Check for updates

J London Math Soc (2) 98 (2018) 12–34 © 2018 London Mathematical Society
doi:10.1112/jlms.12124

n the bordification of uter space

Kai-Uwe Bux, Peter Smillie and Karen Vogtmann

Abstract

We give a simple construction of an equivariant deformation retract of Outer space which is homeomorphic to the Bestvina–Feighn bordification. This results in a much easier proof that the bordification is 2n-5)-connected at infinity, and hence that $Out\ F_n)$ is a virtual duality group.

Introduction

The action of $SL(n,\mathbb{Z})$ on the symmetric space $X_n = SL(n,\mathbb{R})/SO(n)$ is not cocompact, but Siegel [14] showed how to glue affine spaces to X_n to obtain a contractible manifold with corners in such a way that the action extends to a proper cocompact action. In a landmark paper [3], Borel and Serre generalized this to all arithmetic groups Γ in reductive algebraic groups G, and it is now commonly referred to as the *Borel–Serre bordification* of the symmetric space G/K. In the case $G = SL_n$, Grayson [7] later showed how to construct an equivariant deformation retract of X_n with the same properties as the bordification. Leuzinger [11, 12] then defined a similar retract for more general groups, and said it is 'likely that this retract is isomorphic to the Borel–Serre bordification'. These retracts avoid many of the technical problems associated with extending the space and the action, and are generally much easier to understand.

The group $\mathrm{Out}(F_n)$ shares a large number of properties with arithmetic groups, many of which are proved by considering its action on *Outer space* $\mathcal{O}_n$, which serves as a substitute for the homogeneous space G/K. Motivated by the work of Borel and Serre, Bestvina and Feighn [2] constructed a contractible bordification $b\mathcal{O}_n$ such that the action of $Out(F_n)$ on $\mathcal{O}_n$ extends to a proper cocompact action on $b\mathcal{O}_n$. They used their bordification to prove that $Out(F_n)$ is a virtual duality group; the key further ingredient needed for this is to prove that $b\mathcal{O}_n$ is $(2n-5)$-connected at infinity.

In this paper we follow the lead of Grayson and Leuzinger by showing that there is an equivariant deformation retract of $\mathcal{O}_n$ which is cocompact and $(2n-5)$-connected at infinity. We show that this retract is equivariantly homeomorphic to the Bestvina–Feighn bordification and in the process answer a question in their paper about the topology of the pieces $\Sigma(G,g)$ from which their bordification is constructed. The description of the retract is simpler than that of the bordification. We note that our retract is equivalent to a retract sketched briefly without proof, discussed mainly for $n = 3$, in [1]. We use our description to give a different, considerably simpler proof of the connectivity result in [2]. In a sequel we will also use it to study the boundary.

1. *Background: Outer space*

In this section we briefly describe Outer space and its decomposition into open simplices $\mathring{\sigma}(G,g)$, in order to introduce the notation needed for this paper. For a somewhat more detailed quick introduction to these ideas, see [15], and for detailed proofs see the original paper [6].

Received 8 October 2017; revised 16 January 2018; published online 28 March 2018.

2010 *Mathematics Subject Classification* 20E36, 57M07 (primary), 20F65, 20E05 (secondary).

K. Vogtmann was partially supported by the Humboldt Foundation and a Royal Society Wolfson Award. K.-U. Bux was supported by the German Science Foundation via the CRC 701.

Outer space $\mathcal{O}_n$ is a contractible, $(3n-4)$-dimensional space with a proper action of the group $Out(F_n)$ of outer automorphisms of the free group F_n. A point of $\mathcal{O}_n$ is determined by a metric graph G together with a homotopy equivalence g, called a *marking*, from a fixed n-petaled rose R_n to G. The graphs G must be connected with no univalent or bivalent vertices, and we will also assume they have no separating edges; this is sometimes called *reduced* Outer space. The metric on G must have *volume* 1, that is, the sum of the edge lengths is equal to 1. The pair (G, g) is called a *marked graph*. Different marked graphs determine the same point of $\mathcal{O}_n$ if they are isometric by an isometry which commutes with the marking up to homotopy.

There is a natural decomposition of $\mathcal{O}_n$ as a disjoint union of open simplices. The simplex containing the point (G, g) is obtained by simply varying the (positive) edge lengths of G while keeping the volume equal to 1; it is denoted $\mathring{\sigma}(G, g)$ and its closure in $\mathcal{O}_n$ is denoted $\bar{\sigma}(G, g)$. If G has k edges, then $\mathring{\sigma}(G, g)$ is an open $(k-1)$-simplex. A simplex $\mathring{\sigma}(G\,, g\,)$ is a face of $\bar{\sigma}(G, g)$ if G can be obtained from G by shrinking some edges to points, and g is homotopic to g composed with the collapse. Note $\bar{\sigma}(G, g)$ is not a closed simplex since some of its faces are missing, namely those approached by shrinking all edges of a subgraph that contains loops.

The open simplices $\mathring{\sigma}(G, g)$ form a partially ordered set (*poset*) where the partial order is the face relation, that is, $\mathring{\sigma}(G\,, g\,) \leqslant \mathring{\sigma}(G, g)$ if $\mathring{\sigma}(G\,, g\,) \subseteq \bar{\sigma}(G, g)$. The geometric realization of this poset is called the *spine* K_n of $\mathcal{O}_n$. The spine has a natural embedding into $\mathcal{O}_n$ as an equivariant deformation retract (the case $n = 2$ is illustrated in Figure 3).

A key notion in the paper [2] is that of a *core subgraph* of G. By a *subgraph* we mean the closure in G of a set of edges. An (open) edge e of a subgraph H *separates* H if $H - e$ has an additional component, and a subgraph H is *core* if none of its edges separates H. In particular, G is a core subgraph of itself. In general core subgraphs need not be connected, and they may contain bivalent vertices (but not univalent vertices because removing the unique edge to a univalent vertex would separate the subgraph). Note that every subgraph of G contains a unique maximal core subgraph. If this maximal core is empty the subgraph is a union of trees, that is, a *forest* in G.

2. *Jewels and the retract*

In this section we find a compact cell $J(G, g)$ (a *jewel*) inside the closure $\bar{\sigma}(G, g)$ of $\mathring{\sigma}(G, g)$ in $\mathcal{O}_n$, then glue these cells together to form an equivariant deformation retract of $\mathcal{O}_n$. The construction of $J(G, g)$ will be independent of the marking g, so we temporarily eliminate the marking from our notation.

We consider $\mathring{\sigma}(G)$ to be the interior of a closed simplex $\sigma(G)$. The faces of $\sigma(G)$ which are not in $\bar{\sigma}(G)$ are said to be *at infinity*. We view $\sigma(G)$ as a regular Euclidean simplex; the lengths of edges in G give barycentric coordinates on $\sigma(G)$. A face at infinity is obtained by setting the edge lengths of H equal to 0 for some subgraph H that contains a loop. In particular all vertices of $\sigma(G)$ are at infinity, assuming $n > 1$. An equivalent way of saying this is that $\operatorname{rank}(H_1(G//H^c)) < \operatorname{rank}(H_1(G))$, where $G//H^c$ denotes the graph obtained by collapsing all edges of H^c to points.

The jewel $J(G)$ is a convex polytope, obtained by shaving off some faces of $\sigma(G)$ that are at infinity. Specifically, label each vertex of $\sigma(G)$ by the corresponding edge of G. If a set of edges forms a core subgraph of G, we shave off the opposite face, that is, the face spanned by the remaining edges. We shave deeper for larger core graphs.

More precisely: if G has edges $e_0, \ldots, e_m$, realize $\sigma(G)$ as the set of points in the positive orthant of $\mathbb{R}^{m+1}$ whose coordinate sum is equal to 1:

$$\sigma(G) = \{(x_0, \ldots, x_m) \in \mathbb{R}^{m+1} \mid x_i \geqslant 0,\, x_0 + \cdots + x_m = 1\}.$$

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

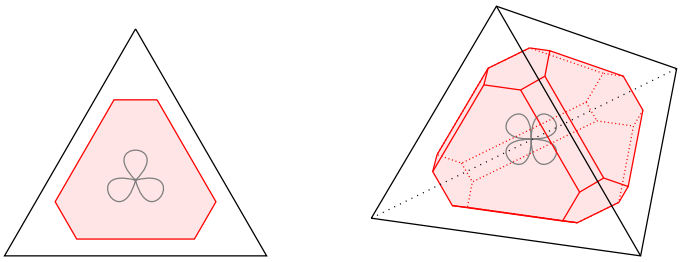

FIGURE 1 colour online). *Permutohedra in rose faces for $n = 3$ and $n = 4$.*

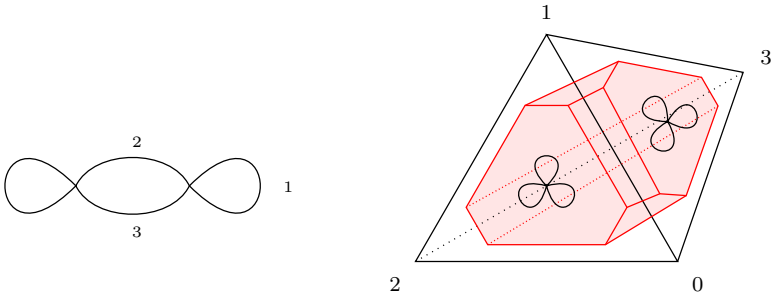

FIGURE 2 colour online). *Example of a graph which is not a rose and its associated jewel.*

If A is a subgraph of G, then there is a natural inclusion $\sigma(A) \subset \sigma(G)$. For each core subgraph A we shave the opposite face $\sigma(A^c)$ by a factor c_A; this is accomplished by taking the intersection of $\sigma(G)$ with the half-space $\sum_{e_i \in A} x_i \geqslant c_A$. Here the constants c_A should be chosen to be small positive numbers which increase quickly with the size of A. Specifically, we require all $c_A \ll 1$ and $c_A > 2c_B$ if A properly contains B. The cell $J(G)$ is the result of shaving all faces opposite core faces of $\sigma(G)$.

REMARK 2.1. The jewel $J(G)$ has also found applications in the context of Feynman integrals, where core subgraphs are called 1-*particle irreducible subgraphs* and $J(G)$ is called the *graph polytope* (see, for example, [4]). In this context the constants c_A depend only on the number of edges in A. In our context it will be more convenient to use constants c_A that depend on the rank of $H_1(A)$.

If G is a rose, then every proper subset of edges is a core subgraph, so every proper face of $\sigma(G)$ is shaved. The resulting convex polytope $J(G)$ is called a *permutohedron* of rank n. The cases $n = 3$ and $n = 4$ are illustrated in Figure 1.

An example of $J(G)$ when G is not a rose is shown in Figure 2. Here a face $\sigma(H)$ of $\sigma(G)$ is identified with the subset of $\{0, 1, 2, 3\}$ indexing the edges of H. The core subgraphs properly contained in G are spanned by the sets $\{e_0\}$, $\{e_1\}$, $\{e_2, e_3\}$, $\{e_0, e_1\}$, $\{e_1, e_2, e_3\}$ and $\{e_0, e_2, e_3\}$, so the faces that are shaved are $\{1, 2, 3\}$, $\{0, 2, 3\}$, $\{0, 1\}$, $\{2, 3\}$, $\{0\}$ and $\{1\}$, respectively. Faces of $\sigma(G)$ obtained by collapsing a maximal tree are called *rose faces*; note that these are not shaved. In our example $\sigma(G)$ has two rose faces, $e_2 = 0$ and $e_3 = 0$. Since the rank of G is 3, each rose face contains a permutohedron with six vertices. These are the only vertices of $J(G)$, that is, $J(G)$ is the convex hull of the permutohedra contained in the rose faces of $\sigma(G)$. This description of $J(G)$ is general:

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

PROPOSITION 2.2. *For any (G, g) in $\mathcal{O}_n$, the jewel $J(G)$ is the convex hull of the permutohedra contained in the rose faces of $\sigma(G)$.*

Proof. Since $J(G)$ is a convex polytope, it is the convex hull of its vertices. Each vertex lies in the interior of some (not necessarily proper) face τ of $\sigma(G)$, so it suffices to show that τ must be a rose face. The face τ corresponds to a subgraph $H \subseteq G$, that is, $\tau = \sigma(H)$. If τ were at infinity, then the complement H^c would contain a non-trivial core subgraph C. Since the face $\sigma(C^c)$ opposite $\sigma(C)$ is truncated when forming $J(G)$ and $H \subseteq C^c$, the face $\tau = \sigma(H) \subseteq \sigma(C^c)$ would not intersect $J(G)$, so would not contain any vertices of $J(G)$. Therefore τ is not at infinity, that is, H^c is a forest and $\sigma(G//H^c) < \sigma(G)$. If we choose the constants carefully we in fact have $J(G//H^c) = J(G) \cap \tau$. This reduces the problem to showing that for any graph G, if $J(G)$ has a vertex in the interior of $\sigma(G)$ then G is a rose.

