwise it is *non-degenerate*. We denote by $ND(G)$ the subset of Σ_G consisting of non-degenerate metrics.

2.1. **The Bordification of** Σ_G**.** We now construct a space $\hat{\Sigma}_G$. It will be homeomorphic to a closed cell (at least modulo the Poincaré Conjecture) and it will naturally map onto Σ_G. Further, this map will be a homeomorphism over the non-degenerate $\mathbb{R}$-metrics. Thus $\hat{\Sigma}_G$ can be viewed as a more refined compactification of the space of non-degenerate $\mathbb{R}$-metrics on G of volume 1.

We assume that G is a core graph. For every core subgraph $C \in \mathcal{C}(G)$ define the *coordinate restriction map* $r_C : Int\ \Sigma_G \to Int\ \Sigma_C \subset \Sigma_C$ by restricting the given $\mathbb{R}$-metric to C and then rescaling so the volume is 1. The product map

$$\prod_{C \in \mathcal{C}(G)} r_C : Int\ \Sigma_G \to \prod_{C \in \mathcal{C}(G)} \Sigma_C$$

is an embedding, as one of the coordinates is inclusion $Int\ \Sigma_G \hookrightarrow \Sigma_G$. Define $\hat{\Sigma}_G$ as the closure of the image of this product map. By definition, $\hat{\Sigma}_G$ is a compactification of $Int\ \Sigma_G$. The natural map

$$\pi : \hat{\Sigma}_G \to \Sigma_G$$

is defined by restricting the projection map. Equivalently, the definition can be made with $ND(G)$ replacing $Int\ \Sigma_G$. It will sometimes be convenient to take this point of view.

Example 1. When G is the circle with n vertices and edges, $\hat{\Sigma}_G = \Sigma_G$ is the $(n-1)$-simplex.

Example 2. When G is the bouquet of two circles, $\hat{\Sigma}_G = \Sigma_G$ is an arc.

Example 3. When G is the theta graph (the bipartite graph with 2 vertices and 3 edges), $\hat{\Sigma}_G$ is a hexagon.

Example 4. When G is the bouquet of three circles, $\hat{\Sigma}_G$ is a hexagon, and $\pi : \hat{\Sigma}_G \to \Sigma_G$ collapses every other face to a point. The hexagon is depicted in the figure. A representative point in each open face is indicated by a decorated bouquet. A point is determined by the $\mathbb{R}$-metrics associated to each of its core subgraphs. In the figure, solid edges have non-zero ℓ_G-length. Dashed edges have non-zero length with respect to the metric on the core graph of non-solid edges. Finally, dotted edges have non-zero length with respect to the metric on the core graph of dotted edges. The metrics on the other core graphs can be reconstructed from these. This will be made more precise in the remainder of this section starting with Definition 2.7. For example, if a is the edge of G displayed with vertical symmetry, and b and c are the other edges in counter-clockwise order then at the topmost point of the hexagon, $\ell_G(a) = 1$, $\ell_G(b) = 0$, and $\ell_G(c) = 0$. If H_1 is the core graph $b \cup c$ then at this same point $\ell_{H_1}(b) = 1$ and $\ell_{H_1}(c) = 0$. If H_2 is the core graph with only edge c then $\ell_{H_2}(c) = 1$. At the point immediately to the right of this one $\ell_{H_1}(b) = \ell_{H_1}(c) = 1/2$. At the central point $\ell_G(a) = \ell_G(b) = \ell_G(c) = 1/3$. This example will be expanded as we progress through the paper.

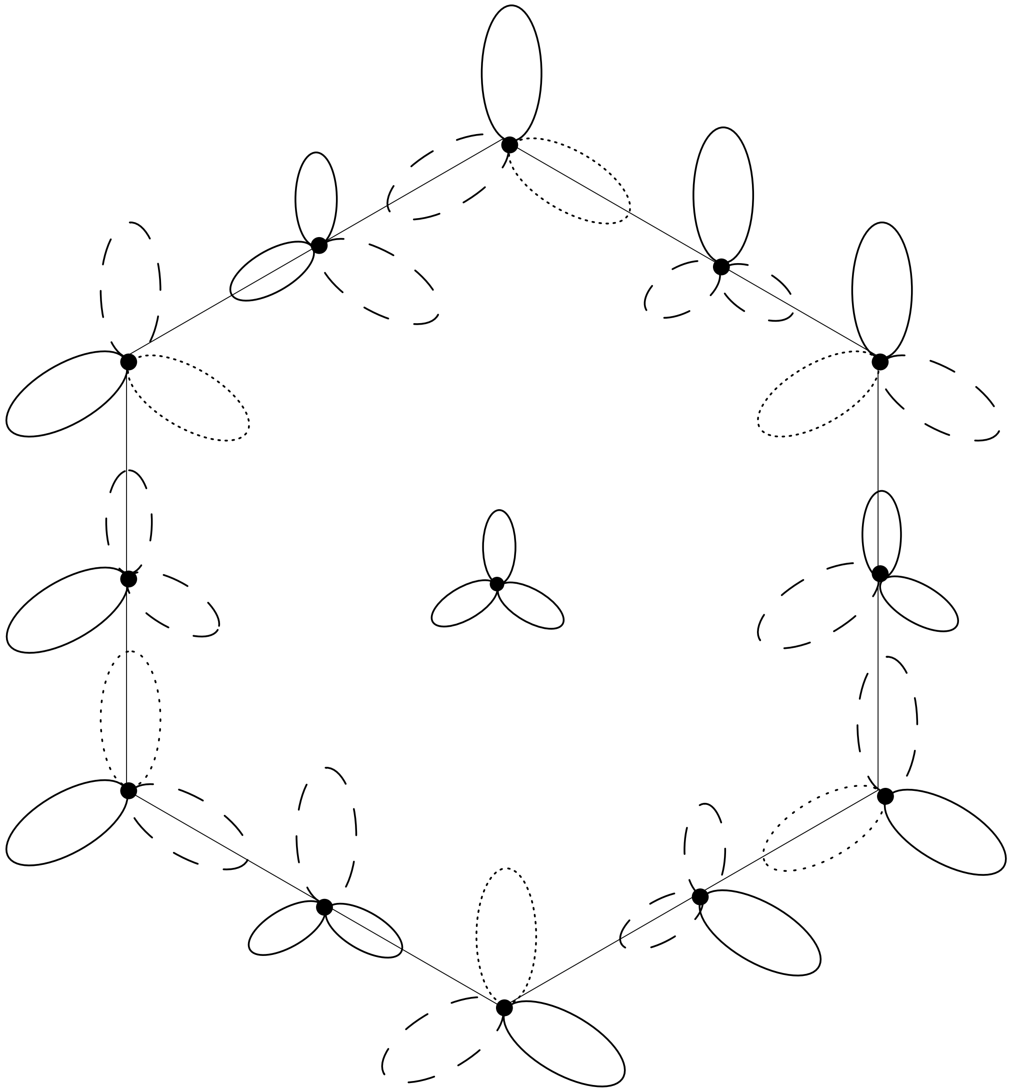

Figure 1

Remark 2.2. An interesting variation of the above construction results if one drops the requirement that core graphs have no separating edges. The reader may verify that in that case the bordification $\hat{\Sigma}_G$ for the case when G consists of two circles joined by an arc is a quadrilateral and the projection collapses one face to a point. We choose to include the non-separation requirement as the resulting space is then somewhat more efficient and our 'Morse function' will have fewer critical points.

A point in the product $\prod_{C\in\mathcal{C}(G)}\Sigma_C$ assigns an $\mathbb{R}$-metric ℓ_C of volume 1 to every nonempty core subgraph C of G. We will often refer to ℓ_G as the *topmost metric.*

Lemma 2.3. *A point $\{\ell_C\}_{C\in\mathcal{C}(G)}$ belongs to $\hat{\Sigma}_G$ iff the following* compatibility condition *holds.*
$(*)$ If $C_1 \subset C_2$ are two core subgraphs and ℓ_{C_2} does not vanish on all edges of C_1, then ℓ_{C_1} is the rescaled restriction of ℓ_{C_2}.

Proof. Condition $(*)$ is a closed condition and it clearly holds on the image of $Int\ \Sigma_G$ (and also $ND(G)$). Thus $(*)$ holds for every point in $\hat{\Sigma}_G$.

Conversely, assume $\{\ell_C\}_{C\in\mathcal{C}(G)}$ satisfies $(*)$. Define a decreasing sequence of core subgraphs of G as follows. Set $G_0 = G$. If the collection of edges of ℓ_G-length 0 forms a forest, then stop. Otherwise, let G_1 be the maximal nonempty core graph on which ℓ_G vanishes. If the collection of edges on which ℓ_{G_1} vanishes forms a forest, stop; otherwise define G_2 as the core of the subgraph of G_1 on which ℓ_{G_1} vanishes. Continuing in this fashion, we obtain a finite sequence $G = G_0 \supset G_1 \supset \cdots \supset G_m$ of core graphs, and ℓ_{G_m} vanishes on a forest.

Claim 2.4. In the presence of the compatibility condition $(*)$, metrics ℓ_{G_i} uniquely determine all ℓ_C for $C \in \mathcal{C}(G)$.

[Indeed, let i be the largest index such that $C \subseteq G_i$. Then ℓ_{G_i} doesn't vanish on every edge of C (or else $C \subseteq G_{i+1}$), and hence ℓ_C is determined by ℓ_{G_i} via $(*)$.]

For $\epsilon > 0$ define the $\mathbb{R}$-metric ℓ_ϵ on G by

$$\ell_\epsilon(e) = \epsilon^j \ell_{G_j}(e)$$

where j is the largest index such that $e \subset G_j$. After rescaling, ℓ_ϵ represents a point in $ND(G)$, and as $\epsilon \to 0$, the corresponding point converges to $\{\ell_C\}_{C\in\mathcal{C}(G)} \in \prod_{C\in\mathcal{C}(G)}\Sigma_C$. The last assertion is immediate on the coordinates of the form Σ_{G_i} and follows from Claim 2.4 for all others. $\square$

Corollary 2.5. *$\hat{\Sigma}_G$ is a semi-algebraic set, hence an ANR.* $\square$

Proposition 2.6. *$\hat{\Sigma}_G$ is contractible and $\pi : \hat{\Sigma}_G \to \Sigma_G$ has point inverses which are AR's (in particular, contractible).*

Proof. We argue by induction on the number of edges of G. First observe that the second statement implies the first, since proper maps with contractible point inverses between ANR's are homotopy equivalences [Sma57] (see also [Lac77]). A point inverse of $\pi : \hat{\Sigma}_G \to \Sigma_G$ consists of those points in $\hat{\Sigma}_G$ that have a fixed ℓ_G. In view of Claim 2.4, the point inverse can be identified with $\hat{\Sigma}_H$ where H

is the core of the subgraph of G on which ℓ_G vanishes (if H is empty, the point inverse is a single point). By induction, it follows that the point inverses are AR's. □

Definition 2.7. The *hierarchy of scales* associated to a point $\{\ell_C\}_{C\in\mathcal{C}(G)} \in \hat{\Sigma}_G$ is the sequence of core graphs $G = G_0 \supset G_1 \supset \cdots \supset G_m$ constructed in the proof of Lemma 2.3.

Intuitively, a point of $\hat{\Sigma}_G$ represents the limit of a sequence of non-degenerate $\mathbb{R}$-metrics on G of volume 1, where we keep track of ratios of lengths on different scales.

It will also be necessary to understand a neighborhood of the point inverse of a vertex under the map π above. Let v be a vertex of Σ_G. For $0 < \epsilon < 1$ let Q_ϵ (resp. $Q_{\leq\epsilon}$) be the set of points in the simplex Σ_G whose barycentric coordinate associated to v is $1-\epsilon$ (resp. $\geq 1-\epsilon$), thus Q_ϵ is obtained from the face opposite to v by a homotethy with scaling factor ϵ centered at v (and $Q_{\leq\epsilon}$ is the closed neighborhood of v determined by Q_ϵ). Let $\hat{Q}_\epsilon = \pi^{-1}(Q_\epsilon)$ and similarly $\hat{Q}_{\leq\epsilon} = \pi^{-1}(Q_{\leq\epsilon})$. Fix some $\epsilon \in (0,1)$ and define a path $\alpha_x(s)$ that joins a point $x \in \hat{Q}_\epsilon$ to a point in $\pi^{-1}(v)$. For a point $y \in \hat{\Sigma}_G$ we denote by ℓ_C^y the metric on the core graph $C \subset G$ associated with y. Let e_0 be the edge of G of length 1 under the metric given by v. We then define:

$$\ell_C^{\alpha_x(s)} = \ell_C^x \quad \text{if } e_0 \not\subseteq C$$

and otherwise, when $e_0 \subseteq C$,

$$\begin{aligned} \ell_C^{\alpha_x(s)}(e_0) &= 1 - \epsilon + s\epsilon \\ \ell_C^{\alpha_x(s)}(e) &= (1-s)\ell_C^x(e), \quad \text{if } e \neq e_0. \end{aligned} \tag{1}$$

Of course, in view of the compatibility condition, we could have stated the definition for $\ell_G^{\alpha_x(s)}$ only, instead of the general core graph C with $e_0 \subseteq C$. The path α_x maps by π to the straight line segment joining $\pi(x)$ to v.

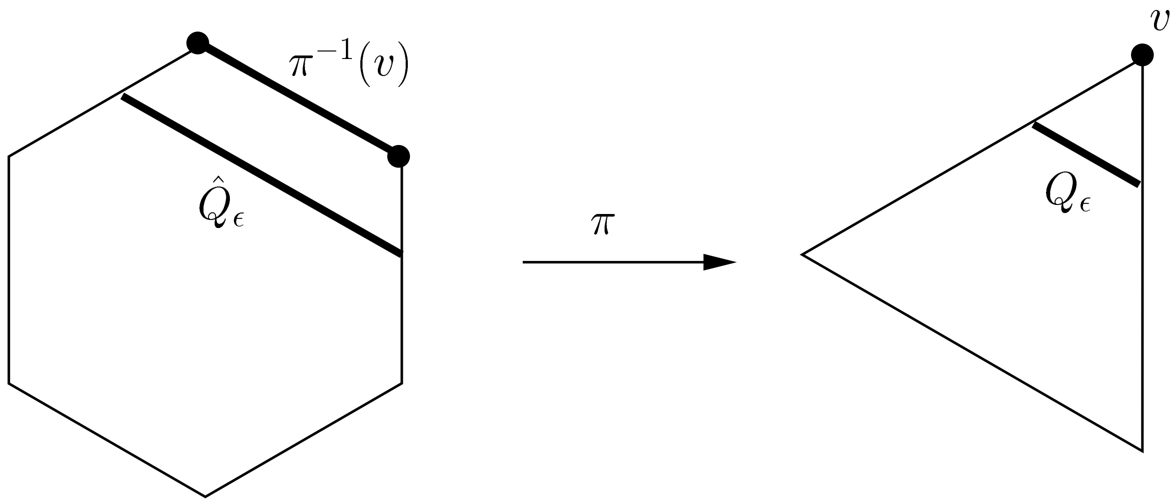

Proposition 2.8. *The paths α_x define a mapping cylinder structure on $\hat{Q}_{\leq\epsilon}$. The point inverses of the associated map $\rho : \hat{Q}_\epsilon \to \pi^{-1}(v)$ are simplices.*

Proof. Each α_x is continuous in s since by definition each ℓ_C changes continuously with s, and for the same reason the paths α_x vary continuously in x. If $x \neq x'$, the paths α_x and $\alpha_{x'}$ can intersect only at the terminal endpoint, since x can be recovered from $\alpha_x(s)$ for every $s < 1$. Each path is contained in $\hat{Q}_{\leq\epsilon}$ and each point of $\hat{Q}_{\leq\epsilon}$ is on some such path. Consequently, $\hat{Q}_{\leq\epsilon}$ has the structure of the mapping cylinder of a map $\rho : \hat{Q}_\epsilon \to \pi^{-1}(v)$. To describe the point inverses of ρ, note that $\ell^x_C = \ell^{\rho(x)}_C$ when $e_0 \not\subseteq C$, so the only indeterminacy is ℓ_G. In particular, π is injective on the point inverses of ρ. The only restrictions on ℓ_G come from the volume normalization and the compatibility conditions. Since $\ell^x_G(e_0) = 1 - \epsilon$ for every $x \in \hat{Q}_\epsilon$, ℓ^x_G is uniquely determined by the ℓ_G-volume of $H = core(G \setminus e_0)$ and the ℓ_G-lengths of the edges in $G \setminus (e_0 \cup H)$, subject to the condition that the sum of these numbers is ϵ. Thus the point inverses are simplices. $\square$

2.2. **Local structure.** A core subgraph C_0 fails to occur in the hierarchy of scales for $\{\ell_C\}_{C\in\mathcal{C}(G)}$ if and only if there is a core subgraph C_1 that properly contains C_0 such that ℓ_{C_1} doesn't vanish on C_0. In particular, 'not occurring in the hierarchy of scales' is an open condition.