In the following we identify faces of $\sigma(G)$ with subsets of $\{0, \dots, m\}$, and call such a subset *core* if it corresponds to a core subgraph. Suppose $y = (y_0, \dots, y_m)$ is a vertex of $J(G)$ with all y_i positive. Then y lies on bounding hyperplanes $\mathcal{H}_A$ for some collection $\mathcal{S}(y)$ of core subsets $A \subset \{0, \dots, m\}$, that is, y satisfies $\sum_{i \in A} y_i = c_A$ for all $A \in \mathcal{S}(y)$.

Note that the union of two core subgraphs is always a core subgraph. Thus if A and B are in $\mathcal{S}(y)$, then $(A \cup B)^c$ is shaved. So for all $x \in J(G)$

$$\sum_{k \in A \cup B} x_k \geqslant c_{A \cup B}.$$

Since $\sum_{i \in A} y_i = c_A$ and $\sum_{j \in B} y_j = c_B$ we have

$$\sum_{k \in A \cup B} y_k \leqslant c_A + c_B.$$

If A and B are both proper subsets of $A \cup B$, then each of c_A and c_B is less than half of $c_{A \cup B}$, so $c_A + c_B < c_{A \cup B}$, giving

$$\sum_{k \in A \cup B} y_k < c_{A \cup B},$$

This contradicts the assumption that y is in $J(G)$.

Thus the subsets of $\mathcal{S}(y)$ are nested, that is, form a flag. Since a vertex is the intersection of at least m hyperplanes, $\mathcal{S}(y)$ contains at least m proper subsets, so up to permutation the flag is $\{0\}, \{0, 1\}, \{0, 1, 2\} \dots, \{0, 1, \dots m-1\}$. Thus $\{e_0\}$ is a core subgraph, so it must be a loop. Since $\{e_0, e_1\}$ also forms a core subgraph, and core subgraphs have no separating edges, e_1 is also a loop. Continuing, we get that $e_0, \dots, e_{m-1}$ are all loops, which implies that e_m too is a loop, since G has no separating edges. Since G is connected and all edges are loops, we conclude that G is a rose (and $n = m + 1$).

2.1. *Fitting jewels together to form $\mathcal{J}_n$*

Suppose $(G\ , g\ )$ can be obtained from (G, g) by collapsing the edges of some subforest Φ of G to points. Then for appropriate truncating constants, $J(G\ , g\ )$ is a face of $J(G, g)$. Specifically, we need $c_A\ = c_A$ where $A = \text{core}(A\ \cup \Phi)$ is the largest core graph in G mapping to $A\ $. To make the constants c_A consistent over all $J(G, g)$ containing $J(G\ , g\ )$, we assume that c_A depends only the rank of $H_1(A)$. This works because $\operatorname{rank} H_1(A\ ) = \operatorname{rank} H_1(A)$ and if $A \quad B$ are core graphs one needs to remove at least one edge of B to get A; this decreases the rank of $H_1(B)$ since B has no separating edges.

Let $\mathcal{J}_n$ denote the union of the cells $J(G, g)$ for all marked graphs (G, g) in $\mathcal{O}_n$. We claim that this a closed subspace of $\mathcal{O}_n$ which is a deformation retract. To see this, recall that $\mathcal{O}_n$ is the union of the simplices $\mathring{\sigma}(G, g)$ glued together using the same face relations, each cell $J(G, g)$

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

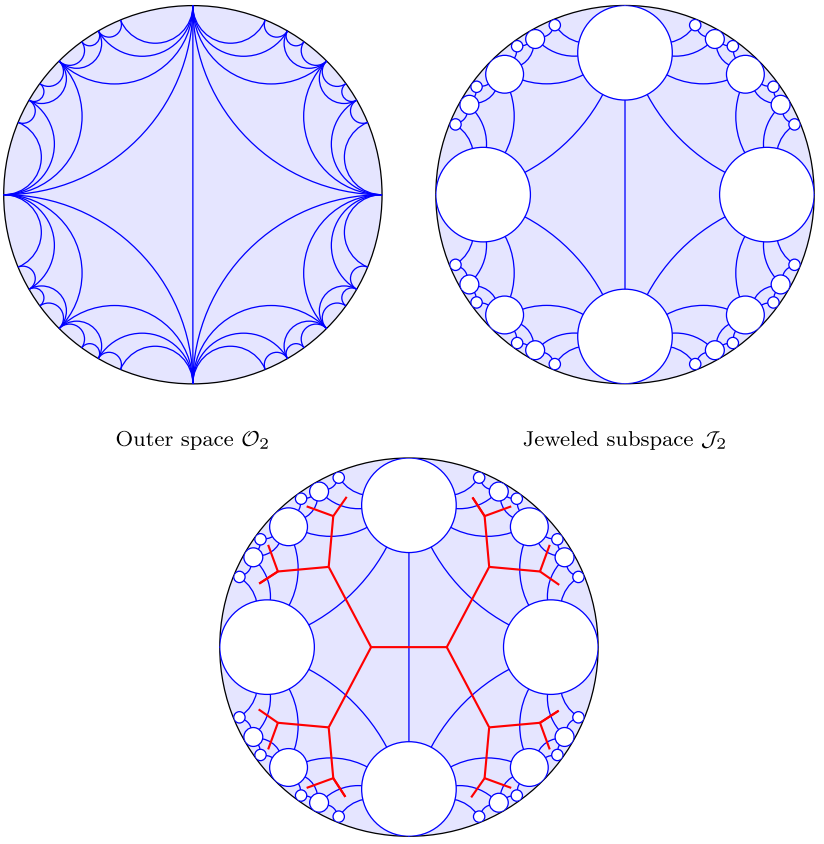

FIGURE 3 colour online). *Outer space the subspace of jewels and the spine for* $n = 2$.

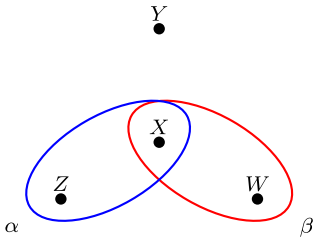

FIGURE 4 colour online). *Key Lemma.*

is evidently a deformation retract of the closure $\bar{\sigma}(G, g)$ of $\mathring{\sigma}(G, g)$ in $\mathcal{O}_n$, and the deformation retraction can be taken to restrict to deformation retractions of all faces $\bar{\sigma}(G\ , g\ )$. Figure 3 shows the relation between the spaces $\mathcal{O}_n, \mathcal{J}_n$ and the spine K_n of $\mathcal{O}_n$ for the case $n = 2$. Here the Euclidean simplices have been deformed for artistic convenience so that they fit into the Poincaré disk as hyperbolic triangles.

The following statement is an immediate corollary of Proposition 2.2.

COROLLARY 2.3. *The vertices of* $\mathcal{J}_n$ *are the marked ordered roses* $(g, R, e_1, \ldots, e_n)$*, where* (R, g) *is a marked rose and* $e_1, \ldots, e_n$ *is an (ordered) list of the petals of* R*.*

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

3. *Simplicial completion of $\mathcal{J}_n$ and a further retract*

We next want to investigate the connectivity properties of $\mathcal{J}_n$ at infinity, in particular to prove

THEOREM 3.1. *The space $\mathcal{J}_n$ is $(2n-5)$-connected at infinity.*

For this it is convenient to replace $\mathcal{J}_n$ by a simplicial complex so that we may use simplicial Morse theoretic arguments. We do this in two stages. First, replace each cell $J(G,g)$ by a simplex $s(G,g)$ with the same vertices to obtain a simplicial complex $\boldsymbol{S}_n$ (this was also done in [2], at the end of the paper).

LEMMA 3.2. *$\boldsymbol{S}_n$ is contractible, and* $\mathrm{Out}(F_n)$ *acts properly and cocompactly.*

Proof. The spaces $\boldsymbol{S}_n, \mathcal{J}_n$ and $\mathcal{O}_n$ are all complexes of contractible spaces (that is, the spaces $s(G,g), J(G,g)$ and $\bar{\sigma}(G,g)$, respectively) and all have the same nerve, namely the spine K_n. Since K_n is contractible, $\boldsymbol{S}_n, \mathcal{J}_n$ (and $\mathcal{O}_n$) are also contractible (see, for example, [8, section 4.G]). The action permutes the cells of $\boldsymbol{S}_n$ and $\mathcal{J}_n$, and there is a compact fundamental domain since K_n is cocompact and the cells of $\boldsymbol{S}_n$ and $\mathcal{J}_n$ are compact.

Now let $\boldsymbol{R}_n$ be the simplicial complex whose vertices are the marked roses in $\mathcal{O}_n$, and where roses are in the same simplex if and only if they can be obtained from a common (G,g) by collapsing maximal trees.

COROLLARY 3.3. *$\boldsymbol{R}_n$ is contractible, and the action of* $\mathrm{Out}(F_n)$ *on it is proper and cocompact.*

Proof. The map sending vertices of $\boldsymbol{S}_n$ to vertices $\boldsymbol{R}_n$ by forgetting the ordering on the petals of the rose extends to a simplicial map and the inverse image of each simplex is itself a simplex, so the map is a homotopy equivalence (see, for example, [9, Corollary 2.7]).

Since both $\mathcal{J}_n$ and $\boldsymbol{R}_n$ are contractible with proper cocompact $Out(F_n)$-actions, Theorem 3.1 is equivalent to the following.

THEOREM 3.4. *The simplicial complex $\boldsymbol{R}_n$ is $(2n-5)$-connected at infinity.*

4. *The Morse function and ascending links*

We are now ready to attack connectivity at infinity, using Morse theory. We begin by defining a *Morse function* μ on the vertices of $\boldsymbol{R}_n$ with values in a certain ordered abelian group.

4.1. *The Morse function*

Let $\rho = (R,g)$ be a vertex of $\boldsymbol{R}_n$, that is, a marked rose, and suppose F_n is generated by $x_1,\dots,x_n$. Let $\mathcal{W}_0$ be the set of conjugacy classes of elements of the form x_i or x_ix_j or $x_ix_j^{-1}$ for $i \neq j$, and let $\mathcal{W} = \{w_1, w_2, w_3, \dots\}$ be a list of *all* the conjugacy classes in F_n.

Given a conjugacy class w and an edge $e \in R$, define $|e|_w$ to be the minimum, over all loops γ homotopic to $g(w)$, of the number of times γ crosses e in either direction. Then define

$$|\rho|_w = \sum_{e \in R} |e|_w$$

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

and

$$|\rho|_0 = \sum_{w \in \mathcal{W}} |\rho|_w.$$

Finally, define

$$\mu(\rho) = (|\rho|_0, |\rho|_{w_1}, |\rho|_{w_2}, \ldots) \in \mathbb{R} \times \mathbb{R}^{\mathcal{W}}$$

Here $\mathbb{R} \times \mathbb{R}^{\mathcal{W}}$ is an ordered abelian group with the lexicographical order.

LEMMA 4.1. (i) $|\rho|_0 > 0$.
(ii) $|\rho|_{w_i} > 0$ *for all* i.
(iii) *If* $\mu(\rho) = \mu(\rho')$ *then* $\rho = \rho'$.
(iv) *At most finitely many* ρ *have* $\mu(\rho) \leqslant N$ *for a given* N.
(v) μ *well-orders the vertices of* $\boldsymbol{R}_n$.

Proof. A proof may be found in [15]. It relies on the fact that a free action on a simplicial tree is uniquely determined by its translation length function, which was proved by Culler and Morgan [5, Theorem 3.7] and independently by Alperin and Bass [1, Theorem 7.13]. The proof from [15] applies here because $|\rho|_0$ is the norm used in the original proof in [6] of the contractibility of Outer space.

4.2. *Ascending links*

In this section we reduce connectivity of $\boldsymbol{R}_n$ at infinity to a local problem. To do this we use the Morse function to arrange all vertices of $\boldsymbol{R}_n$ into an ordered list. The link of each vertex v then has a descending part (spanned by the vertices listed before v) and an ascending part (spanned by the vertices listed after v). A standard argument shows that $\boldsymbol{R}_n$ is $(2n-5)$-connected at infinity provided that the ascending subcomplex of each vertex link is $(2n-5)$-connected. Here are the details.

By Lemma 4.1, the map

$$\begin{aligned} h\colon \boldsymbol{R}_n^{0)} &\to \mathbb{N} \\ v &\mapsto \{u \in \boldsymbol{R}_n^{0)} \mid \mu(u) < \mu(v)\} \end{aligned}$$

is a well-defined bijection between the 0-skeleton $\boldsymbol{R}_n^{0)} \subseteq \boldsymbol{R}_n$ and the natural numbers $\mathbb{N}$; it just counts the number of vertices that come before v in the well-ordering of roses given by μ. Thus we have a list $\rho_0, \rho_1, \rho_2, \ldots$ of all the roses in ascending order of μ-values.

Let $\boldsymbol{R}_{\ i}$ be the subcomplex of $\boldsymbol{R}_n$ spanned by the vertices $\rho_i, \rho_{i+1}, \rho_{i+2}, \ldots$ The *ascending link* of the rose ρ_i is defined to be

$$\mathrm{lk}^+(\rho_i) := \mathrm{lk}(\rho_i) \cap \boldsymbol{R}_{\ i+1}.$$

Recall that a simplicial complex is k-*spherical* if it is k-dimensional and $(k-1)$-connected. In Section 6 we will prove

THEOREM 4.2. *For every rose* ρ, *the ascending link* $\mathrm{lk}^+(\rho)$ *is* $(2n-4)$-*spherical.*

From this we can deduce Theorem 3.4 using the following argument from [2, Theorem 5.3]. Since every compact subset of $\boldsymbol{R}_n$ is disjoint from $\boldsymbol{R}_{\ i}$ for sufficiently large i, it suffices to show that each $\boldsymbol{R}_{\ i}$ is $(2n-5)$-connected. This is done by induction.

The base case $i = 0$ is the statement that $\boldsymbol{R}_n$ is contractible. Assuming that $\boldsymbol{R}_{\ i}$ is $(2n-5)$-connected, observe that $\boldsymbol{R}_{\ i}$ is obtained from $\boldsymbol{R}_{\ i+1}$ by adding the vertex ρ_i and coning

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

off its $(2n-5)$-connected ascending link $\mathrm{lk}(\rho_i) \cap \boldsymbol{R}_{i+1}$. By the theorems of Hurewicz and van-Kampen, it follows that $\boldsymbol{R}_{i+1}$ is again $(2n-5)$-connected.

5. *Blowups and ideal edges*

It remains to prove Theorem 4.2. For the proof we use the technology introduced in [6] relating graphs, maximal trees, roses and the norm μ. In this section we give a brief review of this technology. A more careful discussion can be found in [15].

Let G be a graph of rank n, let $t_1, \ldots, t_k$ be the edges of a maximal tree T in G and let $e_1, \ldots, e_n$ be the remaining edges of G. Choose an orientation for each e_i and let $\overline{e}_i$ denote the same edge with the opposite orientation. Removing any t_i cuts T into two subtrees (either of which may be a point), and determines a partition α_i of the set $E = \{e_1, \overline{e}_1, \ldots, e_n, \overline{e}_n\}$ into two pieces according to which subtree contains the terminus of the oriented edge. Partitions α_i and α_j determined by t_i and t_j are *compatible* in the sense that one side of α_i is disjoint from one side of α_j.

Collapsing T produces a rose R_n with oriented petals $E = \{e_1, \overline{e}_1, \ldots, e_n, \overline{e}_n\}$. The graph G can be uniquely reconstructed from the set of partitions $\{\alpha_j\}$ of E.

Now let $\rho = (R, g)$ be a marked rose, and let $E = \{e_1, \overline{e}_1, e_2, \overline{e}_2, \ldots, e_n, \overline{e}_n\}$ be the oriented petals of R. A partition of E into two parts A and $A^c = E - A$ *splits* e_i if e_i and $\overline{e}_i$ are on different sides of the partition. An *ideal edge* of ρ is a partition of E into two sets, each with at least two elements, which splits some e_i. A set of pairwise-compatible ideal edges is called an *ideal tree*; it corresponds to a maximal tree in a graph that has no leaves or bivalent vertices and no separating edges.

Note that an ideal edge $\alpha = (A, A^c)$ is determined by either of its sides, which we call *representatives* for α.

5.1. *Ideal edges, star graphs and μ*

Let R be a rose and E its set of oriented edges, equipped with the involution $e \mapsto \overline{e}$, sending e to the same edge with opposite orientation. If $\gamma = a_1 \ldots a_k$ is a cyclically reduced edge path in R the *star graph* of γ is the graph with vertices E and an edge from a_i to $\overline{a}_{i+1}$ for each $i = 1, \ldots, k$ (setting $k+1 = 1$).

Now fix a marking $g\colon R_n \to R$. Each conjugacy class w of F_n can be represented by an edge path in R_n, and its image $g(w)$ is homotopic to a unique cyclically reduced edge path $\gamma(w)$. We define $st(w)$ to be the star graph of $\gamma(w)$, and $st(\mathcal{W}_0)$ to be the superposition of $st(w)$ for all $w \in \mathcal{W}_0$. Thus the sequence $(\mathcal{W}_0, w_1, w_2, w_3, \ldots)$ gives an infinite sequence of star graphs associated to $\rho = (R, g)$.

For disjoint subsets X and Y of E, define $(X \cdot Y)_w$ to be the number of edges of $st(w)$ with one vertex in X and one vertex in Y, and

$$X \cdot Y = ((X \cdot Y)_{\mathcal{W}}, (X \cdot Y)_{w_1}, (X \cdot Y)_{w_2}, \ldots) \in \mathbb{R} \times \mathbb{R}^{\mathcal{W}}.$$

If $X \subset E$, define $|X|_w = (X \cdot X^c)_w$ and $|X| = X \cdot X^c$. In particular, if $X = \{e\}$ this agrees with our previous definition of $|e|_w$.

Note that the dot product is commutative:

$$X \cdot Y = Y \cdot X$$

and the following 'distributive law' holds for pairwise disjoint subsets $X, Y, Z \subseteq E$:

$$(X \sqcup Y) \cdot Z = X \cdot Z + Y \cdot Z$$

If α is an ideal edge with sides A and A^c, define $|\alpha| = |A| = |A^c|$.

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

LEMMA 5.1 (Positivity of the dot product and absolute value). *For any non-empty disjoint subsets $X, Y \subseteq E$, we have $X \cdot Y > 0$. In particular:*

(i) $|e| > 0$ *for all* $e \in E$.
(ii) $e \cdot f > 0$ *for all* $e \neq f$ *in* E.
(iii) $|\alpha| > 0$ *for all ideal edges* α.

Proof. By distributivity, we can reduce to singletons $X = \{e\}$ and $Y = \{f\}$. Let w_e be the conjugacy class represented by the loop e in ρ. Define w_f analogously.

If $f = \overline{e}$ then $st(w_e)$ is a single edge from e to f, so the w_e-coordinate of $e \cdot f$ is equal to $1 \neq 0$.

Otherwise $st(w_e w_f^{-1})$ consists of two edges, one joining e and f and one joining $\overline{e}$ and $\overline{f}$. Therefore the $w_e w_f^{-1}$ coordinate of $e \cdot f = 1 \neq 0$.

LEMMA 5.2 (Consequences of equality).

(i) *If* $|e| = |f|$ *then* $f = \overline{e}$ *or* $f = e$.
(ii) *If* α, β *are two distinct ideal edges then* $|\alpha| \neq |\beta|$.
(iii) *If* α *is an ideal edge and* $e \in E$ *then* $|\alpha| \neq |e|$.

Proof. Let w_e denote the conjugacy class corresponding to the loop e.

(i) If $f \notin \{e, \overline{e}\}$, then the w_e coordinates of $|e|$ and $|f|$ are different.
(ii) If there is an edge e split by only one of α or β then the w_e-coordinates of α and β are different. If not, choose sides A and B so that $A \cap B, A - (A \cap B)$ and $B - (A \cap B)$ are all non-empty.

If there are $e \in A \cap B$ and $f \in A - (A \cap B)$ which are both split, then $st(ef)$ crosses α but not β. If there are $e \in A \cap B$ and $f \in A - (A \cap B)$ neither of which is split, then $st(ef)$ crosses β but not α. In either case the w_{ef} coordinates of $|\alpha|$ and $|\beta|$ are different.

If $e \in A \cap B$ is split but $f \in A - (A \cap B)$ is not, then choose $z \in B - (B \cap A)$. If z is split, we are in a previous case by symmetry, so we may assume $\overline{z} \in B - (A \cap B)$. Then $st(fez\overline{e})$ crosses α but not β.

It remains to consider the case that $e \in A \cap B$ is not split but $f \in A - (A \cap B)$ is split. Since B^c cannot be a singleton, we can find $z \neq f$ in $A - (A \cap B)$ or in $(A \cap B)^c$. In the first case we may assume $\overline{z} \in B - (A \cap B)$; otherwise we can reduce to a previous case by exchanging the roles of z and f. Then $st(ef\overline{z})$ crosses α but not β. In the second case we may assume $\overline{z} \in (A \cap B)^c$; otherwise we could again reduce to a previous case by exchanging f and z. Then $st(efz\overline{f})$ crosses α but not β.

(iii) If α does not split e then the w_e-coordinates of $|e|$ and $|\alpha|$ are different. If α splits both e and f, then the w_{ef}-coordinate is different. If e is the only edge split by α, then choose $f, \overline{f}$ on one side of α and $h, \overline{h}$ on the other side. Then the coordinate of w_{fh} is different.

6. *Proof of Theorem 4.2*

In this section we first show that the ascending link $\mathrm{lk}^+(\rho)$ is homotopy equivalent to a simplicial complex $\boldsymbol{Z}(\rho)$ whose vertices are certain ideal edges of ρ, then we prove that $\boldsymbol{Z}(\rho)$ is $(2n-4)$-spherical in Theorem 6.7.

6.1. *Tools of the trade*

We recall here a few standard tools which will be useful in our proof. We start with some elementary observations about k-spherical complexes. For details see, for example, [**13**].

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

LEMMA 6.1. *A k-spherical complex is either contractible or homotopy equivalent to a non-trivial wedge of k-spheres.*

LEMMA 6.2. *Let K be a k-spherical subcomplex of a simplicial complex L and v a vertex of L. If $\mathrm{lk}(v) \cap K$ is $(k-1)$-spherical, then the subcomplex spanned by K and v is k-spherical.*

LEMMA 6.3. *If K is k-spherical and L is ℓ-spherical, then the simplicial join $K * L$ is $(k+\ell+1)$-spherical.*

We will use a simple version of Quillen's Poset Lemma. A map $f\colon X \to Y$ between posets is a *poset map* if $x \leqslant x'$ implies $f(x) \leqslant f(x')$ for all $x, x' \in X$. The *geometric realization* of a poset is the simplicial complex with one vertex for each element of the poset and a k-simplex for each totally ordered chain of $k+1$ elements. A poset map induces a simplicial map of geometric realizations, and when we use topological terms to describe posets and poset maps we are referring to the geometric realizations.

LEMMA 6.4 [Poset Lemma **13**]. *Let $f\colon X \to X$ be a poset map, with $f(x) \leqslant x$ for each $x \in X$ or $f(x) \geqslant x$ for each $x \in X$). Then the image of X is a deformation retract of X.*

6.2. *Complexes of ideal edges, ideal trees and the ascending link*

Let $\mathcal{I} = \mathcal{I}(\rho)$ be the set of ideal edges of ρ. Recall that a *flag complex* is a simplicial complex that is determined by its 1-skeleton: $(k+1)$ vertices span a k-simplex if and only if every pair spans an edge. Let $\boldsymbol{I} = \boldsymbol{I}(\rho)$ be the flag complex whose vertices are the elements of $\mathcal{I}$, and whose 1-simplices are pairs $\{\alpha, \beta\}$ of compatible ideal edges.