If S is a collection of edges in G, we denote by Δ_S the simplex $\{(t_s)_{s\in S} | t_s \geq 0, \sum t_s \leq 1\}$, and by Σ_S the simplex $\{(t_s)_{s\in S} | t_s \geq 0, \sum t_s = 1\}$.

Definition 2.9. Let $\underline{F} = (G = C_0 \supset C_1 \supset \cdots \supset C_q)$ be a sequence of nonempty core graphs. It is sometimes convenient to put $C_{q+1} = \emptyset$. The *chart* $U_{\underline{F}}$ corresponding to $\underline{F}$ is the open set of points of $\hat{\Sigma}_G$ with the property that each graph in the hierarchy of scales is one of the C_i's. The *chart map*

$$\Phi = \Phi_{\underline{F}} : U = U_{\underline{F}} \to \Delta = \Delta_{\underline{F}} = \Delta_{C_0 \setminus C_1} \times \Delta_{C_1 \setminus C_2} \times \cdots \times \Delta_{C_{q-1} \setminus C_q} \times \Sigma_{C_q}$$

is defined by

$$\{\ell_C\}_{C \in \mathcal{C}(G)} \mapsto (\ell_{C_0}|_{C_0 \setminus C_1}, \ell_{C_1}|_{C_1 \setminus C_2}, \cdots, \ell_{C_{q-1}}|_{C_{q-1} \setminus C_q}, \ell_{C_q}).$$

We view Δ as a product polytope. In particular, by a *face* of Δ we mean the product of faces of the coordinate simplices.

The same formula defines an extension $\hat{\Phi} : \hat{\Sigma}_G \to \Delta$ of Φ.

Proposition 2.10. *$\hat{\Phi}$ has AR's as its point inverses. Moreover, a single point $\{\ell_C\} \in \hat{\Sigma}_G$ is a point inverse if each graph in the hierarchy of scales for $\{\ell_C\}$ is one of C_i's.*

Proof. We induct on the number of edges of G. Fix a point $(x_0, x_1, \cdots, x_q) \in \Delta_{C_0 \setminus C_1} \times \Delta_{C_1 \setminus C_2} \times \cdots \times \Sigma_{C_q}$. We first note that all points in $\hat{\Phi}^{-1}(x_0, x_1, \cdots, x_q)$ have the same ℓ_G. Indeed, x_0 determines ℓ_G on edges in $C_0 \setminus C_1$. If $x_0 \in \Sigma_{C_0 \setminus C_1}$, then ℓ_G vanishes on C_1, and otherwise x_1 determines ℓ_G on edges in $C_1 \setminus C_2$, by rescaling. If $x_1 \in \Sigma_{C_1 \setminus C_2}$ then ℓ_G vanishes on C_2, and otherwise x_2 determines ℓ_G on $C_2 \setminus C_3$, etc.

Denote by H the core of the subgraph of G on which ℓ_G vanishes (if ℓ_G vanishes on a forest, there is nothing left to prove). If $H = C_i$ for some i, then $x_0, x_1, \cdots, x_{i-1}$ are determined by ℓ_G by rescaling and carry no additional information, so we can identify $\hat{\Phi}^{-1}(x_0, x_1, \cdots, x_q)$ with $\hat{\Phi}^{-1}_{\underline{F'}}(x_i, x_{i+1}, \cdots, x_q)$, where $\hat{\Phi}_{\underline{F'}} : \hat{\Sigma}_{C_i} \to \Delta_{C_i \setminus C_{i+1}} \times \Delta_{C_{i+1} \setminus C_{i+2}} \times \cdots \times \Sigma_{C_q}$ is the similarly defined map and $\underline{F'} = (H = C_i \supset C_{i+1} \supset \cdots \supset C_q)$. The proposition now follows from the inductive assumption applied to $\hat{\Phi}_{\underline{F'}}$.

Now suppose $H \neq C_i$ for any i. In this situation we need to argue that the point inverse is an AR. Note that if $C_i \not\subseteq H$, then ℓ_G determines x_i by rescaling. Let i be the smallest index with $H \supset C_i$ (if there is no such i then the point inverse can be identified $\hat{\Sigma}_H$ and we are done). Then the point inverse under consideration can be identified with the inverse of $(x_i, x_{i+1}, \cdots, x_q)$ under the composition

$$\hat{\Sigma}_H \overset{\hat{\Phi}_{\underline{F'}}}{\to} \Delta_{H \setminus C_i} \times \Delta_{C_i \setminus C_{i+1}} \times \cdots \times \Sigma_{C_q} \to \Delta_{C_i \setminus C_{i+1}} \times \cdots \times \Sigma_{C_q}$$

with the second map the projection and $\underline{F'} = (H \supset C_i \supset C_{i+1} \supset \cdots \supset C_q)$. Thus this point inverse maps onto the simplex $\Delta_{H \setminus C_i}$ with AR point inverses, and it is an ANR (it is a semi-algebraic set), so by [Sma57] it is an AR. $\square$

Remark 2.11. There are other degenerate point inverses, e.g. when H is a single edge (loop) that does not equal to any C_i.

Proposition 2.12. 1. *$\Phi : U \to \Delta$ is 1-1.*
2. *Φ is a homeomorphism onto its image.*
3. *The image of Φ is the complement of the union of a collection of faces of Δ.*
4. *$\hat{\Sigma}_G$ is a compact contractible manifold whose interior is $Int\ \Sigma_G$.*

Proof. It follows from Proposition 2.10 that Φ is injective and that for open $V \subset U$, $\hat{\Sigma}_G = \hat{\Phi}(V) \coprod \hat{\Phi}(\hat{\Sigma}_G \setminus V)$. In particular, being the complement of a compact set, $\Phi(V) = \hat{\Phi}(V)$ is open and Φ is an open map. So, Φ is a homeomorphism onto its image. We need to argue that $\Phi(U)$ is the complement of the union of some faces of Δ. Let $0 = n_0 < n_1 < n_2 < \cdots < n_k \leq q$ and consider the subset X of $\hat{\Sigma}_G$ consisting of those points whose hierarchy of scales is $G \supset C_{n_1} \supset C_{n_2} \supset \cdots \supset C_{n_k}$. We now argue that $\Phi(X)$ is the union of some faces. For each point in X denote by D_i the subgraph of G consisting of edges on which $\ell_{C_{n_i}}$ vanishes, so that $core(D_i) = C_{n_{i+1}}$. Thus X can be viewed as the union over all different possibilities for the D_i's. The Φ-image of the subset of X with the fixed choice of the D_i's is the union of faces of Δ given by the equations that the sum of the coordinates of $x_{n_i} \in \Delta_{C_{n_i-1} \setminus C_{n_i}}$ is 1 for all i, and that the coordinates of $x_{n_i}, x_{n_i+1}, \cdots, x_{n_{i+1}-1}$ corresponding to the edges of $D_i \setminus C_{n_{i+1}}$ vanish. To prove (4) observe that $Int\ \Sigma_G \subset U$, that $\Phi(Int\ \Sigma_G) = Int\ \Delta$ and that the charts cover $\hat{\Sigma}_G$. $\square$

Remark 2.13. One can show that $\hat{\Sigma}_G$ has the structure of a real analytic manifold with corners (see [BS73]). Indeed, note that each $\{\ell_C\}_{C \in \mathcal{C}(G)}$ is contained in a unique minimal chart, given by the hierarchy of scales of $\{\ell_C\}$. To study the transition maps, it suffices to consider the case of inclusion $U_{\underline{F'}} \subset U_{\underline{F}}$, where $\underline{F'}$ is obtained from $\underline{F}$ by omitting one of the core graphs, say C_i. For concreteness assume $i \neq m$. Define the following map I on the complement of a face of $\Delta_{\underline{F'}}$ into $\Delta_{\underline{F}}$. The map is the projection to the coordinates other than $\Delta_{C_{i-1} \setminus C_i} \times \Delta_{C_i \setminus C_{i+1}}$. If the $\Delta_{C_{i-1} \setminus C_{i+1}}$-coordinate of $x \in \Delta_{\underline{F'}}$ is $(t_1, t_2, \cdots, t_k, s_1, s_2, \cdots, s_l)$ (with k edges in $C_{i-1} \setminus C_i$ and l edges in $C_i \setminus C_{i+1}$), then the $\Delta_{C_{i-1} \setminus C_i} \times \Delta_{C_i \setminus C_{i+1}}$-coordinates of $I(x)$ are

$$((t_1, \cdots, t_k), (\frac{s_1}{1 - t_1 - t_2 - \cdots - t_k}, \cdots, \frac{s_l}{1 - t_1 - t_2 - \cdots - t_k})).$$

The inverse map J is defined on all of $\Delta_{\underline{F}}$, and the formula in interesting coordinates is given by

$$((t_1, \cdots, t_k), (s_1, \cdots, s_l)) \mapsto (t_1, \cdots, t_k, (1-t_1-\cdots-t_k)s_1, \cdots, (1-t_1-\cdots-t_k)s_l)$$

Remark 2.14. It follows from the Poincaré Conjecture that $\hat{\Sigma}_G$ is homeomorphic to the closed ball when its dimension is not 4.

Question. Is $\hat{\Sigma}_G$ a PL ball in all dimensions (with the PL structure induced from the analytic manifold with corners structure described above)?

As we are interested in homotopy-theoretic issues, we will not pursue this.

2.3. **Cell decomposition.** Fix a descending sequence

$$\underline{K} = (G = K_0 \supset K_1 \supset K_2 \supset \cdots \supset K_{m+1})$$

of subgraphs of G such that K_{i+1} is a proper subgraph of the core of K_i (K_{m+1} is allowed to be empty). The *associated open cell* is the subset $W(\underline{K})$ of $\hat{\Sigma}_G$ consisting of those points whose hierarchy of scales (G_i) consists of the cores $G_i = core(K_i)$ $(i \leq m)$, and further, the metric ℓ_{G_i} vanishes exactly on K_{i+1}. Each $W(\underline{K})$ is visibly an open cell, canonically identified with the interior of $\Sigma_{K_0 \setminus K_1} \times \Sigma_{core(K_1)\setminus K_2} \times \cdots \times \Sigma_{core(K_m)\setminus K_{m+1}}$ and these cells form a partition of $\hat{\Sigma}_G$.

Example 4 continued. When G is a bouquet of three circles, $\hat{\Sigma}_G$ is a hexagon, and the partition consists of 13 cells: the vertices, the interiors of the sides, and the interior of the hexagon. Referring back to Figure 1, the sequence $\underline{K}$ is obtained from a displayed bouquet as follows. $K_0 = G$, If K_1 is non-empty then it is the union of the dashed and dotted edges, and if K_2 is non-empty then it is the union of the dotted edges. A vertex corresponds to $\underline{K} = (G, 2\ edges, 1\ edge, \emptyset)$. There are two types of 1-cells, namely $\underline{K} = (G, 2\ edges, \emptyset)$ and $\underline{K} = (G, 1\ edge, \emptyset)$. The 2-cell corresponds to $\underline{K} = (G, \emptyset)$.

Note that the chart $U_{\underline{F}}$ with $\underline{F} = (G = C_0 \supset C_1 \supset \cdots \supset C_q)$ contains the open cell $W(\underline{K})$ with $\underline{K} = (G = K_0 \supset K_1 \supset K_2 \supset \cdots \supset K_{m+1})$ if and only if it intersects it if and only if each $core(K_i)$ equals one of the C_j's, say $C_{j(i)} = core(K_i)$. In that case, the image

$$\Phi_{\underline{F}}(W(\underline{K})) \subset \Delta_{\underline{F}} = \Delta_{C_0\setminus C_1} \times \Delta_{C_1 \setminus C_2} \times \cdots \times \Delta_{C_{q-1}\setminus C_q} \times \Sigma_{C_q}$$

is the open cell

$$\begin{aligned} & Int\ \Delta_{C_0\setminus(C_1\cup K_1)} \times Int\ \Delta_{C_1\setminus(C_2\cup K_1)} \times \cdots \times Int\ \Sigma_{C_{j(1)-1}\setminus(C_{j(1)}\cup K_1)} \\ & \times Int\ \Delta_{C_{j(1)}\setminus(C_{j(1)+1}\cup K_2)} \times Int\ \Delta_{C_{j(1)+1}\setminus(C_{j(1)+2}\cup K_2)} \\ & \qquad\qquad\qquad \times \cdots \times Int\ \Sigma_{C_{j(2)-1}\setminus(C_{j(2)}\cup K_2)} \\ & \times \cdots \times Int\ \Delta_{C_{j(m)}\setminus(C_{j(m)+1}\cup K_{m+1})} \times Int\ \Delta_{C_{j(m)+1}\setminus(C_{j(m)+2}\cup K_{m+1})} \\ & \qquad\qquad\qquad \times \cdots \times Int\ \Sigma_{C_q\setminus(C_{q-1}\cup K_{m+1})}. \end{aligned} \tag{2}$$

In particular, the topological frontier of each open cell is the union of open cells of smaller dimension and the closure $\overline{W}(\underline{K})$ of $W(\underline{K})$ is a compact topological manifold whose interior is homeomorphic to an open ball. Thus $\hat{\Sigma}_G$ is equipped with a homotopy-cell decomposition, and the restriction of this decomposition to a chart coincides with the decomposition into the faces. We will refer to the $\overline{W}(\underline{K})$'s as the '(closed) cells' in this decomposition.

For a fixed $\underline{K} = (G = K_0 \supset K_1 \supset K_2 \supset \cdots \supset K_{m+1})$ it is sometimes convenient to consider the collection of subgraphs $K \subset G$ such that $core(K_i) \subseteq K \subseteq K_i$ for some i. Call these *$\underline{K}$-null subgraphs.* They are characterized by the property that for every core graph $C \not\subseteq K$ the metric ℓ_C^x vanishes on K for each point $x \in W(\underline{K})$. An open cell $W(\underline{K}')$ is contained in $\overline{W}(\underline{K})$ if and only if every $\underline{K}$-null subgraph is a $\underline{K}'$-null subgraph. For graphs $K \supset L$, let K/L denote the graph obtained by collapsing each component of L to a point.

Lemma 2.15. *$\overline{W}(\underline{K})$ is canonically identified with $\prod_{i=0}^m \hat{\Sigma}_{Core(K_i)/K_{i+1}}$.*

Proof. Let $\{\ell_{\overline{C}}\}_{\overline{C}\in\mathcal{C}(Core(K_i)/K_{i+1})} \in \hat{\Sigma}_{Core(K_i)/K_{i+1}}$. If C is core graph of G such that $Core(K_{i+1}) \not\subset C \subset Core(K_i)$ then set $\ell_C(e) = 0$ if $e \subset K_{i+1}$. Otherwise, set $\ell_C(e) = \ell_{\overline{C}}(\overline{e})$ where $\overline{C} = C/K_{i+1}$ and $\overline{e} = e/K_{i+1}$. It is easy to verify that the induced map $\prod_{i=0}^m \hat{\Sigma}_{Core(K_i)/K_{i+1}} \to \prod_{C\in\mathcal{C}(G)} \Sigma_C$ gives a homeomorphism onto $\overline{W}(\underline{K})$. $\square$

Question. Are the $\overline{W}(\underline{K})$'s closed PL cells?