Recall that an *ideal tree* is a set of pairwise-compatible ideal edges. Let $\mathcal{F} = \mathcal{F}(\rho)$ be the collection of all ideal trees in ρ and $\boldsymbol{F} = \boldsymbol{F}(\rho)$ the flag complex whose vertices are the elements of $\mathcal{F}$, with an edge from $\mathcal{A}$ to $\mathcal{B}$ if $\mathcal{A} \cup \mathcal{B} \in \mathcal{F}$. Then $\boldsymbol{F}$ contains $\boldsymbol{I}$ as a subcomplex.

An ideal edge α is *ascending for* e_i if α splits e_i and $|\alpha| > |e_i|$. It is *ascending* if it is ascending for some e_i. A rose ρ' in the link of ρ in $\boldsymbol{R}_n$ is *ascending* if $\mu(\rho') > \mu(\rho)$, where μ is the Morse function μ defined in section 4. In this section we describe the homotopy type of the subcomplex $\mathrm{lk}^+(\rho)$ spanned by these ascending roses in terms of ideal edges.

For a fixed marked rose $\rho = (R, g)$ let $\mathcal{Z}(\rho)$ be the set of ascending ideal edges of ρ and $\boldsymbol{Z}(\rho)$ the flag complex with vertices $\mathcal{Z}(\rho)$ and edges compatible pairs $\{\alpha, \beta\}$. In other words, $\boldsymbol{Z}(\rho)$ is the full subcomplex of $\boldsymbol{I}(\rho)$ spanned by ascending edges.

PROPOSITION 6.5. *The ascending link $\mathrm{lk}^+(\rho)$ is homotopy equivalent to $\boldsymbol{Z}(\rho)$.*

Proof. Every rose in $\mathrm{lk}^+(\rho)$ is obtained from ρ by blowing up some ideal tree $\mathcal{A}$ and then collapsing a maximal tree T to obtain a new rose, denoted $\rho^{\mathcal{A}}_T$. We may assume that T contains none of the blown-up edges, that is, that T is a subset of the edges E of ρ. Then Lemma 5.2 implies that the sets $\mathcal{A} \in \mathcal{F}$ and $T \subset E$ are uniquely determined by ρ', and in particular the map $f\colon \mathrm{lk}^+(\rho) \to \boldsymbol{F}(\rho)$ sending $\rho' = \rho^{\mathcal{A}}_T$ to $\mathcal{A}$ is well defined. Since $\rho^{\mathcal{A}}_T$ and $\rho^{\mathcal{B}}_S$ have a common blowup if and only if $\mathcal{A} \cup \mathcal{B}$ is an ideal tree, f is a simplicial map. The inverse image of the simplex spanned by $\mathcal{A}_0, \ldots, \mathcal{A}_k$ in $\boldsymbol{F}$ consists of roses with the common blowup $\rho^{\mathcal{A}}$, where $\mathcal{A} = \mathcal{A}_0 \cup \cdots \cup \mathcal{A}_k$. Since such roses span a simplex in $\mathrm{lk}^+(\rho)$, f is a homotopy equivalence onto its image $\mathrm{im}(f) \subseteq \boldsymbol{F}$ (see, for example, [**9**, Corollary 2.7]).

Let $\boldsymbol{F}^+(\rho)$ denote the subcomplex of $\boldsymbol{F}(\rho)$ spanned by forests consisting entirely of ascending edges. If $\mathcal{A} = \{\alpha_0, \ldots, \alpha_k\}$ is any ideal tree in $\mathrm{im}(f)$, then by the Factorization Lemma (see, for example, [**15**, Proposition 7.1]), at least one of the α_i must be ascending. (The Factorization Lemma says that one can match the α_i to edges $e_i \in T$ in such a way that α_i splits e_i for

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

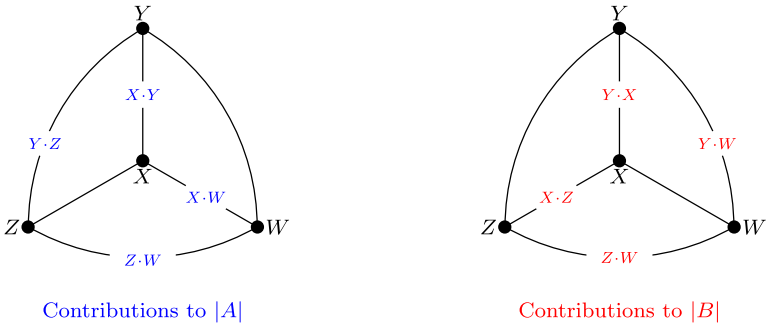

FIGURE 5 colour online). *Proof of Key Lemma.*

all i.) Inclusion makes $\mathcal{F}(\rho)$ into a poset, and the map sending $\mathcal{A} \in \mathrm{im}(f)$ to the subforest $\mathcal{A}^+$ of ascending edges in $\mathcal{A}$ is a poset map satisfying the hypotheses of the Poset Lemma. So it is a homotopy equivalence onto its image $\mathrm{im}(f)^+ \subset \boldsymbol{F}^+(\rho)$.

Every ideal tree consisting of a single ascending edge is in $\mathrm{im}(f)^+$, so $\mathrm{im}(f)^+$ contains $\boldsymbol{Z}(\rho)$. In fact $\boldsymbol{Z}(\rho)$ is a deformation retract of $\mathrm{im}(f)^+$. To see this, define a map $\phi\colon \mathrm{im}(f)^+ \quad \boldsymbol{Z}(\rho)$ by sending a vertex $\mathcal{A}$ to the barycenter of the simplex of $\boldsymbol{Z}(\rho)$ spanned by the elements of $\mathcal{A}$, and extending linearly over simplices of $\mathrm{im}(f)^+$. The deformation retraction is given by sending x to $(1-t)x + t\phi(x)$.

To understand $\boldsymbol{Z}(\rho)$, we need to know when ideal edges are ascending. We will do this using the dot product and norm. The following lemma is a slight variation of [**15**, Lemma 11.2].

LEMMA 6.6 (Key Lemma). *Suppose E is partitioned into four sets X, Y, Z and W. Let $A = X \cup Z$ and $B = Y \cup Z$. Then*

$$|A| + |B| = |X| + |Y| + 2(Z \cdot W).$$

In particular, if $Z \cdot W > 0$ and $|X|, |Y| \geqslant \lambda \in \mathbb{R} \times \mathbb{R}^{\mathcal{W}}$, then $|A| > \lambda$ or $|B| > \lambda$.

Proof. This is a straightforward computation:

$$\begin{aligned} |A| + |B| &= (X \cdot Y + Y \cdot Z + X \cdot W + Z \cdot W) + (Y \cdot X + X \cdot Z + Y \cdot W + Z \cdot W) \\ &= X \cdot (Y \cup Z \cup W) + Y \cdot (X \cup Z \cup W) + 2(Z \cdot W) \\ &= |X| + |Y| + 2(Z \cdot W). \end{aligned}$$

Figure 5 may be helpful for following this computation.

THEOREM 6.7. *For $n \geqslant 2$ and any rose $\rho = (R, g)$ the complex $\boldsymbol{Z}(\rho)$ is $(2n-4)$-spherical.*

We will prove the theorem by induction on n, but the induction requires considering more general complexes. Order the edges of R so that $|e_1| > |e_i|$ for $i > 1$. Suppose E decomposes into the disjoint union of subsets $X_1, \overline{X}_1, \ldots, X_m, \overline{X}_m, Y_1, \ldots, Y_k$ with

- $m \geqslant 1$;
- $e_i \in X_i$;
- $\overline{e}_i \in \overline{X}_i$;
- $|X_i|, |\overline{X}_i| \geqslant |e_i|$.

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

DEFINITION 6.8. A *V-ideal edge* is a partition of

$$V = \{X_1, \overline{X}_1, \ldots, X_m, \overline{X}_m, Y_1, \ldots, Y_k\}$$

into two pieces, each with at least two elements, which separates some X_i from $\overline{X}_i$ (recall $m \geqslant 1$).

Note that a V-ideal edge is also an ideal edge by our old definition, that is, V-ideal edges still split some e_i and can still be classified as ascending or descending. Let $\boldsymbol{Z}(V)$ be the complex whose simplices are sets of pairwise-compatible ascending V-ideal edges. By the above remark, $\boldsymbol{Z}(V)$ is a subcomplex of $\boldsymbol{Z}(\rho)$.

THEOREM 6.9. *If $2m + k - 4 \geqslant 0$ then the complex $\boldsymbol{Z}(V)$ of ascending V-ideal edges is $(2m + k - 4)$-spherical.*

The case $m = n, k = 0$ with $X_i = \{e_i\}, \overline{X}_i = \{\overline{e}_i\}$ is Theorem 6.7. The proof of the theorem is by induction on the pair (m, k), ordered lexicographically. We will need the following lemma, which is basically a restatement of the Key Lemma, to get started with $m = 1$. For $m = 1$ an ideal edge separates X_1 from $\overline{X}_1$ so one side is of the form $X_1 \cup P$, where P is a proper subset of $\mathbb{Y} = \{Y_1, \ldots, Y_k\}$.

LEMMA 6.10. *Suppose $m = 1$ and $k \geqslant 2$, and set $\mathbb{Y} = \{Y_1, \ldots, Y_k\}$. Then for proper subsets P and Q of $\mathbb{Y}$ we have*

(i) *if $X_1 \cup P$ is descending for e_1 then $X_1 \cup (\mathbb{Y} - P)$ is ascending for e_1;*
(ii) *Suppose $X_1 \cup P$ and $X_1 \cup Q$ are descending for e_1 and $P \cap Q = \emptyset$. If $Q \neq \mathbb{Y} - P$ then $X_1 \cup P \cup Q$ is descending for e_1;*
(iii) *Suppose $X_1 \cup P$ and $X_1 \cup Q$ are descending for e_1 and $P \cup Q = \mathbb{Y}$. If $P \cap Q$ is non-empty, then $X_1 \cup (P \cap Q)$ is descending for e_1.*

Proof. Set $\lambda = |e_1|$. Then (1) is the Key Lemma with $X = X_1$, $Y = \overline{X}_1$, $Z = P$ and $W = \mathbb{Y} - P$. Note that $Z \cdot W > 0$ by Lemma 5.1 as P and $\mathbb{Y} - P$ are both non-empty.

For (2), set $X = X_1$, $Z = P$, $W = Q$, $Y = (X \cup Z \cup W)^c$, and $\lambda = |e_1|$. By assumption $|X_1| \geqslant \lambda$. If $|X_1 \cup P \cup Q| \geqslant \lambda$ as well, then by the Key Lemma at least one of $|X \cup Z| = |X_1 \cup Q| > \lambda$ or $|X \cup W| = |X_1 \cup P| > \lambda$, giving a contradiction.

Statement (3) is proved in the same way with $X = \overline{X_1}$, $Z = \mathbb{Y} - P$, $W = \mathbb{Y} - Q$, $Y = (X \cup Z \cup W)^c$, and $\lambda = |\overline{e}_1| = |e_1|$.

For $m = 1$ Theorem 6.9 takes the following stronger form.

PROPOSITION 6.11. *Let $m = 1$ and $k \geqslant 2$, If there are no descending V-ideal edges then $\boldsymbol{Z}(V)$ is a $(k - 2)$-sphere; otherwise it is a contractible subset of this sphere.*

Proof. We have $V = \{X_1, \overline{X}_1, Y_1 \ldots, Y_k\}$. The proof is by induction on k. If $k = 2$ there are only two ideal edges, represented by $X_1 \cup Y_1$ and $X_1 \cup Y_2$. These are incompatible and at least one of them is ascending, by Lemma 6.10 (1). Thus $\boldsymbol{Z}(V)$ is either a point or a 0-sphere.

For any ideal edge α call the side of α containing X_1 the *inside* of α, and the side containing $\overline{X}_1$ the *outside*. The cardinality of the inside is the *size* of the ideal edge. For any $k \geqslant 2$, intersecting the inside of an ideal edge with $\mathbb{Y} = \{Y_1 \ldots, Y_k\}$ gives a one-to-one correspondence between ideal edges and proper subsets of $\mathbb{Y}$. In the remainder of the proof we use this correspondence, referring to a proper subset of $\mathbb{Y}$ as an ideal edge. With this convention, the complex of ideal edges can be identified with the barycentric subdivision of the boundary

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

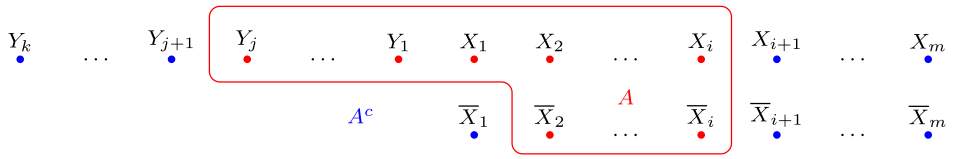

FIGURE 6 colour online). *Ideal edge* $\alpha =$ A A) *in* $\boldsymbol{Z}$ $\boldsymbol{Z}$ *of size* i j).

of a $(k-1)$-simplex Δ. If there are no descending ideal edges, then $\boldsymbol{Z}(V) = \partial\Delta$, which is a $(k-2)$-sphere. Otherwise, let D be a descending ideal edge of minimal size. Then $A = \mathbb{Y} - D$ is ascending by Lemma 6.10 (1).

Let $\boldsymbol{Z}_A$ be the subcomplex of $\boldsymbol{Z}$ spanned by ascending ideal edges B which do not contain D, that is, such that $B \cup A$ is a proper subset of $\mathbb{Y}$. We claim that $B \cup A$ is ascending, so the poset maps B $B \cup A$ A give a contraction of $\boldsymbol{Z}_A$ to the point A. If $B \subseteq A$ then $B \cup A = A$ is certainly ascending. If B intersects D note that $B \cap D = (B \cup A) \cap D$. If $B \cup A$ were descending, then $B \cap D$ would be descending, by Lemma 6.10 (2). But $B \cap D$ is a proper subset of D and D is minimal, so $B \cap D$ is ascending.

We now add the remaining ascending ideal edges to $\boldsymbol{Z}_A$ in order of increasing size, and show that the complex remains contractible after each addition. Ideal edges of the same size are not compatible, so we may add them independently. Let $\boldsymbol{Z}_\ell$ denote the subcomplex of $\boldsymbol{Z}$ spanned by $\boldsymbol{Z}_A$ and all ascending ideal edges of size at most ℓ, and suppose B has size $\ell + 1$, that is, $B \cap \mathbb{Y}$ contains ℓ elements. We need to show that the link of B in $\boldsymbol{Z}$ intersects $\boldsymbol{Z}_\ell$ in a contractible subset. But this intersection consists of ascending subsets of B (ascending sets *containing* B are not in $\boldsymbol{Z}_\ell$ since they contain D and are larger than B). Therefore, replacing $\overline{X}_1$ by $\overline{X}_1 \cup (\mathbb{Y} - B) = E - B$, we can identify $\mathrm{lk}(B) \cap \boldsymbol{Z}_\ell$ with $\boldsymbol{Z}(V\,)$ for the partition $V = \{X_1, E - B, Y_1, \ldots, Y_\ell\}$. By induction, $\boldsymbol{Z}(V\,)$ is either an $(\ell-2)$-sphere or contractible. But it is not the whole sphere because it contains a descending ideal edge, namely D. Therefore the link of B intersects $\boldsymbol{Z}_\ell$ in a contractible set.

We are now in a position to prove Theorem 6.9 for all values of m and k.

Proof. We prove the theorem by induction on pairs (m, k), ordered lexicographically. Lemma 6.11 establishes the theorem for $m = 1$ so we may assume $m \geqslant 2$.

Let $\boldsymbol{Z}_0$ be the subcomplex of $\boldsymbol{Z} = \boldsymbol{Z}(V)$ spanned by ideal edges which are ascending for some e_i with $i > 1$. Since it is irrelevant whether ideal edges in $\boldsymbol{Z}_0$ separate X_1 from $\overline{X}_1$, we can identify $\boldsymbol{Z}_0$ with $\boldsymbol{Z}(V\,)$, with

$$V = \{X_2, \overline{X}_2, \ldots X_m, \overline{X}_m, Y_1, \ldots, Y_k, Y_{k+1} = X_1, Y_{k+2} = \overline{X}_1\},$$

which is $(2m + k - 4)$-spherical by induction.

Ideal edges in $\boldsymbol{Z} - \boldsymbol{Z}_0$ split e_1 but do not split any e_i for $i > 1$, since if they did they would be ascending for that e_i and hence in $\boldsymbol{Z}_0$. Therefore the inside of such an ideal edge is of the form

$$A = X_1 \cup X_2 \cup \overline{X}_2 \cup \cdots \cup X_i \cup \overline{X}_i \cup Y_1 \cup \cdots \cup Y_j$$

(see Figure 6). For the remainder of the proof we specify an ideal edge by giving its inside. We will start adding these to $\boldsymbol{Z}_0$ in order of increasing size, where the *size* of A is the pair (i, j) ordered lexicographically. Ideal edges of the same size are not compatible, so we may add them independently.

Let $\alpha \in \boldsymbol{Z} - \boldsymbol{Z}_0$ with inside A as above, and let $\boldsymbol{Z}_{<\alpha}$ denote the complex spanned by $\boldsymbol{Z}_0$ and all ideal edges of size less than (i, j). We may assume that $\boldsymbol{Z}_{<\alpha}$ is $(2m + k - 4)$-spherical

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

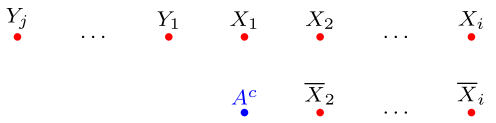

FIGURE 7 colour online). *The decomposition V of E.*

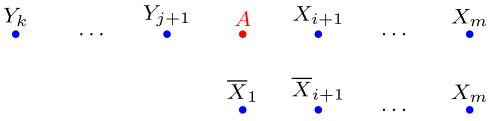

FIGURE 8 colour online). *The decomposition V of E.*

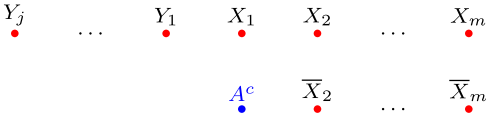

FIGURE 9 colour online). *The decomposition V of E when $i = m$.*

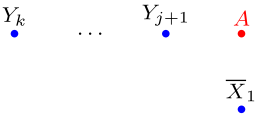

FIGURE 10 colour online). *The decomposition V of E when $i = m$.*

by induction. To prove that it is still $(2m + k - 4)$-spherical after adding α we will show that $\mathrm{lk}(\alpha) \cap \boldsymbol{Z}_{<\alpha}$ is $(2m + k - 5)$-spherical.

The complex $\mathrm{lk}(\alpha)$ is the join of two subcomplexes, namely the subcomplex spanned by ascending sets contained in A and that spanned by ascending sets containing A. Ascending sets contained in A are exactly $\boldsymbol{Z}(V\ )$, where

$$V\ = \{X_1, A^c, X_2, \overline{X}_2 \ldots, X_i, \overline{X}_i, Y_1, \ldots, Y_j\}$$

(see Figure 7). Since all ascending sets contained in A are either in $\boldsymbol{Z}_0$ or are of smaller size they are all in $\boldsymbol{Z}_{<\alpha}$, that is, $\boldsymbol{Z}_{<\alpha} \cap \boldsymbol{Z}(V\ ) = \boldsymbol{Z}(V\ )$, which is $(2i + j - 4)$-spherical by induction.

Ascending sets containing A are exactly $\boldsymbol{Z}(V\ )$, where

$$V\ = \{A, \overline{X}_1, X_{i+1}, \overline{X}_{i+1} \ldots, X_m, \overline{X}_m, Y_{j+1}, \ldots, Y_k\}$$

(see Figure 8). If $i < m$ then $\boldsymbol{Z}_{<\alpha} \cap \boldsymbol{Z}(V\ )$ is isomorphic to $\boldsymbol{Z}_0(V\ )$: an ascending V -ideal edge separating only A and $\overline{X}_1$ has size larger than α and does not belong to $\boldsymbol{Z}_{<\alpha}$; the other ascending V -ideal edges already belong to $\boldsymbol{Z}_0(V)$. Since $\boldsymbol{Z}_{<\alpha} \cap \boldsymbol{Z}(V\ ) \cong \boldsymbol{Z}_0(V\ )$ is $(2(m - i + 1) + (k - j) - 4)$-spherical by induction, the entire complex $\mathrm{lk}(\alpha) \cap \boldsymbol{Z}_{<\alpha}$ is $((2i + j - 4)) + (2(m - i + 1) + (k - j) - 4) + 1 = (2m + k - 5)$-spherical, as required.

This does not work if $i = m$ since in that case $V\ = \{A, \overline{X}_1, Y_{j+1}, \ldots, Y_k\}$, so $\boldsymbol{Z}_0(V\ )$ is empty. The trick is to now add the remaining edges in order of *decreasing* size, as follows:

Let $\boldsymbol{Z}_1(V)$ be the subcomplex of $\boldsymbol{Z}(V)$ spanned by $\boldsymbol{Z}_0$ and all ideal edges in $\boldsymbol{Z} - \boldsymbol{Z}_0$ added so far, that is, of size less than $(m, 0)$. Suppose $\alpha \in \boldsymbol{Z}(V) - \boldsymbol{Z}_1(V)$. Then the *outside* A^c of α is of the form $\overline{X}_1 \cup Y_{j+1} \cup \cdots \cup Y_k$. We will add these to $\boldsymbol{Z}_1(V)$ in order of increasing number of Y s, checking at each stage that $\mathrm{lk}(\alpha) \cap \boldsymbol{Z}_{<\alpha}$ is $(2m + k - 5)$-spherical. As before, edges of the same size are not compatible so we may consider them separately. As before, the link of α is the join of $\mathrm{lk}(\alpha) \cap \boldsymbol{Z}(V\ )$, corresponding to subsets contained in A (whose complement

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

contains A^c), and $\mathrm{lk}(\alpha) \cap \boldsymbol{Z}(V\ )$, corresponding to subsets containing A (whose complement is contained in A^c).

Note that $\boldsymbol{Z}_{<\alpha}$ contains everything in $\boldsymbol{Z}(V\ )$ except the subsets $A^c \cup P$ for $P \subset \{Y_1, \dots, Y_j\}$ as we are adding ideal edges in order of increasing Y s-count. Thus, $\boldsymbol{Z}_{<\alpha} \cap \boldsymbol{Z}(V\ )$ is equal to $\boldsymbol{Z}_1(V\ )$, which is $(2m + j - 4)$-spherical by induction.

Similarly, we have already added all subsets $\overline{X}_1 \cup P$ with P a proper subset of $\{Y_{j+1}, \dots, Y_k\}$, so $\boldsymbol{Z}_{<\alpha} \cap \boldsymbol{Z}(V\ ) = \boldsymbol{Z}(V\ )$, and so is $(k - j - 2)$-spherical. Thus $\mathrm{lk}(\alpha) \cap \boldsymbol{Z}_{<\alpha}$ is $(2m + j - 4) + (k - j - 2) + 1 = (2m + k - 5)$-spherical as required.

7. *Relation with the Bestvina–Feighn bordification*

In this section we recall Bestvina and Feighn's construction of the bordification $b\mathcal{O}_n$, then prove that it is equivariantly homeomorphic to $\mathcal{J}_n$.

7.1. *Cells of the BF bordification*

Bestvina and Feighn also constructed their bordification one cell at a time, producing an enlargement $\Sigma(G, g)$ of each (open) simplex $\mathring{\sigma}(G, g)$ in $\mathcal{O}_n$ and showing these $\Sigma(G, g)$ are compatible with face relations. Thus to describe their construction we can drop the marking g from the notation, as we did in Section 2.