A slight generalization of Proposition 2.6 will be needed in the sequel:

Proposition 2.16. *The restriction $\pi' : \overline{W}(\underline{K}) \to \pi(\overline{W}(\underline{K}))$ of $\pi : \hat{\Sigma}_G \to \Sigma_G$ has point inverses which are AR's.*

Proof. Fix a point ℓ_G of Σ_G and consider the point inverse $\pi^{-1}(\ell_G)$. If H is the core of the subgraph of G on which ℓ_G vanishes, then $\pi^{-1}(\ell_G)$ can be canonically identified with $\hat{\Sigma}_H$. The point inverse intersects the cell $W(\underline{K})$ if and only if K_1 is the subgraph on which ℓ_G vanishes, and in that case the intersection $\pi^{-1}(\ell_G) \cap W(\underline{K})$ is canonically identified with the cell $W(H \supset K_2 \supset K_3 \supset \cdots \supset K_{m+1})$. Therefore, intersections with the cells of $\hat{\Sigma}_G$ induces a cell decomposition of the point inverse $\pi^{-1}(\ell_G)$ that coincides with the cell decomposition of $\hat{\Sigma}_H$. In particular, the nonempty intersections of the point inverse with the closed cells in $\hat{\Sigma}_G$ can be identified with the closed cells in $\hat{\Sigma}_H$, and they are AR's. $\square$

Remark 2.17. The mapping cylinder lines α_x of Proposition 2.8 have the property that with the exception of the terminal endpoint they are contained in a single open cell. Therefore, the mapping cylinder neighborhood restricts to mapping cylinder neighborhoods in each closed cell. More precisely, if π' is the map of Proposition 2.16, then ${\pi'}^{-1}(v)$ has a mapping cylinder neighborhood in $\overline{W}(\underline{K})$ whose frontier is $\hat{Q}_\epsilon \cap \overline{W}(\underline{K})$. The associated map $\rho' : \hat{Q}_\epsilon \cap \overline{W}(\underline{K}) \to {\pi'}^{-1}(v)$ is the restriction of $\rho : \hat{Q}_\epsilon \to \pi^{-1}(v)$. Its point inverses are simplices as well: a point in ${\rho'}^{-1}(y)$ is uniquely determined by specifying ℓ_G-lengths of edges outside $K_1 \cup core(G \setminus e_0) \cup e_0$ and the ℓ_G-volume of $core(G \setminus e_0)$, subject to the condition that these numbers add to ϵ.

2.4. **Functoriality.** Suppose G' is obtained from G by collapsing a forest, and let $\pi : G \to G'$ be the collapsing map. Any $\mathbb{R}$-metric on G' induces an $\mathbb{R}$-metric on G by the requirement that π be an isometry: explicitly, the collapsed edges have length 0, and the other edges have the same length as in G'. If C is a core graph in G, then $\pi(C)$ is a core graph in G', and an $\mathbb{R}$-metric on $\pi(C)$ determines one on C. It is easy to see that if the metrics $\{\ell_{C'}\}_{C' \in \mathcal{C}(G')}$ satisfy the compatibility condition, then so do the pulled back metrics on the core graphs in G. We thus have a natural map $\pi^* : \hat{\Sigma}_{G'} \to \hat{\Sigma}_G$.

It is immediate that π^* is injective. Further, the intersection of a chart $U_{\underline{F}}$ with $\pi^*(\hat{\Sigma}_{G'})$ is sent by $\Phi_{\underline{F}}$ to the intersection of $\Phi(U_{\underline{F}})$ with the face of $\Delta_{\underline{F}}$ given by the condition that the coordinates corresponding to the edges in the collapsed forest are 0. We also see that π^* respects the cell decompositions. We summarize this discussion. If K' is a subgraph of G', define the corresponding subgraph K of G to consist of the edges of K' pulled back to G together with the edges in the collapsed forest.

Lemma 2.18. *π^* is injective, and it sends each pair $(\overline{W}(\underline{K}'), W(\underline{K}'))$ homeomorphically to the corresponding pair $(\overline{W}(\underline{K}), W(\underline{K}))$, where $\underline{K}$ is obtained from $\underline{K}'$ by applying the construction described in the sentence preceding this lemma to each subgraph.*

2.5. **The group $Out(N, K)$ and its Outer Space.** Our arguments will be inductive and we are led to consider more general versions of Culler-Vogtmann's Outer Space [CV86]. Our graphs will have distinguished points and may be disconnected.

Fix an integer $p > 0$ (the number of components), a vector of positive integers $N = (n_1, n_2, \cdots, n_p)$ (ranks of components), and another vector of non-negative integers $K = (k_1, k_2, \cdots, k_p)$ (numbers of distinguished points in each component). Let G be a core graph with p components $G(1), G(2), \cdots, G(p)$ such that the rank of $G(i)$ is n_i. A *labeling* of $G(i)$ is a map $\{1, 2, \cdots, k_i\} \to G(i)$ whose image is in the vertex set of $G(i)$. Elements of the domain of a labeling are *labels*; elements of the image are labeled points. A *labeling* of G is a labeling of each $G(i)$. We say that a core graph G equipped with a labeling has type (N, K) if each valence 2 vertex is in the image of the labeling map. To avoid silly issues, we always assume that each component of rank 1 contains at least two labels, i.e. if $n_i = 1$ then $k_i \geq 2$. By analogy with cutting a surface along essential simple closed curves, this assumption corresponds to the usual requirement that all components have negative Euler characteristic. A map $f : G \to G'$ between two graphs of the same type is *label-preserving* if it sends

$G(i)$ to $G'(i)$ for each i and the labeling of $G'(i)$ is obtained from the labeling of $G(i)$ by composing with $f|G(i)$.

Fix a *model graph* $\mathcal{G} = \mathcal{G}(N, K)$ of type (N, K) in which the labeling of each component is injective (so the labeled points are all distinct). A *marking* of a graph G of type (N, K) is a homotopy equivalence $\mathcal{G} \to G$ which is label-preserving, modulo homotopy through label-preserving maps. The group $Out(N, K)$ is defined to be the group of markings $\mathcal{G} \to \mathcal{G}$. In the case that the number of components $p = 1$, we suppress the parentheses in the vectors N and K. For example:

$$Out(n, 0) \cong Out(F_n)$$
$$Out(n, 1) \cong Aut(F_n)$$
$$Out(1, k) \cong \mathbb{Z}^{k-1} \rtimes \mathbb{Z}/2$$

where in the last case the action of $\mathbb{Z}/2$ on $\mathbb{Z}^{k-1}$ is given by $x \mapsto -x$.

This family of groups was first considered by Hatcher [Hat95] in his work on the homological stability of $Out(F_n)$, where it arose naturally in the course of an inductive argument. This will be the case here.

There is an important alternative description of marked labeled core graphs. For simplicity, we restrict to the case of one component, say of rank n and with k labels; in general the discussion is applied componentwise. First, let $\mathcal{T}$ be the universal cover of the model graph $\mathcal{G}$, together with a choice of lifts of the labeled points – so $\mathcal{T}$ is a labeled minimal free simplicial F_n-tree. Let G be any labeled core graph with marking $f : \mathcal{G} \to G$, and let $\tilde{f} : \mathcal{T} \to T$ be a lift of f to the universal covers. This lift is unique up to post-composition by deck transformations of T. Define the labeling of T so that $\tilde{f}$ is label-preserving. Conversely, if T is a labeled free simplicial F_n-tree, there is a label-preserving F_n-equivariant map $\mathcal{T} \to T$, unique up to homotopy through such maps, and it projects to a marking $\mathcal{G} \to G = T/F_n$. Of course, the quotient graph G may not be a core graph, and we see that there is a 1-1 correspondence between marked labeled core graphs and labeled minimal free simplicial F_n-trees with quotient graph a core graph.

A point $\{\ell_C\}$ of $\hat{\Sigma}_G$ with hierarchy of scales $G = G_0 \supset G_1 \supset \cdots \supset G_m$ can be viewed as giving the structure of a Λ-tree (see [AB87]) to the associated tree T, where Λ is the ordered abelian group $\oplus_{i=0}^m \mathbb{R}$ (with the lexicographic order), and the Λ-length of a lift of an edge $e \in G_i \setminus G_{i+1}$ is the element of Λ with i^{th} coordinate $\ell_{G_i}(e)$ and all other coordinates 0. All edges of T with Λ-length 0 are collapsed.

An element of $Out(n, k)$ with $k > 0$ can, from this point of view (after choosing the first labeled point as the basepoint in $\mathcal{G}$), be regarded as an isomorphism $\alpha : F_n \to F_n$ together with an α-equivariant map $\tilde{f} : \mathcal{T} \to \mathcal{T}$ that preserves the first label, and sends every other label to a translate of itself by a group element, modulo homotopies of $\tilde{f}$ through such maps. Thus, an element of $Out(n, k)$ for

$k > 0$ is a k-tuple $(\alpha, g_2, \cdots, g_k)$ with $\alpha \in Aut(F_n)$ and $g_2, \cdots, g_k \in F_n$, and we see that $Out(n,k)$ in this case is the semi-direct product

$$Out(n,k) = F_n^{k-1} \rtimes Aut(F_n)$$

with the diagonal action of $Aut(F_n)$ on F_n^{k-1}.

Define the category $\mathcal{C}(N,K)$ as follows. Objects are marked graphs $f : \mathcal{G} \to G$ of type (N,K) modulo the equivalence relation: $f : \mathcal{G} \to G$ is equivalent to $f' : \mathcal{G} \to G'$ if there is a label-preserving homeomorphism $h : G \to G'$ such that hf is homotopic to f' through label-preserving maps. Morphisms are induced by collapses of forests $\pi : G \to G'$ (where the labeling and marking of G' are induced by π from the labeling and the marking of G).

The group $Out(N,K)$ acts on the right on the category $\mathcal{C}(N,K)$ via precomposition. Object stabilizers are finite.

The *Spine of Culler-Vogtmann's Outer Space $SCV(N,K)$ for $Out(N,K)$* is defined to be the geometric realization of the category $\mathcal{C}(N,K)$. That is to say, $SCV(N,K)$ is the (underlying space of the) simplicial complex with a vertex for every object of $\mathcal{C}(N,K)$ and a simplex for every chain of morphisms. A point in $SCV(N,K)$ can be thought of as a marked graph of type (N,K) equipped with a non-degenerate $\mathbb{R}$-metric in which all edges have length ≤ 1, and the edges of length < 1 form a forest. The natural (right, simplicial) action of $Out(N,K)$ on $SCV(N,K)$ is cocompact and has finite stabilizers.

Definition 2.19. The *Bordification of Culler-Vogtmann's Outer Space* is the colimit $BCV(N,K)$ of the contravariant functor that to an object $\mathcal{G} \to G$ in $\mathcal{C}(N,K)$ assigns $\hat{\Sigma}_G$ and to a morphism the corresponding inclusion.

In plain English, start with the disjoint union of $\hat{\Sigma}_G$'s over marked graphs G of type (N,K), and then identify $\hat{\Sigma}_{G'}$ with the appropriate subcomplex of $\hat{\Sigma}_G$ whenever G' is obtained from G by collapsing a forest.

In this language, Culler-Vogtmann's Outer Space $CV(N,K)$ for $Out(N,K)$ is the colimit of the contravariant functor that to an object $\mathcal{G} \to G$ in $\mathcal{C}(N,K)$ assigns the open subset of Σ_G that consists of non-degenerate $\mathbb{R}$-metrics.

It is sometimes advantageous to think about $BCV(n,k)$ as a space of Λ-trees. For example, $BCV(1,k)$ for $k > 0$ can be thought of as the space of labelings of $\mathbb{R}$ by k labels, with two labelings considered equivalent if there is a translation that takes one labeling to the other. Thus $BCV(1,k)$ is identified with $\mathbb{R}^k/(a, a, \cdots, a) \cong \mathbb{R}^{k-1}$.

The natural action of $Out(N,K)$ on $CV(N,K)$ and $BCV(N,K)$ is properly discontinuous. It is cocompact on the latter (but not on the former).

Observe that the 3 spaces $CV(N,K)$, $SCV(N,K)$, and $BCV(N,K)$ are homotopy equivalent to each other – the latter two are (equivariantly) proper homotopy equivalent. This is a consequence of the following fact.

Proposition 2.20. [BS73, Theorem 8.2.1] *Let X be a metrizable space and $\mathcal{F} = \{F_i\}_{i\in\Lambda}$ a locally finite closed covering of X which is closed under the operation of taking finite non-empty intersections. Further assume that*

- *the order of the cover is finite,*
- *each F_i is a non-empty absolute retract.*

Then X is homotopy equivalent to the (geometric realization of the) poset of the cover. Further, if each F_i is compact, then X and the poset are proper homotopy equivalent.

The poset of the cover is the simplicial complex whose vertices are the elements of the cover and whose simplices are chains of the F_i's.

Theorem 2.21. *$CV(N,K)$, $SCV(N,K)$, and $BCV(N,K)$ are contractible.*

Proof. It suffices to prove this for the case of a single component, as it is clear that $SCV(N,K)$ is the product $\prod_i SCV(n_i,k_i)$. When there are no labels, this is the theorem of Culler-Vogtmann [CV86]. In the presence of labels, one can either imitate Culler-Vogtmann's argument, or one of the other proofs of contractibility of Outer Space, or else consider the "forget the labels" map $CV(n,k) \to CV(n,0)$ and apply Proposition 2.20 to the cover of $CV(n,0)$ by intersections with Σ_G's and its pull-back. $\square$

Remark 2.22. The set $BCV(N,K) \setminus CV(N,K)$ is a Z-set in $BCV(N,K)$. We will not use this fact, but the proof is straightforward. In the cases of arithmetic groups [BS73] and mapping class groups of compact surfaces [Har86] the bordification is a manifold and the set of ideal points forms the boundary. The notion of a Z-set is a homotopy-theoretic replacement for a "boundary" in a non-manifold setting. See the introduction for definitions.

Question. Is there is an equivariant homeomorphism between bordified Outer Space $BCV(N,K)$ and the regular neighborhood of the spine $SCV(N,K)$ in Outer Space $CV(N,K)$?

3. Morse theory

We now aim to understand the topology at infinity of $BCV(N,K)$ via Morse theory, as explained in the introduction.

Definition 3.1. A *homotopy class* in a graph G of type (N,K) is a map $f: J \to G$ from a compact 1-manifold J that sends ∂J to the labeled points and is essential on circle components, up to homotopy rel ∂J.

Via the marking, any homotopy class in the model graph $\mathcal{G} = \mathcal{G}(N,K)$ induces a homotopy class in a marked graph of type (N,K). We say that $f: J \to G$ is an *immersion* if it is either locally injective or constant on each component, and further, all components on which f is constant are arcs.

Any homotopy class can be represented by an immersion (uniquely up to reparametrization).

The simplest function one could define on $BCV(N,K)$ would be the length (with respect to ℓ_G) of a particular homotopy class. Such a function would not be "Morse" in that it would be constant on large subcomplexes. We are thus led to a refinement of this idea as follows.

3.1. **Hierarchy of lengths of a homotopy class.** Let $\alpha : J \to G$ be an immersion representing a homotopy class, where G is a marked graph of type (N,K). Suppose $\{\ell_C\}_{C\in\mathcal{C}(G)}$ is a point in $\hat{\Sigma}_G \subset BCV(N,K)$, and say its hierarchy of scales is $G = G_0 \supset G_1 \supset \cdots \supset G_m$. Note that m is smaller than the *total rank* $|N| = n_1 + n_2 + \cdots + n_p$. For $i = 0, 1, \cdots, m$ define $l_i^\alpha(\{\ell_C\})$ to be the total ℓ_{G_i}-length of edges of G_i crossed by α, counted with multiplicity. For $m < i < |N|$ define $l_i^\alpha(\{\ell_C\}) = 0$. Note that $l_0^\alpha : BCV(N,K) \to \mathbb{R}_+ = [0,\infty)$ is continuous (and factors through $\pi : \hat{\Sigma}_G \to \Sigma_G$), while l_i^α are generally discontinuous for $i > 0$.

Fix a (possibly finite) sequence $\mathcal{A} = (\alpha_1, \alpha_2, \cdots)$ of homotopy classes and define

$$L = L_{\mathcal{A}} : BCV(N,K) \to (\mathbb{R}_+ \times \mathbb{R}_+ \times \cdots)^{|N|}$$

by

$$L = L_0 \times L_1 \times \cdots \times L_{|N|-1}$$

where

$$L_j = l_j^{\alpha_1} \times l_j^{\alpha_2} \times \cdots$$

We will use the more suggestive notation $\mathbb{R}_+^{\mathcal{A}}$ for $\mathbb{R}_+ \times \mathbb{R}_+ \times \cdots$.

Definition 3.2. A sequence $\mathcal{A}$ of homotopy classes in the model graph $\mathcal{G} = \mathcal{G}(N,K)$ is *ample* provided that $L_{\mathcal{A}}$ is injective.

Note that when $\mathcal{A}$ is ample, L_0 already determines ℓ_G.

Lemma 3.3. *If every homotopy class occurs in $\mathcal{A}$ at least once, then $\mathcal{A}$ is ample.*

Proof. From the point of view of Λ-trees, these are [AB87, Theorem 7.13] and [AB87, 3.14] (for the case when there is only one component and at most one labeled point). The general case easily follows from these. □

Remark 3.4. That finite ample sequences do not exist follows from [SV92]. Say that a sequence of homotopy classes is *weakly ample* if for every marked graph G and any two edges in G there is a homotopy class in the sequence that crosses the two edges with distinct multiplicities. The existence of a finite weakly ample sequence would simplify the proof somewhat.

Question. Do finite weakly ample sequences exist?

With respect to the natural (product) topology on $(\mathbb{R}_+^{\mathcal{A}})^{|N|}$, L is discontinuous. However, we view L as a "Morse function" and study the super-level sets in the following sense. We order $\mathbb{R}_+^{\mathcal{A}}$ and then $(\mathbb{R}_+^{\mathcal{A}})^{|N|}$ lexicographically. For $t \in (\mathbb{R}_+^{\mathcal{A}})^{|N|}$ consider the "intervals" $(t, \cdot) = \{s \in (\mathbb{R}_+^{\mathcal{A}})^{|N|} | s > t\}$ and $[t, \cdot) = \{s \in (\mathbb{R}_+^{\mathcal{A}})^{|N|} | s \geq t\}$. We are interested in understanding the change in the homotopy type of $L^{-1}(t, \cdot)$ and $L^{-1}[t, \cdot)$ as t changes.

Define the *weight* of an edge $e \subset G$ as the $\mathcal{A}$-tuple of non-negative integers

$$w(e) = (\mu_e(\alpha_1), \mu_e(\alpha_2), \cdots)$$

where $\mu_e(\alpha)$ is the number of times the immersion representing α crosses e. We order the weights lexicographically.

Example 4 continued. Let G be a bouquet of three circles, and consider the subset $\hat{\Sigma}_G$ of $BCV(3, 0)$ corresponding to the identity marking (we are taking G for the model graph). Let $\mathcal{A}$ consist of the single homotopy class $a^2b^4c^6$. Of course, this $\mathcal{A}$ is not ample. In Figure 3, the values of L at the representative points are given.

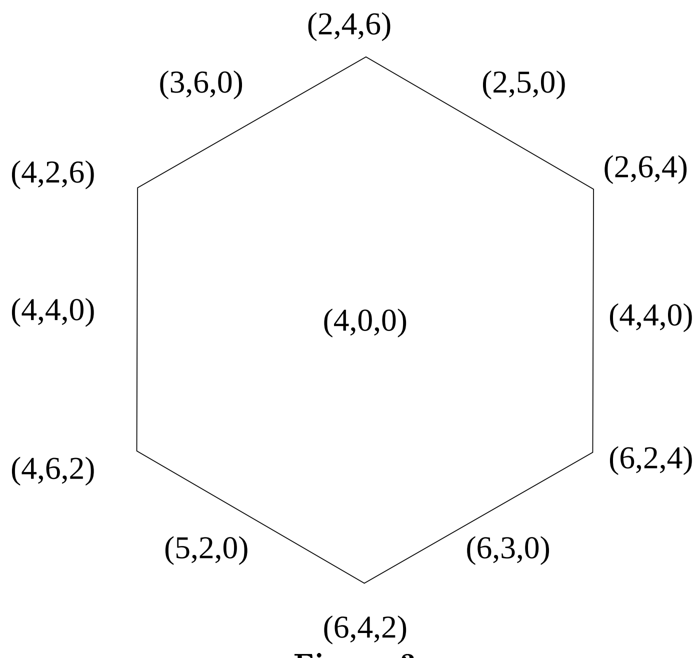

Figure 3

3.2. **The behavior on cells.** The intersections with $\overline{W}(\underline{K})$ with super-level sets. The following lemma will be used to show these sets are semi-algebraic.

Lemma 3.5. *Let $f : \mathbb{R}^n \to (\mathbb{R}^{\mathbb{N}})^M$ with each coordinate $(f_i)_j : \mathbb{R}^n \to \mathbb{R}$ polynomial. Order $\mathbb{R}^{\mathbb{N}}$ and then $(\mathbb{R}^{\mathbb{N}})^M$ lexicographically, and for $t \in (\mathbb{R}^{\mathbb{N}})^M$ let $(t, \cdot)$, respectively $[t, \cdot)$, denote the subset of $(\mathbb{R}^{\mathbb{N}})^M$ consisting of elements $> t$, respectively $\geq t$. Then $f^{-1}(t)$, $f^{-1}(t, \cdot)$, and $f^{-1}[t, \cdot)$ are semi-algebraic.*

Proof. We do the case $M = 1$, the general case follows easily. The set $f^{-1}(t)$ is algebraic by Hilbert's Basis Theorem, see for example [AM69]. In particular, there is a C such that for all $x \in \mathbb{R}^{\mathbb{N}}$ if $f_j(x) = t_j$ for $j \leq C$ then $f_i(x) = t_i$ for all $i \in \mathbb{N}$.

Let X_i denote $\cap_{j=1}^{i} f_j^{-1}(t_j)$. We have

$$f^{-1}(t, \cdot) = \bigcup_{k=1}^{\infty} \left(f_k^{-1}(t_k, \infty) \cap X_{k-1}\right).$$

By the above comment, $f_k^{-1}(t_k, \infty) \cap X_{k-1}$ is empty for large k. Hence, $f^{-1}(t, \cdot)$ is a finite union of semi-algebraic sets and so is itself semi-algebraic. The case of $f^{-1}[t, \cdot)$ is similar. ◻

Proposition 3.6. *Let $\Phi = \Phi_{\underline{F}} : U = U_{\underline{F}} \to \Delta = \Delta_{\underline{F}}$ be a chart map as in Definition 2.9. For every $t \in (\mathbb{R}_+^{\mathcal{A}})^{|N|}$ the sets*

$$\Phi(L^{-1}(t, \cdot) \cap U) \text{ and } \Phi(L^{-1}[t, \cdot) \cap U)$$

are semi-algebraic (viewed as subsets of Δ, which in turn is viewed as a polyhedral subset of Euclidean space). In particular, the super-level sets intersect the cells $\overline{W}(\underline{K})$ in sets which are locally semi-algebraic (and hence they are ANR's).

Proof. Say two points of U are equivalent if they have exactly the same hierarchy of scales. Let E be the equivalence class corresponding to $\underline{K} = (G = C_0' \supset C_1' \supset \cdots \supset C_{q'}')$ (which is a subsequence of $\underline{F}$). Since the union of a finite collection of semi-algebraic sets is semi-algebraic, it suffices to show that $\Phi(L^{-1}(t, \cdot) \cap E)$ and $\Phi(L^{-1}[t, \cdot) \cap E)$ are semi-algebraic.

Note that E is the set of points in $U_{\underline{K}}$ with $\Phi_{\underline{K}}$-image in $\Sigma_{\underline{K}} = \Sigma_{C_0' \setminus C_1'} \times \cdots \times \Sigma_{C_{q'-1}' \setminus C_{q'}'} \times \Sigma_{C_{q'}'}$ and that each component of the restriction of $L\Phi_{\underline{K}}^{-1}$ to $\Phi_{\underline{K}}(E)$ extends to a linear function on $\Delta_{\underline{K}}$. From Proposition 2.12 and Lemma 3.5 it follows that $\Phi_{\underline{K}}(L^{-1}(t, \cdot) \cap E)$ and $\Phi_{\underline{K}}(L^{-1}[t, \cdot) \cap E)$ are semi-algebraic subsets of $\Delta_{\underline{K}}$. The proposition now follows from Proposition 2.12 because the transition map $\Delta_{\underline{F}} \to \Delta_{\underline{K}}$ is polynomial (see Remark 2.13). ◻

Theorem 3.7. *Fix a (homotopy-)cell $\overline{W}(\underline{K}) \subseteq \hat{\Sigma}_G$ and assume that $\mathcal{A}$ is ample. The intersections $\overline{W}(\underline{K}) \cap L^{-1}(t, \cdot)$ and $\overline{W}(\underline{K}) \cap L^{-1}[t, \cdot)$ are either contractible or empty.*

Example 4 continued. Before proving the theorem, we continue with our example with $\mathcal{A} = \{a^2b^4c^6\}$. Recall that this $\mathcal{A}$ is not ample, but it illustrates the point. The super-level sets in this case are: empty, single point (vertex where the function attains the maximum on $\hat{\Sigma}_G$), arc (possibly with an endpoint missing), 2-cell (possibly with an arc in the boundary missing) and the wedge of a 2-cell (possibly with a boundary arc missing) and an arc along boundary points. The set $L^{-1}(t,\cdot)$ is pictured below with $t = (4,0,0)$.

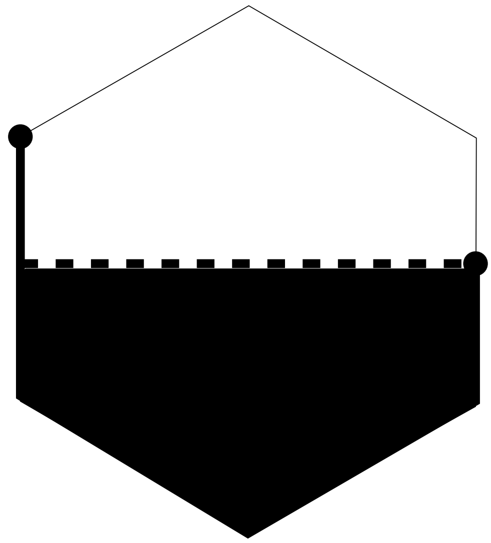

Figure 4

Proof of Theorem 3.7. Denote either one of the two sets by W_t.

We argue by induction on the total rank $|N|$. If $|N| = 1$, then G is a circle with k labeled points, $\hat{\Sigma}_G = \Sigma_G$ is the $(k-1)$-simplex, and each component function of L is linear, so the result in this case is immediate.

Now assume $|N| > 1$. Let $\pi : \hat{\Sigma}_G \to \Sigma_G$ be the natural projection to the simplex Σ_G. The image of $\overline{W}(\underline{K})$ under π is a face of Σ_G. Denote by $\pi' : \overline{W}(\underline{K}) \to \pi(\overline{W}(\underline{K}))$ the restriction of π.

Write $t = (t_0, t_1, \cdots, t_{|N|-1})$ and observe that L_0 factors through $\pi(\overline{W}(\underline{K}))$:

$$\begin{tikzcd} \overline{W}(\underline{K}) \arrow[rr, "L_0"] \arrow[dr, "\pi'"'] & & \mathbb{R}_+^{\mathcal{A}} \\ & \pi(\overline{W}(\underline{K})) \arrow[ur, "\ell"'] & \end{tikzcd}$$

and we have

$$\pi'^{-1}(\ell^{-1}(t_0,\cdot)) \subseteq W_t \subseteq \pi'^{-1}(\ell^{-1}[t_0,\cdot))$$

The difference between the two sets that sandwich W_t is $\pi'^{-1}\ell^{-1}(t_0)$. Note that $\ell^{-1}(t_0)$ is at most a point (by the ampleness assumption).

The proof of the theorem will be broken up into a series of claims.

Claim 1. Both sets sandwiching W_t are contractible (or empty) ANR's.

Indeed, it follows from Proposition 3.6 that they are ANR's. Also, they are the preimages of convex sets under a proper PL map (see [Hir75]) with contractible point-inverses (by Proposition 2.16), and the claim follows from [Sma57].

We show that this sandwich is generically a Z-sandwich (see the introduction for definitions).

Claim 2. If $\ell^{-1}(t_0)$ is not a vertex, then $\pi'^{-1}\ell^{-1}(t_0)$ is a Z-set in $\pi'^{-1}\ell^{-1}[t_0,\cdot)$.

We show that $\pi'^{-1}\ell^{-1}(t_0)$ is locally Z at all points. Indeed, fix a point $z \in \pi'^{-1}\ell^{-1}(t_0)$. Say e_1 and e_2 are distinct edges of G with nonzero ℓ_G-length (at z). For concreteness, say $w(e_1) > w(e_2)$ (ampleness guarantees that they are not equal). For a small $s \geq 0$ and a point $\{\ell_C\} \in \pi'^{-1}\ell^{-1}[t_0,\cdot)$ near z define $F_s(\{\ell_C\})$ by adding s to the ℓ_G-length of e_1 and subtracting s from the ℓ_G-length of e_2. This gives an instant deformation of a neighborhood of z into $\pi'^{-1}\ell^{-1}(t_0,\cdot)$ and proves the claim.

If $\ell^{-1}(t_0)$ is not a vertex, the theorem follows from the Z-sandwich fact. Assume that $\ell^{-1}(t_0)$ is a vertex v. If $W_t \cap \pi'^{-1}(v) = \emptyset$ then $W_t = \pi'^{-1}\ell^{-1}(t_0,\cdot)$ and the theorem again follows.

Now assume that $\ell^{-1}(t_0)$ is a vertex v and that $W_t \cap \pi'^{-1}(v) \neq \emptyset$. Denote by e_0 the unique edge of nonzero length (length 1) in the $\mathbb{R}$-metric corresponding to v.

Claim 3. Suppose that there is an edge e of G such that $e \not\subseteq K_1 \cup core(G \setminus e_0)$ and $w(e) > w(e_0)$ (K_1 is a graph in $\underline{K}$ that defines the cell $W(\underline{K})$). Then $\pi'^{-1}(v)$ is a Z-set in $\pi'^{-1}\ell^{-1}[t_0,\cdot)$.

The desired instant deformation of a neighborhood of $\pi'^{-1}(v)$ is obtained by decreasing $\ell_G(e_0)$ and increasing $\ell_G(e)$ by $s \geq 0$. This deformation is well-defined since $e \not\subseteq core(G \setminus e_0)$ and we leave ℓ_C unchanged for $C \subseteq core(G \setminus e_0)$. The condition $e \not\subseteq K_1$ guarantees that the deformation takes place in $\overline{W}(\underline{K})$, and finally $w(e) > w(e_0)$ means that the function L_0 increases during the deformation, so that the deformation takes place inside $\pi'^{-1}(\ell^{-1}[t_0,\cdot))$ and proves the claim

Under the hypotheses of the Claim 3, the theorem follows from the Z-sandwich as before.

We now finish the proof of the theorem under the assumption $W_t \cap {\pi'}^{-1}(v) \neq \emptyset$ plus the assumption that there are no edges e as in Claim 3. We may assume that W_t is not contained in a proper sub-cell of $\overline{W}(\underline{K})$, since otherwise the statement of the theorem follows by induction on the dimension of the cell. In particular, there are interior points of $\pi(\overline{W}(\underline{K}))$ with larger L_0-value than v. It is convenient to represent W_t as the union of two open sets ${\pi'}^{-1}(\ell^{-1}(t_0, \cdot))$ and $W_t \cap \hat{Q}_{<\epsilon}$, where $\hat{Q}_{<\epsilon} = \hat{Q}_{\leq\epsilon} \setminus \hat{Q}_\epsilon$ (see Proposition 2.8). The first set is contractible by Proposition 2.16 and [Sma57], and so is the intersection ${\pi'}^{-1}(\ell^{-1}(t_0, \cdot) \cap Q_{<\epsilon})$ where $Q_{<\epsilon} = Q_{\leq\epsilon} \setminus Q_\epsilon$. So it remains to prove the following claim.

Claim 4. $W_t \cap \hat{Q}_{<\epsilon}$ is contractible.

Denote by $Q^0_{<\epsilon}$ the intersection of $Q_{<\epsilon}$ with the face of Σ_G where the ℓ_G-lengths of all edges e outside $e_0 \cup core(G \setminus e_0)$ is 0. Likewise, let $\hat{Q}^0_{<\epsilon}$ be the intersection of $\hat{Q}_{<\epsilon}$ with the preimage of that face under π.

We first note that there is a deformation retraction of $W_t \cap \hat{Q}_{<\epsilon}$ onto $W_t \cap \hat{Q}^0_{<\epsilon}$. The deformation does not change ℓ_C when C is a core graph that does not contain e_0. It suffices to specify the deformation on ℓ_G. On ℓ_G, the deformation is given by decreasing the lengths of edges e outside $e_0 \cup core(G \setminus e_0)$ with constant speed and increasing $\ell_G(e_0)$ to insure that the volume is 1. During this deformation, the value of L_0 can only increase due to our assumption that $w(e) < w(e_0)$ for all affected edges e. Note also that if a point moves during this deformation, it does not land into ${\pi'}^{-1}(v)$ (points in that set have minimal L_0).