They first define an embedding

$$i_G \colon \mathring{\sigma}(G) \qquad \prod_A \sigma(A),$$

where A runs over all core subgraphs of G (including G itself), and $\sigma(A)$ is a closed regular Euclidean simplex with vertices the edges of A.

A point in $\mathring{\sigma}(G)$ is a volume 1 metric on G, and the image of this point in the term indexed by A is a volume 1 metric on A, obtained by restricting the metric on G to A and then rescaling. The space $\Sigma(G, g)$ is then defined to be the closure of the image of $\mathring{\sigma}(G)$ in $\prod_A \sigma(A)$. Bestvina and Feighn prove that $\Sigma(G, g)$ is contractible, and even that it is homeomorphic to a cell; however in the course of proving this last fact they must appeal to the Poincaré conjecture.

Bestvina and Feighn then re-introduce the markings and show that the spaces $\Sigma(G, g)$ fit together as expected, so that the nerve of the cover is equal to the spine of Outer space. This shows that the union $b\mathcal{O}_n$ of the cells $\Sigma(G, g)$ is contractible. They then extend the action of $\mathrm{Out}(F_n)$ on $\mathcal{O}_n$ to $b\mathcal{O}_n$ and check that the action is proper and cocompact.

7.2. *The faces of $J(G)$*

In order to give an explicit description of the faces of $J(G)$ we first establish some notation.

Index the edges of G by the set $\Delta = \{0, \dots, m\}$, and let be the collection of core subsets of Δ, corresponding to core subgraphs of G. The *core* of an arbitrary subset $U \subset \Delta$ is the maximal element of contained in U. Note that the core of U is unique (though it may be empty), since the union of two core subgraphs is a core subgraph.

Let $\mathscr{F}$ be the set of singletons $\{i\}$ which are not in (that is, the corresponding edges are not loops), and set $\mathscr{S} = \mathscr{F} \cup - \{\Delta\}$.

Recall that the jewel $J(G)$ is defined to be the intersection of the standard simplex $\sigma(G) \subset \mathbb{R}^{m+1}$ with the half spaces $\sum_{i \in A} x_i \geqslant c_A$ for each core subset A.

Let $e_0, \dots, e_m$ be the standard basis of $\mathbb{R}^{m+1}$. For any subset S of Δ, define the vector $e_S = \sum_{i \in S} e_i$ and let x_S be the function on $\sigma(G)$ defined by $x_S(\mathbf{x}) = \langle \mathbf{x}, e_S \rangle = \sum_{i \in S} x_i$, that

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

is, x_S is the total volume of the subgraph corresponding to S with edge lengths x_i. If we set $c_S = 0$ for $s \in \mathscr{F}$, then

$$J(G) = \{(x_0, \dots, x_m) \in \mathbb{R}^{m+1} \mid x_0 + \cdots + x_m = 1 \text{ and } x_S \geqslant c_S \text{ for all } S \in \mathscr{S}\}.$$

We are now ready to characterize the faces of $J(G)$.

PROPOSITION 7.1. *Let G be a core graph and $P = J(G)$. With , $\mathscr{F}$ and $\mathscr{S}$ as above, each codimension k face of P is given by k equations $x_{S_i} = c_{S_i}$ with $S_i \in \mathscr{S}$. Suppose the S_i are ordered so that the size never decreases and the elements of $\mathscr{F}$ come first. Let $U_j = \cup_{i=1}^{j} S_i$. Then for some t between 0 and k we have the following.*

(i) *If $i \leqslant t$ then $S_i \in \mathscr{F}$ and the core of U_i is empty. In particular, U_t is a forest in G.*
(ii) *If $i > t$ then $S_i \in$ and S_i is the core of U_i. In particular, as the U_i form an increasing sequence, their cores also increase; that is, $S_{t+1} \subset S_{t+2} \subset \cdots \subset S_k$.*
(iii) *U_k is a proper subset of Δ.*

Conversely, any sets $S_1, \dots, S_k$ satisfying these conditions determine a face of P.

Proof. By Proposition 2.2 the vertices of $J(G)$ lie on rose faces, and there is one vertex for each ordering of the petals. A rose is obtained by collapsing all of the edges in a maximal tree T of G; each of these edges gives a singleton in $\mathscr{F}$. The remaining edges are ordered $e_1, \dots, e_r$, where r is the rank of $H_1(G)$. What this means for the functions x_S at a vertex of $J(G)$ is that $x_{S_1} = c_{S_1}, \dots, x_{S_m} = c_{S_m}$, where

(i) $S_1, \dots, S_t$ are singletons and no union of these is in (these singletons are the edges of T; no core graph is contained in T).
(ii) S_{t+i} is the core of the graph spanned by the edges in $T \cup \{e_1, \dots, e_i\}$. These form a chain $S_{t+1} \subset \cdots \subset S_m$ of core graphs.

Any subset of a set of equations $x_{S_i} = c_{S_i}$ that determine a vertex, determines a face containing that vertex. Since every face has vertices such subsets in fact give all of the faces.

It remains to show that any collection of sets $S_1 \dots, S_t, S_{t+1}, \dots, S_k$ satisfying conditions (1)–(3) can be enlarged to a full collection of m sets satisfying the same conditions. Such a full set of hyperplanes will define a vertex of the jewel, which is given by a spanning tree T in G and an ordering of the edges in the complement $G - T$. To ensure compatibility with the given sets $S_1 \dots, S_t, S_{t+1}, \dots, S_k$, we choose the spanning tree T in G so that S_i is the core of $T \cup S_i$ for all $i > t$. One way of finding such a spanning tree runs as follows: first enlarge the forest $U_t \cap S_{t+1}$ to a maximal forest in S_{t+1}; add further edges to obtain a maximal forest in S_{t+2}, and continue until you obtain a maximal tree T in G. This construction ensures that $S_i = \text{core}(T \cup S_i)$ for $i > t$.

Now, choose an ordering of the edges outside T compatible with the chain $S_{t+1} \subset S_{t+2} \subset \cdots \subset S_k$, that is, first come edges in $S_{t+1} - (T \cap S_{t+1})$ in some order, then the edges in $S_{t+2} - ((T \cup S_{t+1}) \cap S_{t+2})$ in some order, then those from $S_{t+3} - ((T \cup S_{t+2}) \cap S_{t+3})$, etc. A refinement of the chain $S_{t+1} \subset S_{t+2} \subset \cdots \subset S_k$ can be defined by adding one edge at a time, that is, we find the steps between U_t and S_{t+1} by adding the edges in $S_{t+1} - (T \cap S_{t+1})$ to U_t one at a time in the given order. The steps from S_{t+1} to S_{t+2} are constructed by adding the edges from $S_{t+2} - ((T \cup S_{t+1}) \cap S_{t+2})$ to S_{t+1} again one at a time in the given order. Continuing this way we obtain a chain of length m that defines a the vertex specified by the rose G/T with an ordering of its petals given by the chosen order of edges outside T.

REMARK 7.2. It follows that any vertex of the jewel $J(G)$ satisfies exactly m equations $x_{S_i} = c_{S_i}$ with $S_i \in \mathscr{S}$. The vertex is defined by a spanning tree and an ordering of its complementary

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

edges; and there are only m sets in $\mathscr{S}$ that can occur in a compatible collection: any such compatible collection can be extended to a maximal one, which is unique.

7.3. *The map to* $\Sigma(G)$

Recall that the truncation constants defining $P = J(G)$ depend only on the rank of the core subgraph, that is, $c_A = c_r$ where $r = \operatorname{rank}(H_1(A))$. Recall also that $c_i \ll c_{i+1} \ll 1$ for all i.

A map from P to $\Sigma(G)$ is the same as a family of maps $p_A : P \quad \sigma(A)$ for each $A \in \quad$, compatible in the sense that the image of P under the product of these maps is contained in $\Sigma(G)$. We identify $\sigma(A)$ with the positive cone of the projective space $\mathbb{P}^A = P\mathbb{R}^A$. We will choose each map p_A to be a perturbation of the canonical projection from P to $\sigma(A)$ in which we have stretched the map near the boundary to surject onto $\sigma(A)$. In order to glue our cell-by-cell maps to a map on the whole of $\mathcal{J}_n$, we need to make sure that the various stretches do not interfere with one another. We do this as follows.

For $r \geqslant 1$ let $g_r \colon [c_r, 1] \quad [0,1]$ be a smooth function which is 0 at c_r, strictly increasing between c_r and c_{r+1} and then constantly equal to 1. For $S \in \quad$, set $g_S = g_r$ where $r = \operatorname{rank}(H_1(S))$. For each $A \in \quad$ define $\pi_A \colon P \quad \mathbb{R}^A$ by

$$\pi_A(\mathbf{x})_i = x_i \prod g_S(x_S),$$

where the product is over all $S \in \quad$ which contain i but not all of A.

LEMMA 7.3. *Let $\mathbf{x} \in P$. For any $A \in \quad$, some coordinate $\pi_A(\mathbf{x})_i$ is non-zero.*

Proof. To simplify notation, write $z_i = \pi_A(\mathbf{x})_i$.

We only need to prove that $\pi_A(\mathbf{x}) \neq 0$ for vertices $\mathbf{x}$ of P. This is because vertices satisfy the largest number of relations $x_S = c_S$, making the largest number of z_i equal to 0 (since $g_S(c_S) = 0$). Any other point in the boundary satisfies a subset of these relations.

Since a vertex $\mathbf{x}$ has codimension m we have $x_{S_i} = c_{S_i}$ for some collection of sets $\{S_1, \ldots, S_m\} \in \mathscr{S}$ satisfying the conditions of Proposition 7.1. We observed in Remark 7.2 that for each set $S \in \mathscr{S} - \{S_1, \ldots, S_m\}$, the inequality $x_S \neq c_S$ holds, whence we have $g_S(x_S) > 0$.

Since the U_i are strictly increasing and U_m is proper, it follows that each U_i must have exactly i elements. Without loss of generality we may assume $U_i = \{1, \ldots, i\}$, so that 0 is used by none of the S_i.

If $A \in \quad$, we now want to claim some z_i is non-zero, for $i \in A$. Let k be the smallest index such that U_k contains A; if there is no such k then $0 \in A$ and $z_0 \neq 0$, since only those factors $g_S(x_S)$ contribute to z_0 where $0 \in S$, but $0 \in S$ implies $S \notin \{S_1, \ldots, S_m\}$ whence $g_S(x_S) \neq 0$.

Otherwise $k \in A$, and we claim that $z_k \neq 0$. The expression for z_k does not use any S_i for $i < k$ since those S_i do not contain k. If $i \geqslant k$ then U_i contains A. Since S_i is the core of U_i and A is a core subset, A is contained in S_i, so $g_{S_i}(x_{S_i})$ does not occur in the expression for z_k. Since none of the $g_{S_i}(x_{S_i})$ occur in the expression for z_k, we must have $z_k \neq 0$.

By Lemma 7.3, the image of π_A misses the origin of $\mathbb{R}^A$, so we can compose it with projection to the projective space $\mathbb{P}^A = P(\mathbb{R}^A)$ to obtain a function p_A. Since all the terms $g_S(x_S)$ are non-negative on P the image $p_A(P)$ is actually in $\mathbb{P}^A{}_0$, which is canonically identified with the face $\sigma(A)$ of σ spanned by A.

REMARK 7.4. The $\pi_A(x)_i$ provide homogeneous coordinates for the point $p_A(x)$, that is, the point does not change if we multiply all the $\pi_A(x)_i$ by the same non-zero number. In the

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

interior of the jewel $J(G)$, if we use the number

$$\prod_{S \supseteq A} g_S(x_S) \neq 0$$

as a multiplier, we find that the point $p_A(x)$ is also given by the homogeneous coordinates $\pi_\Delta(x)_i = x_i \prod_{i \in S} g_S(x_S)$. In other words, p_Δ determines p_A in the interior of $J(G)$.

The next proposition shows that the maps p_A are compatible with the face relations in $J(G)$ and $\sigma(G)$. If ϕ is a forest in G, let $\kappa_\phi \colon G \quad G$ be the associated forest collapse. If $\Delta = \{0, \ldots, m\}$ indexes the edges of G, then ϕ corresponds to a subset $\Phi \subset \Delta$ and $\Delta \ = \Delta - \Phi$ indexes the edges of G . Collapsing a forest sends core graphs to core graphs, so κ_ϕ induces a map $\kappa_\phi \colon$ sending $A \mapsto \quad A \ = A - (A \cap \Phi)$.

LEMMA 7.5. *The map $\kappa_\phi \colon$ has a section sending A to the core of $A \cup \Phi$. This core has the same rank as A , and if B is any other core graph in $\kappa_\Phi^{-1}(A\ )$ then* $\mathrm{rank}(B) < \mathrm{rank}(A\ )$.

Proof. The core subsets of Δ that are sent to A by the collapse are partially ordered by inclusion and contain a unique maximal element, namely $A = \mathrm{core}(A \cup \Phi)$. Since A is a core graph, any proper subgraph of A has strictly smaller rank.

PROPOSITION 7.6. *Let ϕ be a forest in G and $\kappa_\phi \colon G \quad G$ the associated forest collapse, so that $J(G\ )$ is a face of $J(G)$ and $\sigma(G\ )$ is a face of $\sigma(G)$. Then the following diagram commutes:*

$$\begin{array}{ccc} J(G') & \hookrightarrow & J(G) \\ \downarrow{\scriptstyle p_{A'}} & & \downarrow{\scriptstyle p_A} \\ \sigma(G') & \hookrightarrow & \sigma(G) \end{array}$$

where $A \ = \kappa_\phi(A) = A - (A \cap \phi)$.

Proof. Since all maps are continuous it suffices to prove the diagram commutes for $\mathbf{x}$ in the interior of $J(G\ )$. For $\mathbf{x} = (x_0, \ldots, x_m)$ in the interior of $J(G\ )$, viewed as a face of $J(G)$, we have that $x_S = c_S$ if and only if $S = \{i\}$ for $i \in \Phi$. Hence the multiplicative factor of Remark 7.4 is non-zero for any core subgraph A of G, and so p_Δ determines p_A on $J(G\ )$. Therefore, it suffices to prove the claim for $A = \Delta$.

In particular, for $\mathbf{x} = (x_0, \ldots, x_m)$ in $J(G\ )$, we have $x_i = 0$ for $i \in \Phi$. Therefore on $J(G\ )$, $x_S = x_{S - \ S \cap \Phi)} = x_S$ for all $S \in \quad$.

We want to show that for each $j \in \Delta - \Phi = \Delta$ we have $(p_\Delta\ \mathbf{x})_j = (p_\Delta \mathbf{x})_j$, that is,

$$x_j \prod_{\{S \ \in \quad |j \in S\ \}} g_S\ (x_S\ ) = x_j \prod_{\{S \in \quad |j \in S\}} g_S(x_S).$$

The set breaks into the disjoint union of the sets $\kappa_\phi^{-1}(S\ )$, for $S \ \in \quad$. For the maximal set in preimage $S = \mathrm{core}(S \ \cup \Phi)$, we have $r = \mathrm{rank}(S) = \mathrm{rank}(S\ )$ so $g_S(x_S) = g_r(x_S) = g_r(x_S\ ) = g_S\ (x_S\ )$. If $T \in \kappa_\phi^{-1}(S\ )$ is not equal to S, then $q = \mathrm{rank}(T) < r$, so $g_T(x_T) = g_T(x_T\ ) = g_q(x_S\ ) = 1$ since $x_S \ \geqslant c_S \ = c_r$ and $r \geqslant q + 1$. Thus

$$g_S\ (x_S\ ) = \prod_{T \in \kappa_\phi^{-1}\ S\ )} g_T(x_T) = g_S(x_S)$$

with $S = \mathrm{core}(S \ \cup \Phi)$. Since this is true for all $S \ \in \quad$, the result follows.

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

7.4. *Homeomorphism*

We now define p to be the product of all of the p_A defined in the last subsection, that is,

$$p = \prod_{A\in} p_A \colon P \qquad \prod_{A\in} \mathbb{P}^A{}_0 \cong \prod_{A\in} \sigma(A).$$

In this section we show that p defines a homeomorphism from $P = J(G)$ to the closure of its image in $\prod \mathbb{P}^A{}_0$. Since $J(G)$ is compact and $\prod \mathbb{P}^A{}_0$ is Hausdorff, it suffices to show that p is injective.

We first show that points on different faces have different images.

Let Q be a face of P, determined by subsets $S_1, \dots, S_t, S_{t+1}, \dots, S_k$ of Δ satisfying the conditions of Proposition 7.1. Let $V_0 = \Delta - U_k$ and $V_i = U_{k+1-i} - U_{k-i}$ for $i > 0$. Note that the V_i are pairwise disjoint. For notational convenience in what follows, set $A_0 = \Delta$ and $A_i = S_{k+1-i}$ for $1 \leqslant i \leqslant r = k - t$, so that the A_i are core graphs and $A_0 = \Delta \supset A_1 = S_k \supset \cdots \supset A_r = S_{t+1}$. Since $U_{k+1-i} = U_{k-i} \cup A_i$, we have $V_i \subseteq A_i$.

LEMMA 7.7. *Let x be a point in the interior of Q, let $1 \leqslant \ell \leqslant r = k - t$ and let $z_i = p_{A_\ell}(x)_i$. Then $z_i \neq 0$ if and only if $i \in V_\ell$.*

Proof. If $i \in A_\ell - V_\ell$ then either $i \in U_t$ or $i \in A_{\ell+1} \subset A_\ell$ because $A_\ell - V_\ell \subseteq U_{k-\ell} = U_t \cup A_{\ell+1}$. For $i \in U_t$, we have $x_i = 0$. For $i \in A_{\ell+1}$, the factor $g_S(x_S) = 0$ with $S = A_{\ell+1}$ contributes to z_i. In either case, $z_i = 0$.

Suppose $i \in V_\ell$. Since $\mathbf{x}$ is in the interior of Q all $g_S(x_S)$ are non-zero except when $S = S_j$ for some j. The expression for z_i does not use any singletons in U_t or any A_j for $j > \ell$ since those A_j do not contain i. If $j \leqslant \ell$ then A_j contains A_ℓ so again is not used in the expression for z_i. Therefore $z_i \neq 0$.

PROPOSITION 7.8. *If x and x are in different open faces of P then $p(x) \neq p(x)$.*

Proof. By Lemma 7.7 if x is in the interior of Q then $p_\Delta(x)$ determines $V_0 = \Delta - U_k$ and therefore determines U_k and $\mathrm{core}(U_k) = S_k = A_1$. The map p_{A_1} then determines $V_1 = U_k - U_{k-1}$, hence U_{k-1} and $\mathrm{core}(U_{k-1}) = S_{k-1} = A_2$. Continuing, we see that the maps p_{A_i} determine all A_i and U_t, so determine Q.

We now concentrate on a single face Q.

PROPOSITION 7.9. *Let $\mathcal{Q} \subseteq$ be the set $\{\Delta, A_1, \dots, A_r\}$ of core graphs associated to Q, and $p_{\mathcal{Q}}$ the product map*

$$p_{\mathcal{Q}} = p_\Delta \times p_{A_1} \times \cdots \times p_{A_r} \colon Q \qquad \mathbb{P}^\Delta{}_0 \times \mathbb{P}^{A_1}{}_0 \times \cdots \times \mathbb{P}^{A_r}{}_0.$$

Then the image of Q is contained in $\mathbb{P}^V{}_0 \times \mathbb{P}^{V_1}{}_0 \times \cdots \times \mathbb{P}^{V_r}{}_0$. The boundary of Q maps to the boundary of $\mathbb{P}^V{}_0 \times \mathbb{P}^{V_1}{}_0 \times \cdots \times \mathbb{P}^{V_r}{}_0$.

Proof. It follows immediately from Lemma 7.7 that the interior of Q maps to $\mathbb{P}^V_{>0} \times \mathbb{P}^{V_1}_{>0} \times \cdots \times \mathbb{P}^{V_r}_{>0}$. Since $p_{\mathcal{Q}}$ is continuous, all of Q maps to the closure $\mathbb{P}^V{}_0 \times \mathbb{P}^{V_1}{}_0 \times \cdots \times \mathbb{P}^{V_r}{}_0$. If x is on the boundary of Q then it is in a different open face Q . By Proposition 7.8 the sets V_i on which p_Q is non-zero determine this face, so at least one more coordinate in the sets V_i which determine Q must be zero.

The proof of the following proposition relies on the explicit formula for the function $p_{\mathcal{Q}}$.

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

PROPOSITION 7.10. *The map* $p_{\mathscr{Q}}\colon Q \quad \mathbb{P}^{V}{}_{0} \times \mathbb{P}^{V_1}{}_{0} \times \cdots \times \mathbb{P}^{V_r}{}_{0}$ *is a local diffeomorphism on the interior of* Q,

Proof. Relabeling edges if necessary, we may assume $0 \in V_0$ and $\ell \in V_\ell \subseteq A_\ell = S_{k+1-\ell}$ for $\ell = 1, \ldots, r = k - t$.

If $i \in V_\ell$ set $z_i = p_{A_\ell}(x)_i$, which we may do since $V_\ell \subset A_\ell$. By Lemma 7.7, if x is in the interior of Q, then all such z_i are non-zero; in particular $z_\ell \neq 0$. Thus the image of Q is contained in the affine charts $Z_i = \frac{z_i}{z_\ell}$ on $\mathbb{P}^{V_\ell}$.

We've defined $z_i = x_i \prod g_S(x_S)$ where the product is over all $S \in$ which contain i but not all of A_ℓ. Here, the isolated factor x_i introduces a complication in the formula for the derivative. We can resolve this by defining a new function $\tilde{g}_S$ for any $S \in \mathscr{S}$ (as opposed to $S \in$) by

$$\tilde{g}_S(t) = \begin{cases} g_S(t) & \text{if } S \text{ is not a singleton} \\ t & \text{if } S = \{i\} \in \mathscr{F} \\ t g_{\{i\}}(t) & \text{if } S = \{i\} \text{ is a loop} \end{cases}.$$

We can now rewrite the formula for z_i as

$$z_i = \prod_{\{S \in \mathscr{S} \mid i \in S, A_\ell \quad S\}} \tilde{g}_S(x_S).$$

Let $e_{ij} = e_i - e_j$. For $i \in V_\ell$ we now have

$$Z_i = \frac{z_i}{z_\ell} = \frac{\prod_{i \in S, A_\ell \quad S} \tilde{g}_S(x_S)}{\prod_{\ell \in S, A_\ell \quad S} \tilde{g}_S(x_S)} = \prod_{S \in \mathscr{S}} \tilde{g}_S(x_S)^{\langle e_{i\ell}, e \ \rangle}$$

Note that $\langle e_{i\ell}, e_{S_j} \rangle = 0$ for all $j = 1, \ldots, k$ since each S_j either contains A_ℓ or does not contain any element of V_ℓ, so it suffices to take the product over all $S \neq S_1, \ldots, S_k$. This is to say, Z_i does not depend on x_S for $S \in \{S_1, \ldots, S_k\}$, whence the partial derivative $\partial Z_i / \partial x_S$ vanishes.

The $m - k$ vectors $\{e_{i\ell} \mid i \in V_\ell,\ i \neq \ell\}$ are linearly independent, span an $m - k$-dimensional subspace of $\mathbb{R}^{m+1}$ and are perpendicular to e_{S_j} for all $j = 1, \ldots, k$ and to $(1, \ldots, 1)$, so give a basis for the tangent space TQ to Q. Let m_i be the coordinates on TQ determined by the vectors $e_{i\ell}$, that is, $m_i(e_{j\ell}) = \delta_{ij}$. Note that j determines ℓ via $j \in V_\ell$.