It remains to argue that $W_t \cap \hat{Q}^0_{<\epsilon}$ is contractible. Consider the subset E of ${\pi'}^{-1}(v)$ consisting of the points x with $L_1(x) > L_0(x)$ (these are elements of the ordered abelian group $\mathbb{R}^{\mathcal{A}}$). Since L_0 is constant on ${\pi'}^{-1}(v)$ (and equal to $w(e_0)$), the set E is a super-level set in the cell ${\pi'}^{-1}(v)$, and is thus contractible or empty by induction on the dimension of the cell. The intersection of W_t with ${\pi'}^{-1}(v)$ is also a super-level set in ${\pi'}^{-1}(v)$, and we consider two cases:

Case 1. $E \subseteq W_t \cap {\pi'}^{-1}(v)$.

We construct a deformation of $W_t \cap \hat{Q}^0_{<\epsilon}$ onto $W_t \cap {\pi'}^{-1}(v)$ (which is contractible). The deformation consists of following the mapping cylinder lines of Proposition 2.8. This amounts to increasing the length of e_0 to 1 with linear speed, and decreasing the ℓ_G-lengths of other edges with linear speed, keeping the volume equal to 1. It remains to argue that all points that move during this deformation land into E. Let x be a landing point, and x_δ the point where

$\ell_G(e_0) = 1 - \delta$, $volume(core(G \setminus e_0)) = \delta$, and which lands to x. Then

$$L_0(x_\delta) = (1 - \delta)w(e_0) + \delta L_1(x).$$

Since $L_0(x) = w(e_0)$ and $L_0(x) \leq L_0(x_\delta)$, it follows that $L_0(x) \leq L_1(x)$. Equality cannot hold since the collection $\mathcal{A}$ is ample.

Case 2. $W_t \cap {\pi'}^{-1}(v) \subseteq E$.

In this case, we will construct an instant deformation of $W_t \cap \hat{Q}^0_{<\epsilon}$ into

$${\pi'}^{-1}(\ell^{-1}(t_0, \cdot)) \cap \hat{Q}^0_{<\epsilon} = (W_t \cap \hat{Q}^0_{<\epsilon}) \setminus {\pi'}^{-1}(v).$$

Since the latter set is contractible (either by appealing to [Sma57] again, or noting that it is a deformation retract of the contractible set $(W_t \cap \hat{Q}_{<\epsilon}) \setminus {\pi'}^{-1}(v))$, it will follow that $W_t \cap \hat{Q}^0_{<\epsilon}$ is contractible.

The deformation is the reverse of the deformation of Case 1. Choose an instant deformation h_s of $[0, \epsilon)$ into $(0, \epsilon)$ that moves all points in the positive direction. The deformation we are defining does not change ℓ_C when $e_0 \not\subseteq C$, and ℓ_G changes according to h_s: at time s, $volume(core(G \setminus e_0))$ is equal to h_s applied to the volume of the same graph at the starting point, and $\ell_G(e_0) = 1 - volume(core(G \setminus e_0))$ at all points. The calculation in Case 1 combined with the assumption $W_t \cap {\pi'}^{-1}(v) \subseteq E$ shows that the L_0-value increases during the deformation, and so this deformation takes place in $W_t \cap \hat{Q}^0_{<\epsilon}$. If $core(G \setminus e_0) \not\subseteq K_1$, then $\ell_{core(G \setminus e_0)}$ vanishes on K_1 for all points in $\overline{W}(\underline{K})$, and as a result ℓ_G vanishes on $K_1 \cap core(G \setminus e_0)$ throughout the deformation, so that the deformation takes place in the closed cell. The last case to consider is that $core(G \setminus e_0) \subseteq K_1$. In that case $W_t \cap \hat{Q}^0_{<\epsilon}$ is equal to $W_t \cap {\pi'}^{-1}(v)$, which is contractible by induction. □

3.3. **Critical points.** Denote either of the super-level sets $L^{-1}(t, \cdot)$ or $L^{-1}[t, \cdot)$ by B_t. Consider the closed cover $\mathcal{U}(B_t)$ of B_t by the nonempty intersections of B_t with the cells $\overline{W}(\underline{K})$ of $BCV(N, K)$. Denote by $\mathcal{P}(B_t)$ the poset of this cover.

Proposition 3.8. *B_t is homotopy equivalent to $\mathcal{P}(B_t)$.*

Proof. This follows from Proposition 2.20 and Theorem 3.7 □

Proposition 3.9. *If B_t and $B_{t'}$ are two super-level sets such that $B_t \subset B_{t'}$ but $B_{t'} \setminus B_t$ contains no vertices, then B_t and $B_{t'}$ are homotopy equivalent.*

Proof. The Morse function attains its maximum on every cell at a vertex. It then follows that a cell intersects B_t if and only if it intersects $B_{t'}$, so that the posets $\mathcal{P}(B_t)$ and $\mathcal{P}(B_{t'})$ are isomorphic. □

Proposition 3.9 can be interpreted as saying that the vertices are the only critical points of L.

4. Ascending links

We next turn to the study of the change of the homotopy type of a super-level set as it passes through a vertex v. Intuitively, one expects that the subspace corresponding to the set of directions from v in which L increases gets coned off. This is what we now make precise.

Fix a cell $\overline{W}(\underline{K}) \subseteq \hat{\Sigma}_G \subset BCV(N,K)$ and a vertex v of $\overline{W}(\underline{K})$. An ample sequence $\mathcal{A}$ and the corresponding Morse function L are fixed as well.

The chart associated with the hierarchy of scales at v identifies a neighborhood of v in $\overline{W}(\underline{K}) \subset \hat{\Sigma}_G \subset BCV(N,K)$ with a neighborhood of a vertex of a polytope. We define $Lk(v, \overline{W}(\underline{K}))$ to be the space of germs of unit speed straight line segments leaving the vertex into the polytope. Thus $Lk(v, \overline{W}(\underline{K}))$ is naturally identified with a spherical simplex in the unit sphere (and so is an AR if nonempty). The construction is natural with respect to inclusions of cells, and we define $Lk(v, BCV(N,K))$ to be the union of $Lk(v, \overline{W}(\underline{K}))$'s with identifications induced by inclusions of cells. Thus $Lk(v, BCV(N,K))$ is a spherical cell complex (and hence an ANR). We also refer to $Lk(v, BCV(N,K))$ as the set of *directions* at v.

Definition 4.1. Fix a Morse function $L : BCV(N,K) \to (\mathbb{R}_+^{\mathcal{A}})^{|N|}$ coming from an ample sequence $\mathcal{A}$ of relative homotopy classes. For a vertex $v \in \overline{W}(\underline{K}) \subseteq \hat{\Sigma}_G \subset BCV(N,K)$ define the *ascending link* $Lk_>(v, \overline{W}(\underline{K}))$ to be the subset of $Lk(v, \overline{W}(\underline{K}))$ consisting of the germs of directions d at v with the property that the first coordinate function $l_i^{\alpha_j}$ which is not constant along d is monotonically increasing along d (in a small neighborhood of v). We define $Lk_>(v, BCV(N,K))$ as the subset of $Lk(v, BCV(N,K))$ similarly.

We also define the *descending links* $Lk_<(v, \overline{W}(\underline{K}))$ and $Lk_<(v, BCV(N,K))$ analogously (replacing 'increasing' by 'decreasing').

At a vertex $v \in \hat{\Sigma}_G$, the metric ℓ_{G_i} is determined by edges e_i, $0 \le i < |N|$, such that $\ell_{G_i}(e_i) = 1$. There is a unique chart containing v determined by $\underline{F} = (G = C_0 \supset C_1 \supset \cdots \supset C_{|N|-1})$ where $C_{i+1} = core(C_i \setminus e_i)$ for $i \ge 0$, this is *the chart at* v. Let $x = (x_0, x_1, \cdots, x_{|N|-1}) \in \Delta_v = \Delta_{G \setminus C_1} \times \Delta_{C_1 \setminus C_2} \times \cdots \times \Sigma_{C_{|N|-1}}$ denote the coordinates at v. We may assume that the first coordinate $x_{i,1}$ of $x_i \in \Delta_{C_i \setminus C_{i+1}}$ corresponds to the edge of ℓ_{C_i}-length 1 at v (so all other coordinates of x_i are 0 at v).

The local coordinate chart Δ_v at v has a form that naturally lends itself to inductive arguments. Namely, $\Delta_v = \Delta_{G \setminus C_1} \times \Delta_w$ where Δ_w is the local chart of the vertex w of $\hat{\Sigma}_{C_1}$ with hierarchy of scales $C_1 \supset C_2 \supset \cdots \supset C_{|N|-1}$. The coordinate map Φ_v at v is a product with second coordinate Φ_w.

If α is a homotopy class, then the function l_0^α can be written in the local coordinates $\Delta_{G\setminus C_1} \times \Delta_w$ as

$$l_0^\alpha(x_0, y) = x_0 \cdot a_0 + (1 - S(x_0)) l_0^{[\alpha]}(y) \tag{3}$$

where a_i is the vector $(a_{i,j})$ of multiplicities with which α crosses the edges of $C_i \setminus C_{i+1}$, $\cdot$ is the dot product, $S(x_i)$ denotes the sum of the coordinates of the point x_i, and $l_0^{[\alpha]}$ the function on Δ_w corresponding to the homotopy class $[\alpha]$ obtained from α by deleting edges in $G \setminus C_1$.

More generally, if V is a subset containing v of the chart at v where L_j is constant for $j < i$ then the restriction of l_i^α is given in the coordinates $\Delta_{G\setminus C_1} \times \Delta_{C_1\setminus C_2} \times \cdots \times \Delta_{C_{i-1}\setminus C_i} \times \Delta_{w^i}$ by

$$l_i^\alpha|_V(x_0, x_1, \cdots, x_{i-1}, y) = x_i \cdot a_i + (1 - S(x_i)) l_i^{[\alpha]}(y) \tag{4}$$

where now w^i is the vertex with hierarchy of scales $C_i \supset C_{i+1} \supset \cdots \supset C_{|N|-1}$ and $[\alpha]$ is obtained from α by deleting edges in $G \setminus C_i$.

The following lemma is easily proved inductively from Formulas 3 and 4.

Lemma 4.2. 1. *Each coordinate of L_0 is a polynomial in x.*
2. *Each coordinate of L_i is a polynomial function on each straight line d leaving the vertex v provided L_j is constant on d in a neighborhood of v for $j < i$.*
3. $l_0^\alpha(v) = a_{0,1}$.
4. *The derivative of l_0^α at v is the dot product $(a_{0,1} - a_{1,1}, a_{0,2} - a_{1,1}, a_{0,3} - a_{1,1}, \cdots) \cdot x_0$.*

$\square$

Corollary 4.3. *$Lk(v, BCV(N, K))$ is the disjoint union of the descending link $Lk_<(v, BCV(N, K))$ and of the ascending link $Lk_>(v, BCV(N, K))$.*

Proof. Since $\mathcal{A}$ is ample, there is a minimal i and α such that the restriction of l_i^α to d is not constant. By Lemma 4.2(2), $l_i^\alpha|_d$ is a polynomial. So, l_i^α and hence L is increasing or decreasing along d, and d is accordingly either in the ascending or descending link at v. $\square$

We focus on ascending links, as only those are needed for understanding of the topology at infinity. The next lemma will be used to show that they are ANR's.

Lemma 4.4. *Let $P : \mathbb{R}^n \to \mathbb{R}$ be a polynomial with P_i its homogeneous part of degree i. Let d be a direction at $0 \in \mathbb{R}^n$ represented by the unit vector x. Then P is increasing (respectively constant, decreasing) in the direction d iff $(P_1(x), P_2(x), \cdots)$ is greater than 0 (respectively $= 0$, < 0) in lexicographically ordered $\mathbb{R}^{\mathbb{N}}$.*

Proof. We have $P(tx) - P(0) = \Sigma_{i>0} P_i(tx) = \Sigma_{i>0} t^i P_i(x)$ which is 0 if the vector $(P_1(x), P_2(x), \cdots) = 0$. Otherwise, let k be minimal greater than 0 such that $P_k(x) \neq 0$. Then $P(tx) - P(0) = \Sigma_{i \geq k} t^i P_i(x)$ has a non-zero leading term of degree k which dominates for small t. $\square$

Corollary 4.5. *Let $f : \mathbb{R}^n \to (\mathbb{R}^{\mathbb{N}})^M$ be as Lemma 3.5. Then the set of directions at $0 \in \mathbb{R}^n$ (as a subset of the unit sphere) along which f is increasing is semi-algebraic.*

Proof. Let $((f_k)_j)_i$ denote the homogeneous part of degree i of the component $((f_k)_j$ of f. According to Lemma 4.4, f is increasing in the direction given by the unit vector x iff the vector $(((f_k)_j)_i(x))$ is greater than 0 in lexicographically ordered $((\mathbb{R}^{\mathbb{N}})^{\mathbb{N}})^M$. The result now follows as in Lemma 3.5. $\square$

Corollary 4.6. *$Lk_>(v, \overline{W}(\underline{K}))$ and $Lk_>(v, BCV(N, K))$ are ANR's.*

Proof. Since unions of ANR's over ANR's are ANR's (see Section 1.1), it is enough to show that $Lk_>(v, \overline{W}(\underline{K}))$ is an ANR. As in Proposition 3.6, it is enough to show that $Lk_>(v, W(\underline{K}))$ is semi-algebraic, and this follows as in Proposition 3.6 from Corollary 4.5. $\square$

We will need the following Z-set lemma which follows quickly from definitions.

Lemma 4.7. *Let A, B, and A' be ANRs with $A' \subset A$.*

1. *If A is contractible then $\emptyset * B$ is a Z-set in $A * B$.*
2. *If A' and A differ by a Z-set in A then $A' * B$ and $A * B$ differ by Z-set in $A * B$.*
3. *If both A is contractible and A' and A differ by a Z-set in A then $(A' * B) \setminus (\emptyset * B)$ and $A * B$ differ by a Z-set in $A * B$.*

$\square$

Proposition 4.8.
1. *If $L|\overline{W}(\underline{K})$ attains its unique maximum at v, then $Lk_>(v, \overline{W}(\underline{K})) = \emptyset$.*
2. *Otherwise, there are points in $\overline{W}(\underline{K})$ where L has larger value than at v and $Lk_>(v, \overline{W}(\underline{K}))$ is contractible (and so nonempty). Moreover, there is a vertex of $\overline{W}(\underline{K})$ with larger L-value that spans a 1-cell with v (and the corresponding edge determines a direction d in $Lk_>(v, \overline{W}(\underline{K}))$).*

Proof. By Lemma 2.15, $\overline{W}(\underline{K})$ has the form $\prod_i \hat{\Sigma}_{H_i}$ where each H_i is a quotient of a subgraph of G. The map L can be identified with $\prod_i L_i$ where the target is ordered lexicographically, and L_i is the Morse function on $\hat{\Sigma}_{H_i}$ obtained by restricting the elements of $\mathcal{A}$ in the obvious way. The vertex v is identified with a vertex $(v_1, v_2, \cdots) \in \prod_i \hat{\Sigma}_{H_i}$. Set $v' = (v_2, v_3, \cdots)$.

With these identifications, we also have

$$Lk(v, \overline{W}(\underline{K})) = Lk(v_1, \hat{\Sigma}_{H_1}) * Lk(v', \prod_{i>1} \hat{\Sigma}_{H_i})$$

and

$$Lk_>(v, \overline{W}(\underline{K})) = \left(Lk_>(v_1, \hat{\Sigma}_{H_1}) * Lk(v', \prod_{i>1} \hat{\Sigma}_{H_i})\right) \setminus \left(\emptyset * Lk_<(v', \prod_{i>1} \hat{\Sigma}_{H_i})\right).$$

Let k and α be minimal so that l_k^α is not constant on a neighborhood of v in the cell $\overline{W}(\underline{K})$ (ampleness). Let $Lk_>(l_k^\alpha)$, respectively $Lk_\geq(l_k^\alpha)$, denote the subset of $Lk(v, \overline{W}(\underline{K}))$ along which l_k^α is increasing, respectively non-decreasing. We have a sandwich

$$Lk_>(l_k^\alpha) \subset Lk_>(v, \overline{W}(\underline{K})) \subset Lk_\geq(l_k^\alpha).$$

Claim:

1. If $Lk_>(l_k^\alpha)$ is non-empty then the above sandwich is a Z-sandwich of contractible sets, and there is a vertex in $\overline{W}(\underline{K})$ that spans an edge with v and with larger l_k^α-value than v.
2. If $Lk_>(l_k^\alpha)$ is empty, then $Lk_>(v, \overline{W}(\underline{K}))$ is contained in a proper subcell of $Lk(v, \overline{W}(\underline{K}))$.