We want to show that the derivative $D = Dp_{\mathscr{Q}}$ is non-singular. Since the domain and range both have dimension $m - k$, it suffices to show the kernel of D is zero. For $v = (v_1, \ldots, v_{m-k}) \in TQ$, set $w = ((v_1/Z_1), \ldots, (v_{m-k}/Z_{m-k}))$; then to show $\ker(D) = 0$ it suffices to show $\langle Dv, w \rangle \neq 0$ for all v. We have

$$\frac{\partial Z_j}{\partial m_i} = \sum_{S \neq S_1, \quad , S_k} \frac{\partial Z_j}{\partial x_S} \frac{\partial x_S}{\partial m_i}.$$

So

$$\langle Dv, w \rangle = \sum_{i,j} w_j \frac{\partial Z_j}{\partial m_i} v_i = \sum_{i,j} \sum_{S \neq S_1, \quad , S_k} w_j \left(\frac{\partial Z_j}{\partial x_S} \right) \left(\frac{\partial x_S}{\partial m_i} \right) v_i.$$

Since we are only summing over S with x_S non-zero,

$$\frac{\partial Z_j}{\partial x_S} = \frac{\tilde{g}_S(x_S)}{\tilde{g}_S(x_S)} \langle e_{j\ell}, e_S \rangle Z_j$$

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

so

$$\sum_{i,j}\sum_{S\neq S_1,\ ,S_k} w_j\left(\frac{\partial Z_j}{\partial x_S}\right)\left(\frac{\partial x_S}{\partial m_i}\right)v_i = \sum_{S\neq S_1,\ ,S_k}\sum_{i,j} w_j\frac{\tilde{g}_S(x_S)}{\tilde{g}_S(x_S)}\langle e_{j\ell},e_S\rangle Z_j\langle e_{i\ell},e_S\rangle v_i$$

$$= \sum_{S\neq S_1,\ ,S_k}\sum_{i,j}\frac{v_j}{Z_j}\frac{\tilde{g}_S(x_S)}{\tilde{g}_S(x_S)}\langle e_{j\ell},e_S\rangle Z_j\langle e_{i\ell},e_S\rangle v_i$$

$$= \sum_{S\neq S_1,\ ,S_k}\frac{\tilde{g}_S(x_S)}{\tilde{g}_S(x_S)}(\sum_i\langle v_i e_{i\ell},e_S\rangle)^2.$$

All summands in the final expression are non-negative. To conclude that $\langle Dv,w\rangle\neq 0$, we have to argue that there is some strictly positive contribution. For singletons, g_S is strictly positive. Moreover, there must be some singleton S for which $\langle v_i e_\ell,e_S\rangle$ is not zero, since such e_S are the standard basis for $\mathbb{R}^{m+1}$. This S cannot be one of the $S_1,\dots,S_k$ since these are all orthogonal to TQ.

PROPOSITION 7.11. *$p_{\mathscr{Q}}$ restricted to the interior of Q is a diffeomorphism onto $\mathbb{P}^V_{>0}\times\mathbb{P}^{V_1}_{>0}\times\cdots\times\mathbb{P}^{V_r}_{>0}$.*

Proof. Suppose $K\subset\mathbb{P}^V_{>0}\times\mathbb{P}^{V_1}_{>0}\times\cdots\times\mathbb{P}^{V_r}_{>0}$ is compact. Then $p^{-1}(K)$ is compact in the compact domain Q. By Proposition 7.9, $p_{\mathscr{Q}}$ sends points in ∂Q to the boundary of $\mathbb{P}^V{}_0\times\mathbb{P}^{V_1}{}_0\times\cdots\times\mathbb{P}^{V_r}{}_0$, so $p^{-1}_{\mathscr{Q}}(K)$ is actually contained in the interior of Q. Thus $p_{\mathscr{Q}}$ restricted to the interior of Q is a proper map which is a local diffeomorphism, so it is a covering map. But $\mathbb{P}^V_{>0}\times\mathbb{P}^{V_1}_{>0}\times\cdots\times\mathbb{P}^{V_r}_{>0}$ is simply connected, so $p_{\mathscr{Q}}$ restricted to the interior of Q is a diffeomorphism.

Both the domain P and the codomain $\prod\mathbb{P}^A{}_0$ of p are naturally stratified by the partially ordered set (poset) of their open faces, and the stratification of the codomain induces a stratification of the image.

THEOREM 7.12. *The map p is a stratum-preserving homeomorphism onto its image which restricts to a diffeomorphism on each stratum.*

Proof. The map p is stratified if each $p_A\colon P\to\mathbb{P}^A{}_0$ is stratified. The strata of the simplex $\mathbb{P}^A{}_0$ are defined by the set of coordinates z_i which vanish on them. An open face Q of P is determined by the set of coordinates x_S which are non-zero on Q. As the z_i are given by monomials in the x_S, the set of z_i which vanish is the same for all $x\in Q$, that is, p_A maps strata to strata.

Propositions 7.8 and 7.11 show that p is injective, and the restriction to each stratum is a diffeomorphism onto its image. Since p is an injective continuous map from a compact space to a Hausdorff space, it is a homeomorphism onto its image.

Proposition 7.8 could be rephrased to say that the map p induces an injection from the poset of strata of Q to the poset of strata of $\prod\mathbb{P}^A{}_0$.

As remarked earlier, $\mathbb{P}^A{}_0$ is canonically identified with $\sigma(A)$. The map p_Δ identifies the interior of P with the interior $\mathring{\sigma}(G)$ of $\sigma(G)=\sigma(\Delta)$. If A is a core graph with edges $1,\dots,m$ and x is in the interior of P, we find

$$[p_A(x)_1:\ldots:p_A(x)_m]=[p_\Delta(x)_1:\ldots:p_\Delta(x)_m]$$

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

because $p_\Delta(x)_i = x_i \prod_{S \ni i} g_S(x_S) = p_A(x)_i \big(\prod_{S \supseteq A} g_S(x_S) \big)$ for any edge i in A. Thus we have a commutative diagram

$$\begin{array}{ccccc} & & \mathring{P} & \hookrightarrow & P \\ & \overset{p_\Delta}{\swarrow} & \downarrow{\scriptstyle p_{\mathscr{C}}} & & \approx\downarrow{\scriptstyle p_{\mathscr{C}}} \\ \mathring{\sigma}(G) & \xrightarrow{BF} & \prod_{A \in \mathscr{C}} \sigma(A) & \hookleftarrow & \Sigma(G) \end{array}$$

where BF is the map defined by Bestvina and Feighn, and the image $p\ \ (P)$ is equal to Bestvina and Feighn's cell $\Sigma(G)$.

THEOREM 7.13. *The jewel space $\mathcal{J}_n$ is equivariantly homeomorphic to the bordification $b\mathcal{O}_n$.*

Proof. Theorem 7.12 yields stratum-preserving homeomorphisms

$$p : J(G,g) \qquad \Sigma(G,g)$$

for each marked graph (g,G).

Now consider a forest collapse $\kappa : G \qquad G$. A core graph A in G maps to a core graph $A = \kappa(A)$ in G . The edges of A can naturally be identified with edges in A, whence $\mathbb{P}^A{}_0$ is naturally a face of $\mathbb{P}^A{}_0$. Although several core graphs in G might give rise to the same core graph A in G , the diagonal inclusion

$$\prod_{A \in \quad G\ )} \mathbb{P}^A{}_0 \qquad \prod_{A \in \quad G)} \mathbb{P}^A{}_0$$

is well defined and restricts to an inclusion $\Sigma(G\ , g\ ) \qquad \Sigma(G,g)$ where g is a marking on G and g is the induced marking on G .

The bordification $b\mathcal{O}_n$ is defined by gluing the cells $\Sigma(G,g)$ together, identifying $\Sigma(G\ , g\ )$ with its image in $\Sigma(G,g)$ whenever $(G\ , g\ )$ can be obtained from (G,g) by collapsing a forest. In categorical language, the bordification $b\mathcal{O}_n$ is the colimit of the functor $(G,g) \mapsto \Sigma(G,g)$, defined on the category of marked graphs with forest collapses as morphisms.

Similarly, the jewel $J(G\ , g\ )$ is a face of the jewel $J(G,g)$ and jewel space $\mathcal{J}_n$ can be obtained by gluing the jewels along the given identifications, that is, $\mathcal{J}_n$ is the colimit of the functor $(G,g) \mapsto J(G,g)$. This follows from the description given in Section 2.1 because Outer space $\mathcal{O}_n$ is the colimit of the cells $\bar{\sigma}(G,g)$ and the deformation retractions $\bar{\sigma}(G,g) \qquad J(G,g)$ are compatible with inclusion of faces.

By Proposition 7.6 the diagrams

$$\begin{array}{ccc} J(G',g') & \hookrightarrow & J(G,g) \\ \downarrow & & \downarrow \\ \Sigma(G',g') & \hookrightarrow & \Sigma(G,g) \end{array}$$

are commutative. Thus we can paste the homeomorphisms $J(G,g) \qquad \Sigma(G,g)$ together to obtain an equivariant homeomorphism of colimits.

8. *The boundary of the bordification*

In a sequel to this paper we will use the identification of $\mathcal{J}_n$ with $b\mathcal{O}_n$ to study the boundary of the bordification, that is, the difference between the bordification and its interior. In particular, we show how to cover this boundary by contractible subcomplexes with contractible intersections. This is analogous to Borel and Serre's covering of the bordification of symmetric

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

space by Euclidean spaces $e(P)$ associated to parabolic subgroups P. In the Borel–Serre case the nerve of the covering is homotopy equivalent to the Tits building of subspaces of a rational vector space, which has the homotopy type of a wedge of spheres. The top-dimensional homology of the Tits building is the dualizing module. In our case the nerve of the covering is homotopy equivalent to a subcomplex of the sphere complex in a doubled handlebody (also called the complex of free splittings), and the relation between its homology and the dualizing module is not so clear.

References

1 R. Alperin and H. Bass, 'Length functions of group actions on Λ-trees , *Combinatorial group theory and topology* (Alta, Utah, 1984) 265–378; *Ann of Math Stud* 111 (Princeton University Press, Princeton, NJ, 1987).

2 M. Bestvina and M. Feighn, 'The topology at infinity of Out($_n$) , *Invent Math* 140 (2000) 651–692.

3 A. Borel and J. P. Serre, 'Corners and arithmetic groups , *Comment Math Helv* 48 (1973) 436–491.

4 F. Brown, 'Feynman amplitudes, coaction principle and cosmic Galois group , *Commun Number Theory Phys* 11 (2017) 453–556.

5 M. Culler and J. W. Morgan, 'Group actions on $\mathbb{R}$-trees , *Proc London Math Soc* (3) 55 (1987) 571–604.

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

7 D. Grayson, 'Reduction theory using semistability , *Comment Math Helv* 59 (1984) 600–634.

8 A. Hatcher, *Algebraic Topology* (Cambridge University Press, Cambridge, 2002) xii+544.

9 A. Hatcher and K. Vogtmann, 'Tethers and homology stability for surfaces , *Algebr Geom Topol* 17 (2017) 1871–1916.

10 L. Ji, 'From symmetric spaces to buildings, curve complexes and Outer spaces , *Innov Incidence Geom* 10 (2009) 33–80.

11 E. Leuzinger, 'An exhaustion of locally symmetric spaces by compact submanifolds with corners , *Invent Math* 121 (1995) 389–410.

12 E. Leuzinger, 'On polyhedral retracts and compactifications of locally symmetric spaces , *Differential Geom Appl* 20 (2004) 293–318.

13 D. Quillen, 'Homotopy properties of the poset of Nontrivial p-subgroups of a group , *Adv Math* 28 (1978) 101–128.

14 C. L. Siegel, *Zur Reduktionstheorie quadratischer Formen*, Publications of the Mathematical Society of Japan 5 (Mathematical Society of Japan, Tokyo, 1959) ix+69.

15 K. Vogtmann, 'Contractibility of Outer space: reprise , *Hyperbolic geometry and geometric group theory*, Advanced Studies in Pure Mathematics 73 (Mathematical Society of Japan, Tokyo, 2017) 265–280.

Kai-Uwe Bux
Fakultät für Mathematik
Universität Bielefeld
Postfach 100131
Universitätsstraße 25 D-33501 Bielefeld
Germany

bux@math.uni-bielefeld.de

Peter Smillie
Department of Mathematics
Harvard University
One Oxford St.
Cambridge, MA 02138
USA

smillie@math.harvard.edu

Karen Vogtmann
Mathematics Institute
Zeeman Building University of Warwick
Coventry CV4 7AL
United Kingdom

k.vogtmann@warwick.ac.uk

14697750, 2018, 1, Downloaded from https://londmathsoc.onlinelibrary.wiley.com/doi/10.1112/jlms.12124 by Cambridge University Library, Wiley Online Library on [18/08/2025]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License