To prove the claim, we argue by induction on the dimension of $\overline{W}(\underline{K})$ and the number of edges in G. The claim is clear if the dimension is 1.

If the ℓ_G-metric is the same for all points in the cell then by ampleness L_0 is constant on the cell, and the claim follows by induction after replacing G by C_1. Hence we may assume that $k = 0$. With this assumption, the restriction of l_0^α to $\prod_{i>1} \hat{\Sigma}_{H_i}$ is zero.

It follows that the restriction $l_0^{[\alpha]}$ of l_0^α to $\hat{\Sigma}_{H_1}$ is not constant and that $Lk_>(l_0^\alpha)$, respectively $Lk_\geq(l_0^\alpha)$, is identified with

$$\left(Lk_>(l_0^{[\alpha]}) * Lk(v', \prod_{i>1} \hat{\Sigma}_{H_i})\right) \setminus \left(\emptyset * Lk(v', \prod_{i>1} \hat{\Sigma}_{H_i})\right),$$

respectively $Lk_\geq(l_0^{[\alpha]}) * Lk(v', \prod_{i>1} \hat{\Sigma}_{H_i})$.

By our induction hypothesis, Lemma 4.7, Corollary 4.6, and the discussion above, it is enough prove the claim in the case that $\overline{W}(\underline{K}) = \hat{\Sigma}_G$, i.e. $\underline{K}$ is the hierarchy of scales at v. We will use notation as in Lemma 4.2.

Case 1. The derivative D of l_0^α is non-zero on some direction in $Lk(v, \overline{W}(\underline{K}))$.

Let $D_>$, respectively $D_\geq$, denote the subset of $Lk(v, \overline{W}(\underline{K}))$ where D is positive, respectively non-negative. We can enlarge our sandwich

$$D_> \subset Lk_>(l_0^\alpha) \subset Lk_>(v, \overline{W}(\underline{K})) \subset Lk_\geq(l_0^\alpha) \subset D_\geq.$$

Since D is linear, the outermost sets are visibly convex subsets of $Lk(v, \overline{W}(\underline{K}))$ and the sandwich is a Z-sandwich unless $D_>$ is empty. The latter occurs only if the set of directions where $D = 0$ (and hence $Lk_>(v, \overline{W}(\underline{K}))$) is contained in a proper face of the cell, and we are done by induction on the dimension of the cell.

If D is positive on some direction in $Lk(v, \overline{W}(\underline{K}))$ then by Lemma 4.2(4) either $a_{0,1} - a_{1,1} < 0$ and the direction obtained by scaling up $x_{1,1}$ at the expense of $x_{0,1}$ is in $Lk_>(v, \overline{W}(\underline{K}))$ or one of $a_{0,j} - a_{1,1} > 0$ for some $j > 0$ and the direction obtained by scaling up $x_{0,j}$ at the expense of $x_{0,1}$ is in $Lk_>(v, \overline{W}(\underline{K}))$. (Recall that $x_{0,1}$ doesn't increase along a direction at v whereas the other $x_{0,j}$ don't decrease). In the first case, the desired vertex is obtained from v by switching the roles of the edges $e_{0,1}$ and $e_{1,1}$ and in the latter case by switching the roles of $e_{0,1}$ and $e_{0,j}$. This finishes the first case.

Case 2. The derivative D is zero on $Lk(v, \overline{W}(\underline{K}))$.

It follows from Lemma 4.2(4) and Formula 3 that $a_{1,1} = a_{0,j}$ for all j. Let C be this common value. Then,

$$l_0^\alpha(x) = C + (1 - S(x_0))(l_0^{[\alpha]}(y) - C).$$

So,

$$Lk_>(l_0^\alpha) = \left(Lk(x_0, \Delta_{G\backslash C_1}) * Lk_>(l_0^{[\alpha]})\right) \setminus \left(Lk(x_0, \Delta_{G\backslash C_1}) * \emptyset\right)$$

and

$$Lk_\geq(l_0^\alpha) = Lk(x_0, \Delta_{G\backslash C_1}) * Lk_\geq(l_0^{[\alpha]}).$$

If $Lk_>(l_0^\alpha)$ is non-empty then $Lk_>(l_0^{[\alpha]})$ is non-empty in this case. By induction, $Lk_>(l_0^{[\alpha]})$ and $Lk_\geq(l_0^{[\alpha]})$ are contractible and differ by a Z-set in $Lk_\geq(l_0^{[\alpha]})$. By Lemma 4.7(3), $Lk_>(l_0^\alpha)$ and $Lk_\geq(l_0^\alpha)$ differ by a Z-set. Since $Lk_\geq(l_0^{[\alpha]})$ is contractible so is $Lk_\geq(l_0^\alpha)$ by Lemma 4.7(1). By induction, the desired edge direction can be found in $\emptyset * Lk_>(l_0^{[\alpha]})$.

If $Lk_>(l_0^\alpha)$ is empty, then $Lk_>(l_0^{[\alpha]})$ is empty. By induction, $Lk_>(w, \overline{W}(\underline{K}'))$ where $\underline{K}' = (C_1, C_2, \cdots)$ is contained in a proper face $\overline{W}'$ of $\overline{W}(\underline{K}')$. Then $Lk_>(v, \overline{W}(\underline{K}))$ is contained in the proper face $Lk(x_1, \Delta_{G\backslash C_1}) * Lk(w, \overline{W}')$. This completes the proof of the claim.

To finish the proof of the proposition, we now show that the proposition follows from the claim. Again we argue by induction on the number of edges in G and the dimension of the cell. If $Lk_>(l_k^\alpha)$ is non-empty we are done by the claim. If $Lk_>(l_k^\alpha)$ is empty then $Lk_>(v, \overline{W}(\underline{K}))$ is contained in a proper face and we are done by induction. □

If v is a vertex, denote by $\mathcal{L}(v)$ the poset of the cover of $Lk(v, BCV(N, K))$ by closed simplices $Lk(v, \overline{W}(\underline{K}))$, and by $\mathcal{L}_>(v)$ the poset of the subcollection consisting of those $Lk(v, \overline{W}(\underline{K}))$ with the property that $Lk_>(v, \overline{W}(\underline{K}))$ is nonempty.

Notice that L sends each vertex to a point of $(\mathbb{R}_+^{\mathcal{A}})^{|N|}$ with all coordinates integral.

Proposition 4.9. *Suppose v is a vertex of $BCV(N, K)$ and $L(v) = t \in (\mathbb{R}_+^{\mathcal{A}})^{|N|}$.*

1. *$Lk_>(v, BCV(N, K))$ is homotopy equivalent to $\mathcal{L}_>(v)$.*
2. *$\mathcal{P}(L^{-1}[t, \cdot))$ is isomorphic to $\mathcal{P}(L^{-1}(t, \cdot))$ with a subcomplex $\{P_v\}$ coned off. The subcomplex P_v is isomorphic to $\mathcal{L}_>(v)$.*

Proof. 1. This follows directly from Propositions 2.20 and 4.8.

2. We can regard $\mathcal{P}(L^{-1}(t, \cdot))$ as a subcomplex of $\mathcal{P}(L^{-1}[t, \cdot))$. $\mathcal{L}_>(v)$ naturally embeds as a subcomplex P_v of $\mathcal{P}(L^{-1}(t, \cdot))$. The vertex v gives rise to a vertex c_v of $\mathcal{P}(L^{-1}[t, \cdot))$ not contained in $\mathcal{P}(L^{-1}(t, \cdot))$. Now observe that the cone $c_v * P_v$ is a subcomplex of $\mathcal{P}(L^{-1}[t, \cdot))$ and that

$$\mathcal{P}(L^{-1}[t, \cdot)) = \mathcal{P}(L^{-1}(t, \cdot)) \cup (c_v * \mathcal{L}(v))$$

which collapses onto $\mathcal{P}(L^{-1}(t, \cdot)) \cup (c_v * P_v)$. $\square$

Theorem 4.10. *Each ascending link $Lk_>(v, BCV(N, K))$ is homotopy equivalent to the wedge of finitely many spheres of dimension $2|N| + |K| - 3p - 1$ (where p is the number of components of the model graph).*

By the wedge of 0 spheres we mean a point, so the ascending link is possibly contractible. However, it is never empty (recall that we are assuming that each rank 1 component contains at least two labels, so the above integer is non-negative).

4.1. **Case $|N| = p = 1$.** We start by verifying the theorem in case $|N| = p = 1$. Recall from Section 2.5 that $BCV(1, k)$ is identified with $\mathbb{R}^k / span(1, 1, \cdots, 1) \cong \mathbb{R}^{k-1}$. We claim that $Lk_<(v, BCV(1, k))$ is convex subset of $Lk(v, BCV(1, k)) \cong S^{k-2}$ that contains no antipodal points. It follows that the complement $Lk_>(v, BCV(1, k))$ of $Lk_<(v, BCV(1, k))$ is either S^{k-2} or contractible. Since $\mathbb{R}^{k-1}$ is contractible, it further follows from standard Morse theory that $Lk_>(v, BCV(1, k))$ is contractible except at the unique minimum. We now prove the claim.

Let θ_j be the k^{th} root of unity $e^{2\pi j/k}$. Take for the model graph $\mathcal{G} = \mathcal{G}(1, k)$ the unit circle with labeling $\{1, \cdots, k\} \to \mathcal{G}$ given by $j \mapsto \theta_j$. Let $\beta_j : I \to \mathcal{G}$ the homotopy class with immersed image a circular arc from θ_{j-1} to θ_j (subscripts taken mod k). Then $l_0^{\beta_j}(x) = |x_j - x_{j-1}|$. For general $\alpha \in \mathcal{A}$, l_0^α is a finite sum of such functions. Hence, each coordinate of the Morse function L (which equals L_0 in this case) is convex.

Let $L_<$ denote the subset of points of $BCV(1,k)$ with smaller L-value than v, i.e. the set of $x \in BCV(1,k)$ such that $L(x) < L(v)$. Since each coordinate of L is convex, $L_<$ is convex. Since $v \notin L_<$, $Lk_<(v, BCV(1,k))$ contains no points that are antipodal. The claim follows.

Remark 4.11. The set $L_<$, and hence also $Lk_<(v, BCV(1,k))$ and $Lk_>(v, BCV(1,k))$, need be neither open nor closed.

4.2. **Necklaces and the inward normal.** Consider a vertex $v \in BCV(N,K)$. There is a unique minimal $\hat{\Sigma}_H \subset BCV(N,K)$ that contains v: each component of H is a rose. This $\hat{\Sigma}_H$ is contained in all other $\hat{\Sigma}_G$ that contain v.

Let e_0 be the unique edge of H with $\ell_H(e_0) \neq 0$, and let e_1 be the unique edge with $\ell_{H_1}(e_1) \neq 0$. Consider the 1-cell $\overline{W}(\underline{K})$ determined by $\underline{K} = (H, H_2, H_3, \cdots)$ (only H_1 is omitted from the hierarchy of scales of v). This 1-cell determines a direction in $Lk(v, BCV(N,K))$, called the *inward normal* at v.

We say that an $\mathbb{R}$-graph G is a *necklace* if the graph obtained from G by collapsing all edges of length 0 is homeomorphic to the circle plus a collection of isolated points. By $Neck(\hat{\Sigma}_G)$ denote the subset of $\hat{\Sigma}_G$ consisting of those points that are taken by $\pi : \hat{\Sigma}_G \to \Sigma_G$ to a necklace. This is a subcomplex of $\hat{\Sigma}_G$. Let $Neck = Neck(N,K)$ denote the union of $Neck(\hat{\Sigma}_G)$ over all marked graphs G of type (N,K). Then $Neck$ is a subcomplex of $BCV(N,K)$ that contains all vertices. In particular, $Neck$, $Lk(v, Neck)$, $Lk_>(v, Neck)$, and $Lk_<(v, Neck)$ are all ANR's. Recall that for a point $\{\ell_C\}_{C \in \mathcal{C}(G)} \in \hat{\Sigma}_G$ we refer to ℓ_G as the *topmost metric*.

Lemma 4.12. *Denote by $\beta_1, \cdots, \beta_{|N|-1}$ the homotopy classes of the embedded loops at v of topmost length 0. A point of $BCV(N,K)$ near v is a necklace iff each β_i has topmost length 0.*

Proof. Let G be a graph obtained from v by blowing up a forest. Homotopy classes β_i are represented by embedded loops.

Suppose $\{\ell_C\}_C \in \hat{\Sigma}_G$ represents a necklace. If $\ell_G(\beta_i) > 0$, then $\ell_G(\beta_i) = 1$, and this is impossible in a small neighborhood of v.

Conversely, suppose $\{\ell_C\}_C \in \hat{\Sigma}_G$ does not represent a necklace. Then there are two homology classes γ_1 and γ_2 in $H_1(G; \mathbb{Q})$ such that no nontrivial linear combination $t_1\gamma_1 + t_2\gamma_2$ is represented by a cycle of ℓ_G-length 0. By β_0 denote the homotopy class represented by an embedded loop in v of length 1. Since $\beta_0, \beta_1, \cdots, \beta_{|N|-1}$ form a basis in homology, there is a nontrivial linear combination $t_1\gamma_1 + t_2\gamma_2$ which has the form $s_1\beta_1 + \cdots + s_{|N|-1}\beta_{|N|-1}$, and so not all β_i (for $i > 0$) can have length 0. □

Proposition 4.13. 1. *If the inward normal is in $Lk_>(v, BCV(N,K))$ (i.e. $w(e_1) > w(e_0)$), then $Lk_>(v, BCV(N,K))$ is contractible.*

2. *If the inward normal is not in* $Lk_>(v, BCV(N,K))$ *(i.e.* $w(e_1) < w(e_0)$*), then* $Lk_>(v, BCV(N,K))$ *is homotopy equivalent to* $Lk_>(v, Neck)$.

Proof. 1. The intersection-closed cover of $Lk_>(v, BCV(N,K))$ by the sets of the form $Lk_>(v, \hat{\Sigma}_G)$ has contractible elements and its poset is a cone (there is a unique minimal element corresponding to G being the disjoint union of roses).

2. Consider the intersection-closed cover of $Lk_>(v, BCV(N,K))$ by nonempty sets of the form $Lk_>(v, \hat{\Sigma}_G)$ and the induced intersection-closed cover of the ascending link $Lk_>(v, Neck)$. First notice that both covers consist of contractible sets. For the first cover this follows from Proposition 4.8. As for the second, it follows from Lemma 4.12 that if G is obtained from v by blowing up a forest, then there is a cell $\overline{W}(\underline{K})$ (consisting of those points in $\hat{\Sigma}_G$ where $\beta_1, \cdots, \beta_{|N|-1}$ have topmost length 0) such that $U \cap \hat{\Sigma}_G \cap Neck = U \cap \overline{W}(\underline{K})$ for a small neighborhood U of v. Hence Proposition 4.8 can be applied again.

It remains to show that the two covers have the same poset. To this end, we prove that if $Lk_>(v, \hat{\Sigma}_G)$ is nonempty, then $Lk_>(v, Neck \cap \hat{\Sigma}_G)$ is also nonempty. As usual, denote by e_0 the edge of G whose topmost length at v is 1.

Case 1. There is an edge $e \subseteq G$ such that $w(e) > w(e_0)$ and $e \not\subseteq core(G \setminus e_0)$.

Then $Lk_>(v, Neck \cap \hat{\Sigma}_G)$ is nonempty: the direction determined by $\ell_G(e_0) = 1 - t$, $\ell_G(e) = t$ for $t \geq 0$ (and $\ell_{G_i}(e')$ constant in all other cases) is ascending and is contained in $Neck$.

Case 2. There is no such e.

It then follows from Proposition 4.8(2) that all directions in which the topmost metric ℓ_G at v changes are descending. In particular, all ascending directions are contained in the necklace space, and so in this case $Lk_>(v, \hat{\Sigma}_G) = Lk_>(v, Neck \cap \hat{\Sigma}_G)$. ∎

Proposition 4.14. *Assume that Theorem 4.10 holds for total rank* $< |N|$. *Then* $Lk_>(v, Neck(N,K))$ *is homotopy equivalent to the wedge of finitely many spheres of dimension* $2|N| + |K| - 3p - 1$.

Theorem 4.10 (for total rank $|N|$) follows from Proposition 4.13 and Proposition 4.14. Our next goal is to formulate a proposition (Proposition 4.17) from which Proposition 4.14 will follow. After reordering the components we may assume that e_0 is in the first component.

Case: The first component is a circle. If the first component is the circle, then the component of $Neck$ containing v can be identified with

$$\mathbb{R}^{k_1 - 1} \times BCV((n_2, n_3, \cdots, n_p), (k_2, k_3, \cdots, k_p))$$

with v corresponding to (x, v'). For convenience, we denote

$$BCV((n_2, n_3, \cdots, n_p), (k_2, k_3, \cdots, k_p))$$

by BCV'. The function L induces a Morse function on BCV' by removing L_0 and the components of the homotopy classes contained in the first component. It is with respect to this function that we write e.g. $Lk_>(v', BCV')$. The link $Lk(v, Neck)$ can be identified with the join $S^{k_1-2} * Lk(v', BCV')$. The first component L_0 of the Morse function factors through $\mathbb{R}^{k_1-1}$. If x corresponds to the absolute minimum of L_0, then the ascending link $Lk_>(v, Neck)$ is the complement

$$\Big(S^{k_1-2} * Lk(v', BCV')\Big) \setminus \Big(\emptyset * Lk_<(v', BCV')\Big).$$

That this complement is homotopy equivalent to $S^{k_1-2} * Lk_>(v', BCV')$ follows from the next two lemmas.

Lemma 4.15. *Let A be an ANR, Z be a Z-set in A, and let $Y \subset Z$. Then $A \setminus Y$ is an ANR.*

Proof. This follows from [Han51]. For the sake of completeness, we give a proof in the case that A is finite dimensional. Recall that a finite dimensional space is an ANR iff it is locally contractible. Since $A \setminus Y \subset A$, $A \setminus Y$ is finite dimensional. Since $A \setminus Z \subset A \setminus Y \subset A$ is a Z-sandwich, it follows from the Z-sandwich fact in Section 1.3 that $A \setminus Y$ is also locally contractible. □

Lemma 4.16. *Let X and $\mathcal{F}$ be as in Proposition 2.20. Suppose also that X is finite dimensional. Then*

$$Y = \big(X * Lk(v, BCV(N, K))\big) \setminus \big(\emptyset * Lk_<(v, BCV(N, K))\big)$$

is homotopy equivalent to $Z = X * Lk_>(v, BCV(N, K))$.

Proof. The covers of Y and Z induced by sets of the form $F_i * Lk(v, \overline{W}(\underline{K}))$, $F_i * \emptyset$, and $\emptyset * Lk(v, \overline{W}(\underline{K}))$ are finite intersection closed, have finite order, have identical posets, and by Proposition 4.8 have contractible elements. Thus Y and Z are homotopy equivalent by Proposition 2.20 if the elements of the covers are ANRs. That this is the case follows from Corollary 4.6 and Lemmas 4.7(1) and 4.15. □

Recalling that $k_1 \geq 2$ in this case and using Theorem 4.10 for BCV', we conclude that $Lk_>(v, Neck)$ is homotopy equivalent to the wedge of finitely many spheres of dimension

$$1 + (k_1 - 2) + 2(|N| - 1) + (|K| - k_1) - 3(p - 1) - 1 = 2|N| + |K| - 3p - 1.$$

We saw in Section 4.1 that if x does not correspond to the minimum, the ascending link D at x with respect to the function L_0 is contractible. The set $Lk_>(v, Neck)$ can be identified with the complement in $D * Lk(v', BCV')$ of $\emptyset * Lk_<(v', BCV')$. By Lemma 4.16, the latter set is homotopy equivalent to $D * Lk_>(v', BCV')$, and is therefore contractible.

General case. For the rest of the argument, we assume that the first component has rank > 1.

We start with an informal discussion. Let $G \in Neck$ be near v. The metric ℓ_G sees a circle with a certain number of distinguished points ("beads") plus some additional isolated points ("spare beads"). Some of the beads come from labels in G; these we color blue and retain the labels. Others come from collapsed edges of G; these are colored red. Each β_i of Lemma 3.5 that is contained in the first component determines a red bead; we label this red bead i. We thus have k_1 labeled blue beads and $n_1 - 1$ red beads on the circle (some beads may coincide). One of the edges in the circle has length close to 1; this edge can be identified with e_0. If we choose an orientation on e_0, then the circle is naturally oriented. Removing e_0 produces $n_1 + k_1 - 1$ points on an (oriented) line.

Each red bead represents a core graph of total rank $< |N|$. There are 2 additional labels in this core graph indicating where the corresponding positive length edges are to be attached.

We are thus led to the following definition. Let

$$\mathcal{R} = \mathbb{R}^{n_1-1} \times \mathbb{R}^{k_1}/span\ (1, 1, \cdots, 1) \cong \mathbb{R}^{k_1+n_1-2}$$

be the "configuration space" (as usual we mod out by simultaneous translation of all points).

Consider the partition of $\mathbb{R}^{n_1-1} \times \mathbb{R}^{k_1}$, and the induced partition of $\mathcal{R}$, into cones, as follows. Two points

$$(x_1, x_2, \cdots, x_{n_1-1}, y_1, \cdots, y_{k_1}) \text{ and } (x'_1, x'_2, \cdots, x'_{n_1-1}, y'_1, \cdots, y'_{k_1})$$

belong to the same cone in the partition provided

1. $x_i < x_j \iff x'_i < x'_j$ for all $i, j = 1, 2, \cdots, n_1 - 1$,
2. $x_i < y_j \iff x'_i < y'_j$ for all $i = 1, 2, \cdots, n_1 - 1$ and all $j = 1, 2, \cdots, k_1$, and
3. $x_i > y_j \iff x'_i > y'_j$ for all $i = 1, 2, \cdots, n_1 - 1$ and all $j = 1, 2, \cdots, k_1$.

Informally, one configuration is in the same cone as another if it can be transformed to the other by moving the points around and not allowing the red points to collide or uncollide with any other points. Blue points are allowed to pass through each other freely.

There is a natural map ϕ from a neighborhood of v in $Neck$ to $\mathcal{R}$, described above (collapse all edges of topmost length 0 and remove e_0). Pulling back by ϕ induces a partition of this neighborhood. We study the preimages of higher and higher skeleta inductively. The map ϕ is affine on closed cells (with respect to the structure induced by the chart at v). In particular, ϕ induces a partition of the ascending link at v. We show:

Proposition 4.17. *Assume that Theorem 4.10 is true when the total rank is $< |N|$. If $\mathcal{R}_0$ is a nonempty closed subset of $\mathcal{R}$ which is a union of cones in* *the partition, then the ascending link at v with respect to $\phi^{-1}(\mathcal{R}_0)$ is homotopy equivalent to the wedge of spheres of dimension $2|N| + |K| - 3p - 1$.*

Proposition 4.14 then follows by taking $\mathcal{R}_0 = \mathcal{R}$. The rest of this section is devoted to the proof of Proposition 4.17.

Proof of Proposition 4.17. We proceed by induction on the number of cones in $\mathcal{R}_0$. Consider first the subspace $Neck_0$ of $Neck$ consisting of those graphs where there is a unique edge of length 1 (with respect to the topmost metric). We claim that $Lk_>(v, Neck_0)$ is homotopy equivalent to the wedge of finitely many spheres of dimension $2|N| + |K| - 3p - 1$.

Indeed, (a neighborhood of v in) $Neck_0$ can be identified with (a corresponding neighborhood in)

$$BCV((n_1 - 1, n_2, \cdots, n_p), (k_1 + 2, k_2, \cdots, k_p))$$

and Theorem 4.10 applied to this situation says that $Lk_>(v, Neck_0)$ is homotopy equivalent to the wedge of finitely many spheres of dimension $2(|N|-1)+(|K|+2) - 3p - 1 = 2|N| + |K| - 3p - 1$. Thus the proposition is verified for the case $\mathcal{R}_0$ equals the cone point.

To illustrate the inductive step in the next easiest case, we next prove the proposition for the case when $\mathcal{R}_0$ is a ray.

A generic point on a ray has exactly two distinct coordinates. The ray $r(A, B)$ is given by a partition of the $k_1 + n_1 - 1$ coordinates into two nonempty subsets A and B and consists of equivalence classes of points in $\mathcal{R}$ where all coordinates in A coincide and are less than or equal to all coordinates in B, and these also coincide. The ray $r(A, B)$ is in the partition if and only if either each A and B contain a red coordinate or else one of A, B consists of a single blue coordinate.

Consider the former case first.

The preimage under ϕ of the ray minus the origin is the subset consisting of those elements whose first component from the perspective of the topmost metric looks like a circle of length 1 with points a and b close to each other, with a preceding b (with respect to the orientation of the circle induced by e_0), and with all loops and labeled points in A being at a and all loops and labeled points in B being at b. Denote by $Neck(A, B)$ the closure of the preimage of the open ray. There is a map $Neck(A, B) \to [0, \epsilon)$ that measures the distance between a and b. There is also a "connection" (homeomorphism between fibers of $Neck(A, B) \to [0, \epsilon)$) that changes the lengths of e_0 and of the short edge connecting a and b, and this induces a product structure $W(A, B) \times [0, \epsilon)$ in a neighborhood of v in $Neck(A, B)$. Say A contains k_A labeled points and n_A loops, and k_B, n_B are defined similarly for B. Thus $n_A + n_B = n_1 - 1$, $k_A + k_B = k_1$, and $n_A > 0$, $n_B > 0$.

A neighborhood of v in $W(A, B) \times 0$ can be identified with a neighborhood of $v(A, B)$ ($= H \setminus e_0$ and cut open according to A and B, i.e. the graph underlying

$v(A, B)$ is obtained from the disjoint union of roses representing v by removing the topmost edge in the first component and then dividing the rest of the first component into two components according to A and B) in

$$BCV((n_A, n_B, n_2, \cdots, n_p), (k_A + 2, k_B + 2, k_2, \cdots, k_p)).$$

(The extra 4 points keep track of where the long and the short edges are attached.) By induction, the ascending link $Lk_>(v(A, B), W(A, B) \times \{0\})$ is homotopy equivalent to the wedge of spheres of dimension

$$2(|N| - 1) + (|K| + 4) - 3(p + 1) - 1 = 2|N| + |K| - 3p - 2.$$

If the "inward normal" at v corresponding to the connection above (i.e. the direction from v in which the length of e_0 goes down at the expense of the distance between a and b) is not in the ascending link at v, then the ascending link at v with respect to the preimage of the closed ray is homotopy equivalent to the ascending link of v in the preimage of the origin, and this is homotopy equivalent to the wedge of spheres of dimension $2|N| + |K| - 3p - 1$, as seen above. If this inward normal is in the ascending link, then the ascending link at v with respect to the preimage of the ray is obtained from the ascending link with respect to the preimage of the origin by coning off a subset which is homotopy equivalent to the wedge of spheres of one less dimension, and so is homotopy equivalent to the wedge of spheres of dimension $2|N| + |K| - 3p - 1$.

Now consider the case where B consists of a single blue coordinate. A product structure $W(A, B) \times [0, \epsilon)$ in a neighborhood of v is defined similarly. Then a neighborhood of v in $W(A, B) \times 0$ can be identified with a neighborhood of $v(A, B)$ $(= H \setminus e_0)$ in

$$BCV((n_1 - 1, n_2, \cdots, n_p), (k_1 + 1, k_2, \cdots, k_p)).$$

The arithmetic now is

$$2(|N| - 1) + (|K| + 1) - 3p - 1 = 2|N| + |K| - 3p - 2.$$

The argument in general is similar. Suppose $\dim \mathcal{R}_0 = l$.

Choose an l-dimensional open cone C in $\mathcal{R}_0$. By induction, Proposition 4.17 holds for $\mathcal{R}_0 \setminus C$. By $V(C)$ denote the closure of $\phi^{-1}(C)$. The map ϕ induces a map $V(C) \to \overline{C}$, which in turn induces a product structure $W(C) \times \overline{C}$ in a neighborhood of v in $\phi^{-1}(\mathcal{R}_0)$.

A generic point in C has $l+1$ distinct coordinates and each equivalence class either contains a red bead, or else consists of a single blue bead. Say there are $l(r)$ classes of the first kind, and $l(b)$ of the second, so that $l(r) + l(b) = l + 1$. Denote by $n(j)$ and $k(j)$ the number of red and blue beads respectively in the j^{th} equivalence class of the first kind, $j = 1, 2, \cdots, l(r)$.

A neighborhood of v in $W(C) \times \{0\}$ can be identified with a neighborhood of a certain graph $v(C)$ (obtained from $H \setminus e_0$ by cutting open and assigning labeled points according to C) in

$$BCV((n(1),\cdots,n(l(r)),n_2,\cdots,n_p),\\ (k(1)+2,\cdots,k(l(r))+2,k_2,\cdots,k_p)).$$

The ascending link with respect to $W(C)\times 0$ is therefore homotopy equivalent to the wedge of spheres of dimension

$$2(|N|-1)+(|K|-l(b)+2l(r))-3(p-1+l(r))-1=\\ 2|N|+|K|-3p-l(b)-l(r)=2|N|+|K|-3p-1-l.$$

Consider the commutative diagram

$$\begin{array}{ccc} W(C)\times\overline{C} & \xrightarrow{L} & (\mathbb{R}_+^{\mathcal{A}})^{|N|} \\ \downarrow & & \downarrow \\ \overline{C} & \xrightarrow{\ell} & \mathbb{R}_+^{\mathcal{A}} \end{array}$$

in which the vertical arrows are projections (the right one projects to the first $\mathcal{A}$ coordinates). The function $\ell:\overline{C}\to\mathbb{R}_+^{\mathcal{A}}$ measures the length of α_i's with respect to the topmost metric.

We consider two cases.

Case 1. The subset $Lk_<(0,C)$ of $Lk(0,C)$ of directions at 0 in C along which ℓ is not strictly increasing is not empty.

As in Section 4.1, $Lk_<(0,C)$ is convex and contractible. It follows that in this case the radial projection (with respect to the spherical metric) from a direction in $Lk_<(0,C)$ induces a strong deformation retraction from $Lk_>(v,W(C)\times\overline{C})$ onto $Lk_>(v,W(C)\times(\overline{C}\setminus C))$, and thus from $Lk_>(v,\phi^{-1}(\mathcal{R}_0))$ onto $Lk_>(v,\phi^{-1}(\mathcal{R}_0\setminus C))$, and this proves the proposition in this case.

Case 2. $Lk_>(0,C)=Lk(0,C)$.

We first claim that $Lk_>(v,W(C)\times\overline{C})$ is contractible. Indeed, consider inclusions

$$\big(Lk(v(C),W(C))*Lk(0,C)\big)\setminus\big(Lk(v(C),W(C))*\emptyset\big)\\ \subseteq Lk_>(v,W(C)\times\overline{C})\subseteq Lk(v(C),W(C))*Lk(0,\overline{C}).$$

By Lemma 4.7(1), since $Lk(0,C)$ is contractible, $Lk(v(C),W(C))*\emptyset$ is a Z-set in

$$Lk(v(C),W(C))*Lk(0,C).$$

Also, since $Lk(0,C)$ and $Lk(0,\overline{C})$ differ by a Z-set, the above sequence of inclusions is a Z-sandwich by Lemma 4.7(3). Hence we see that $Lk_>(v,W(C)\times\overline{C})$ is contractible.

Secondly, we claim that inclusion

$$Lk_>(v(C),W(C))*Lk_>(0,\overline{C}\setminus C)\hookrightarrow Lk_>(v,W(C)\times(\overline{C}\setminus C))$$

is a homotopy equivalence. Indeed,

$$\begin{aligned}&Lk_>(v, W(C) \times (\overline{C} \setminus C)) \\ &= \big(Lk(v(C), W(C)) * Lk_>(0, \overline{C} \setminus C)\big) \setminus \big(Lk_<(v(C), W(C)) * \emptyset\big)\end{aligned}$$

which by Lemma 4.16 is homotopy equivalent to

$$Lk_>(v(C), W(C)) * Lk_>(0, \overline{C} \setminus C).$$

As in Section 4.1, from the convexity of ℓ it follows that $Lk_>(0, \overline{C} \setminus C)$ is either contractible (and nonempty!) or the whole $Lk(0, \overline{C} \setminus C) = S^{l-2}$. In the first case, $Lk_>(v, W(C) \times (\overline{C} \setminus C))$ is contractible, and in the second it is homotopy equivalent to the wedge of finitely many spheres of dimension $2|N| + |K| - 3p - 1 - l + (l-2) + 1 = 2|N| + |K| - 3p - 2$. Now $Lk_>(v, \phi^{-1}(\mathcal{R}_0))$ is obtained from $Lk_>(v, \phi^{-1}(\mathcal{R}_0 \setminus C))$ (which is homotopy equivalent to the wedge of $2|N| + |K| - 3p - 1$-spheres) by attaching $Lk_>(v, W(C) \times \overline{C})$ (which is contractible) along $Lk_>(v, W(C) \times (\overline{C} \setminus C))$ (which is homotopy equivalent to the wedge of $(2|N| + |K| - 3p - 2)$-spheres). The proposition finally follows. □

5. The topology at infinity

To study the topology at infinity of $BCV(N, K)$ it is more convenient to replace $BCV(N, K)$ by the simplicial complex $\mathcal{S}(N, K)$, which is proper homotopy equivalent to $BCV(N, K)$.

Definition 5.1. The simplicial complex $\mathcal{S}(N, K)$ is defined to have the same vertex set as $BCV(N, K)$, and a collection $v_0, v_1, \cdots, v_k$ of vertices spans a simplex iff they are contained in some $\hat{\Sigma}_G \subset BCV(N, K)$.

Proposition 5.2.
1. *$\mathcal{S}(N, K)$ is proper homotopy equivalent to $BCV(N, K)$.*
2. *Let $v \in \mathcal{S}(N, K)$ be a vertex and by $Lk_\uparrow(v, \mathcal{S}(N, K))$ denote the link of v in the subcomplex of $\mathcal{S}(N, K)$ spanned by the vertices w with $L(w) \geq L(v)$. Then $Lk_\uparrow(v, \mathcal{S}(N, K))$ is homotopy equivalent to $Lk_>(v, BCV(N, K))$, and thus to the wedge of finitely many spheres of dimension $2|N| + |K| - 3p - 1$.*

Proof. The function that to each $\hat{\Sigma}_G$ assigns the simplex in $\mathcal{S}(N, K)$ spanned by the vertices in $\hat{\Sigma}_G$ induces an isomorphism between the posets of intersection-closed covers by compact sets, and thus (1) follows from Proposition 2.20.

The proof of (2) is similar: Consider the covering of $Lk_>(v, BCV(N, K))$ by nonempty sets of the form $Lk_>(v, \hat{\Sigma}_G)$. To each such set assign the simplex in $Lk_\uparrow(v, \mathcal{S}(N, K))$ whose vertices are those $w \in \hat{\Sigma}_G$ with $L(w) > L(v)$. That this correspondence induces a poset isomorphism follows from Proposition 4.8, and the claim then follows from Proposition 4.9. □

The following theorem is the analog of Lemma 1.3 in our situation.

Theorem 5.3. *Suppose $\mathcal{S}$ is a locally finite r-connected simplicial complex such that*

1. *There is an injective function $L : \mathcal{S}^{(0)} \to \Lambda$ from the vertex set of $\mathcal{S}$ to a totally ordered set Λ.*
2. *There is a collection Δ of simplices in $\mathcal{S}$, called "distinguished", that is closed under intersections and such that every simplex in $\mathcal{S}$ is contained in a distinguished simplex.*
3. *There is a function $\mu : \mathcal{S}^{(0)} \to \mathbb{N} = \{1, 2, 3, \cdots\}$ which is "proper", i.e. $\mu^{-1}(n)$ is finite for all $n \in \mathbb{N}$, and compatible with L in the following sense: If σ is a distinguished simplex and v the vertex of σ with maximal L-value, then $\mu(v) \geq \mu(w)$ for all vertices w of σ.*
4. *For all vertices $v \in \mathcal{S}^{(0)}$ the ascending link $Lk_{\uparrow}(v, \mathcal{S})$ is r-connected (and nonempty). Here the ascending link is the link of v in the subcomplex of $\mathcal{S}$ spanned by the vertices w with $L(w) \geq L(v)$.*

Then $\mathcal{S}$ is r-connected at infinity.

Proof. The first step is to reduce to the case where $\Lambda = \mathbb{N}$ (and then μ can be taken to be L). Consider $\mu \times L : \mathcal{S}^{(0)} \to \mathbb{N} \times \Lambda$ where we order $\mathbb{N} \times \Lambda$ lexicographically. This map is injective and properness of μ implies that the image $(\mu \times L)(\mathcal{S}^{(0)}) \subseteq \mathbb{N} \times \Lambda$ is order isomorphic to $\mathbb{N}$ (assuming without loss of generality that it is infinite). Denote by $Lk'_{\uparrow}(v, \mathcal{S})$ the ascending link with respect to $\mu \times L$, i.e. the link of v in the subcomplex of $\mathcal{S}$ spanned by the vertices w with $(\mu \times L)(w) \geq (\mu \times L)(v)$. We will argue that $Lk_{\uparrow}(v, \mathcal{S})$ and $Lk'_{\uparrow}(v, \mathcal{S})$ are homotopy equivalent, and our first step will then be complete. Note that the compatibility assumption says that L and $\mu \times L$ attain maximal value among the vertices of a distinguished simplex at the same vertex. Thus, if σ is a distinguished simplex that contains v, then $Lk(v, \sigma)$ intersects $Lk_{\uparrow}(v, \mathcal{S})$ if and only if it intersects $Lk'_{\uparrow}(v, \mathcal{S})$. If these intersections are nonempty, then they are contractible (they are simplices). Therefore, the closed covers of $Lk_{\uparrow}(v, \mathcal{S})$ and $Lk'_{\uparrow}(v, \mathcal{S})$ consisting of nonempty intersections with $Lk(v, \sigma)$ (as σ ranges over the distinguished simplices that contain v) are both intersection-closed and consist of AR's. Since they have the same poset, it follows from Proposition 2.20 that $Lk_{\uparrow}(v, \mathcal{S}) \simeq Lk'_{\uparrow}(v, \mathcal{S})$.

For the rest of the proof we assume that $\Lambda = \mathbb{N}$ and that $L : \mathcal{S}^{(0)} \to \mathbb{N}$ is a bijection. For $m \in \mathbb{N}$ denote by $\mathcal{S}_m$ the subcomplex of $\mathcal{S}$ spanned by the vertices w with $L(w) \geq m$. Thus $\mathcal{S}_0 = \mathcal{S}$ and $\mathcal{S}_m$ is obtained from $\mathcal{S}_{m+1}$ by coning off $Lk_{\uparrow}(v, \mathcal{S})$. From the Mayer-Vietoris sequence we see that $\mathcal{S}_m$ and $\mathcal{S}_{m+1}$ have isomorphic homology groups in dimensions $\leq r$. Since $\mathcal{S}_0 = \mathcal{S}$ is r-connected, it follows that homology groups of $\mathcal{S}_m$ vanish in dimensions $\leq r$ for all m. If $r \geq 1$ we see similarly from van Kampen's theorem that each $\mathcal{S}_m$ is simply connected. Therefore, Hurewicz theorem implies that each $\mathcal{S}_m$ is r-connected. Since the closure of $\mathcal{S} \setminus \mathcal{S}_m$ is compact for all m and every compact subset of $\mathcal{S}$ is disjoint from $\mathcal{S}_m$ for large m it follows that $\mathcal{S}$ is r-connected at infinity. $\square$

Proposition 5.4. *Under the hypotheses of Theorem 5.3 and assuming in addition that $\mathcal{S}$ is contractible and that each ascending link is the wedge of $(r+1)$-spheres, we have that compactly supported cohomology*

$$H^i_c(\mathcal{S}) \cong \begin{cases} 0 & \text{if } i \neq r+2 \\ \bigoplus_{v \in \mathcal{S}^{(0)}} \tilde{H}^{r+1}(Lk_\uparrow(v, \mathcal{S})) & \text{if } i = r+2 \end{cases}$$

Proof. The group $H^i_c(\mathcal{S})$ is by definition the direct limit of the sequence $H^i(\mathcal{S}, \mathcal{S}_m) \cong \tilde{H}^{i-1}(\mathcal{S}_m)$. Argue by induction on m that

$$\tilde{H}^j(\mathcal{S}_m) \cong \begin{cases} 0 & \text{if } j \neq r+1 \\ \bigoplus_{v \notin \mathcal{S}^{(0)}_m} \tilde{H}^{r+1}(Lk_\uparrow(v, \mathcal{S})) & \text{if } j = r+1 \end{cases}$$

The inductive step consists of writing down the Mayer-Vietoris sequence for

$$\mathcal{S}_m = \mathcal{S}_{m+1} \cup_{Lk_\uparrow(v,\mathcal{S})} c(Lk_\uparrow(v, \mathcal{S}))$$

and noticing that the portion

$$0 \to \tilde{H}^{r+1}(\mathcal{S}_m) \to \tilde{H}^{r+1}(\mathcal{S}_{m+1}) \to \tilde{H}^{r+1}(Lk_\uparrow(v, \mathcal{S})) \to 0$$

splits. $\square$

Lemma 5.5. *There are finitely many homotopy classes $\mathcal{B}(N, K)$ in $\mathcal{G}(N, K)$ such that if some element $\alpha \in Out(N, K)$ fixes each of them then α is the identity.*

Proof. In the case $Out(F_n =< x_1, \cdots, x_n >) = Out(n, 0)$, the set $\mathcal{B}(n, 0) = \{x_1, \cdots, x_n, x_1x_2, x_1x_3, \cdots, x_1x_n, x_2x_3, x_2x_4, \cdots, x_2x_n\}$ suffices. (Proof: Let $f : F_n \to F_n$ be an automorphism that fixes all these elements up to conjugacy. By conjugating f, we may assume $f(x_1) = x_1$. After a further conjugation by a power of x_1, we may assume that $f(x_2) = \gamma x_2 \gamma^{-1}$ is a reduced word that does not start with a nontrivial power of x_1. Then $f(x_1x_2) = x_1\gamma x_2\gamma^{-1}$ and x_1x_2 can be fixed up to conjugacy only if $\gamma = 1$. Now if $i > 2$ and $f(x_i) = \gamma_i x_i \gamma_i^{-1}$ is a reduced word, then γ_i does not start either with $x_1^{\pm 1}$ or $x_2^{\pm 1}$. In the former case, $f(x_1x_i) = x_1\gamma_i x_i \gamma_i^{-1}$ and x_1x_i can be fixed up to conjugacy only if $\gamma_i = 1$. In the latter case, we consider x_2x_i. Therefore, $f = id$.)

In the case of $Aut(F_n) = Out(n, 1)$ we can take $\mathcal{B}(n, 1) = \{x_1, \cdots, x_n\}$. For $Out(n, k)$ take $\mathcal{B}(n, k) = \mathcal{B}(n, 1) \cup \{y_2, \cdots, y_k\}$ where y_i is a path from the first to the i^{th} labeled point. The general case follows by choosing the classes as above in each component. $\square$

Theorem 5.6. *$\mathcal{S}(N,K)$ is $(2|N|+|K|-3p-2)$-connected at infinity. In particular, $Out(F_n)$ is $(2n-5)$-connected at infinity. Moreover,*

$$H^i_c(\mathcal{S}(N,K)) \cong \begin{cases} 0 & \text{if } i \neq 2|N|+|K|-3p \\ \text{free abelian} & \text{if } i = 2|N|+|K|-3p \end{cases}$$

Proof. We apply Theorem 5.3 with $\mathcal{S}=\mathcal{S}(N,K)$ and $r=2|N|+|K|-3p-2$. The function L is the Morse function with respect to an ample sequence with the first homotopy class α_1 equal to the disjoint union of the homotopy classes from Lemma 5.5. The function μ is the multiplicity with respect to α_1, i.e. $\mu(v)$ is the maximal number of times the class α_1 crosses an edge of the underlying rose when represented by an immersion. Properness of μ is a consequence of Lemma 5.5, and assumption (4) of Theorem 5.3 comes from Proposition 5.2. A simplex of $\mathcal{S}$ is "distinguished" if its vertices form the entire vertex set of some $\hat{\Sigma}_G$. To see that the compatibility condition (3) of Theorem 5.3 is satisfied note that the very first coordinate of $L(v)$ (where v is a vertex of $\hat{\Sigma}_G$) is the multiplicity with which α_1 crosses the unique edge of length 1 (topmost edge) in v. Thus the largest value of L among the vertices of $\hat{\Sigma}_G$ will occur at a vertex v where this multiplicity equals the largest multiplicity with which α_1 crosses an edge of G, and in particular $\mu(v)\geq\mu(w)$ for all other vertices w. Note also that the compatibility condition fails for arbitrary simplices in $\mathcal{S}(N,K)$ since the maximal multiplicity might occur on an edge that is never topmost at the vertices of a non-distinguished simplex. The other properties are immediate. The last sentence follows from Proposition 5.4. □

References

[AB87] R. Alperin and H. Bass, *Length functions of groups actions on Λ-trees*, Combinatorial Group Theory and Topology (S.M. Gersten and J.R. Stallings, eds.), Annals of Math. Studies, vol. 111, Princeton Univ. Press, 1987.

[AM69] M. F. Atiyah and I. G. MacDonald, *Introduction to Commutative Algebra*, Addison-Wesley, 1969.

[BB97] M. Bestvina and N. Brady, *Morse theory and finiteness properties of groups*, Invent. Math **129** (1997), no. 3, 445–470.

[BE73] R. Bieri and B. Eckmann, *Groups with homological duality and Poincaré duality*, Invent. Math **20** (1973), 103–124.

[Bor67] K. Borsuk, *Theory of retracts*, Monografie Matematyczne, Polish Scientific Publishers, Warsaw, 1967.

[Bro82] K. Brown, *Cohomology of groups*, GTM, vol. 87, Springer-Verlag, New York, 1982.

[BS73] A. Borel and J.-P. Serre, *Corners and arithmetic groups. avec un appendice: Arrondissment de variétés á coins, par A. Douady and L. Hérault*, Comment Math. Helv. **48** (1973), 436–491.

[BT68] G. Baumslag and T. Taylor, *The centre of groups with one defining relator*, Math. Ann. **175** (1968), 315–319.

[CV86] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), 91–119.

[GM85] R. Geoghegan and M. Mihalik, *Free abelian cohomology of groups and ends of universal covers*, J. Pure and Appl. Algebra **36** (1985), 123–137.

[Han51] Olof Hanner, *Some theorems on absolute neighborhood retracts*, Ark. Mat. **1** (1951), 389–408.

[Har86] J. Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. **84** (1986), 157–176.

[Hat95] A. Hatcher, *Homological stability for automorphism groups of free groups*, Comment. Math. Helv. **70** (1995), 39–62.

[Hen75] D. Henderson, *Z-sets in ANR's*, Trans. Amer. Math. Soc. **213** (1975), 205–216.

[Hir75] H. Hironaka, *Triangulations of algebraic sets*, Algebraic Geometry, Proc. Sympos. Pure Math., vol. 29, Amer. Math. Soc., Providence, R.I., 1975, pp. 165–185.

[Lac77] R.C. Lacher, *Cell-like mappings and their generalizations*, Bull. Amer. Math. Soc. **83** (1977), 495–552.

[Loj64] S. Lojasiewicz, *Triangulation of semi-analytic sets*, Ann. Scu. Norm. di Pisa **18** (1964), 449–552.

[Ric] J. Rickert, *$Out(F_4)$ is simply-connected at infinity*, Ph.D. Thesis, Vanderbilt University, 1996.

[Sie59] Carl Ludwig Siegel, *Zur Reduktionstheorie quadratischer Formen*, The Mathematical Society of Japan, Tokyo, 1959, Publications of the Mathematical Society of Japan, Vol. 5.

[Sma57] S. Smale, *A Vietoris mapping theorem for homotopy*, Proc. Amer. Math. Soc. **8** (1957), 604–610.

[SV92] J. Smillie and K. Vogtmann, *Length functions and Outer Space*, Mich. Math. J. **39** (1992), 485–493.

[Vog90] K. Vogtmann, *Local structure of some $Out(F_n)$-complexes*, Proc. Edinburgh Math. Soc.(2) **33** (1990), 367–379.

[Vog95] K. Vogtmann, *End invariants of the group of outer automorphisms of a free group*, Topology **34** (1995), 533–545.

Department of Mathematics, University of Utah, Salt Lake City, UT 84112
E-mail address: `bestvinamath.utah.edu`

Department of Mathematics, Rutgers University, Newark, New Jersey 07102
E-mail address: `feighnandromeda.rutgers.edu`