Outer Space for Right-angled Artin Groups

Adrien Abgrall, Ens–Sorbonne Université
advised by
Prof. Vincent Guirardel, Université Rennes 1

July 10, 2022

Contents

1 Introduction

Given any group G, it is of legitimate interest to try to understand the group of automorphisms of G, or rather its quotient $Out(G)$ by the well-behaved subgroup of conjugacy automorphisms. This group arises naturally in the context of geometric group theory : if we realize G as the fundamental group of some path-connected topological space X, based at $x_0 \in X$ and we let $\phi\colon X \to X$ be a homeomorphism, ϕ induces a group isomorphism

$$\phi_*\colon G = \pi_1(X, x_0) \to \pi_1(X, \phi(x_0))$$

The target is isomorphic to G, however the choice of an isomorphism $\psi\colon \pi_1(X, \phi(x_0)) \simeq G$ is non-canonical, *i.e.* depends on the choice of some path

$p\colon x_0 \rightsquigarrow \phi(x_0)$. A change of path from p to q amounts to composing ψ on the left by some conjugacy in G (namely, by the homotopy class of qp^{-1}). All together, the class modulo conjugacy $[\psi \circ \phi_*]$ is a well-defined element of $Out(G)$ induced by ϕ.

However, that process doesn't realize every element of $Out(G)$. Indeed if we represent $F_2 = \langle a, b\rangle$ as the fundamental group of a graph with two edges and one vertex V, any topological automorphism of that graph must send an edge to an edge (possibly flipping it), and the group automorphism :

$$\begin{aligned} a &\mapsto ab \\ b &\mapsto b \end{aligned}$$

cannot be realized up to conjugacy by such a homeomorphism because ab isn't conjugate to any generator or inverse of a generator.

To overcome that obstruction, one needs to consider not only homeomorphisms, but more generally all homotopy equivalences. The right viewpoint for studying them is to introduce several topological objects whose fundamental groups have a fixed isomorphism to G up to conjugacy or equivalently up to change of basepoint (called *marked spaces*), and specific, elementary, homotopy equivalences between them (such as *blow-ups* and *collapses*). In our example with F_2 that would amount to considering all finite graphs of Euler characteristic -1. The goal is then to see those objects as points in some configuration space with a natural action of $Out(G)$, as well as a topological structure (even a cellular one if possible). When the action is proper and the configuration space happens to be contractible, it will be given the name of *Outer Space* for G.

Such a space was built first in the case where $G = F_n$ is a finitely generated free group in [1]. Its cellular structure allows explicit computations of bounds related to the cohomology of $Out(F_n)$. The construction has been recently generalized to the case of right-angled Artin groups (or *RAAGs*) in [2] and [3]. That extension proved to add substantial complexity to the objects considered, hence to the proofs as well.

The scope of this thesis is to give an overview of the contents of [2] and, to some lesser extent, [3] in order to outline the construction of Outer Space and the core arguments of the proof of the main result : its contractibility.

I would like to thank my advisor Pr. Vincent Guirardel, as well as Pr. Nicolas Bergeron for their valuable time and advice.

2 Graphs, partitions and right-angled Artin groups

In this section, we present a number of notations and concept introduced in [2] and [3], somewhat rephrased, which will prove crucial for the construction and understanding of Outer Space for a right-angled Artin group.

2.1 Graph-theoretic preliminaries

Definition 1. In this thesis, a *graph* is the data of a set V of *vertices*, and a family $E = (e_i)_{i\in I}$ of pairs of elements of V or *edges*. An edge of the form $\{v, v\}$ for some $v \in V$ is called a *loop*. The graph is called *finite* if both sets V and I are finite.

A *simple graph* is a graph with no loops and with distinct edges (that is : the map $I \to \mathcal{P}_2(V)$ is injective). The family E will be then identified with a subset of $\mathcal{P}_2(V)$.

Any graph is endowed with an integer-valued *path metric*. For a vertex v of a simple graph, we define its *link*, $lk(v)$ as the set of vertices at distance 1 from v and its *star*, $st(v)$ as the set of vertices at distance at most 1 from v[1].

Notation. For $\Gamma = (V, E)$ a finite simple graph, we denote V^- a disjoint copy of V whose elements will be written formally $\{v^- \mid v \in V\}$. We write $V^\pm := V \sqcup V^-$, for $v \in V$, $v^\pm := \{v, v^-\}$, and

$$E^\pm := \big\{\{a, b\} \mid \exists \{v, w\} \in E,\ a \in v^\pm,\ b \in w^\pm\big\} \subseteq \mathcal{P}_2(V^\pm)$$

The graph $\Gamma^\pm = (V^\pm, E^\pm)$ is still simple and finite, and importantly has no edge linking v and v^- for any $v \in V$.

Definition 2. Let $\Gamma = (V, E)$ be a finite simple graph. Let $v, w \in V$ be vertices.

- We say that w *fold-dominates* v and write $v \lhd_f w$:when $v \neq w$ and $lk(v) \subseteq lk(w)$ holds.
- We say that w *twist-dominates* v and write $v \lhd_t w$:when $v \neq w$ and $st(v) \subseteq st(w)$ holds.
- We say that w *dominates* v and write $v \lhd w$:when one of the two (incompatible) conditions above is satisfied.

Note that $v \lhd w$ if and only if $v \neq w$ and $lk(v) \subseteq st(w)$, and that in that case, $v \lhd_t w$ if and only if v and w are adjacent.

These three relations are strict preorders on V. As such, they define respectively the *fold-equivalence*, *twist-equivalence* and *equivalence* relations on V, whose equivalence classes will be denoted as $[\cdot]_f$, $[\cdot]_t$ and $[\cdot]$, and induce strict partial orders on the quotient sets.

An element $v \in V$ is called *twist-dominant* if $w \lhd v$ for some $w \in V$, and *twist-minimal* otherwise.

Remark 1. If Γ is discrete, respectively complete, every vertex fold-dominates, resp. twist-dominates, every other vertex.

Lemma 1. *Let $\Gamma = (V, E)$ be a finite simple graph, Let $u, v, w \in V$. The following relation is impossible :*

$$u \lhd_t v \lhd_f w$$

Proof. Assume the relation is satisfied. Then u and v are adjacent, hence $u \in lk(v) \subseteq lk(w)$ which yields $w \in lk(u) \subseteq st(v)$. As $v \neq w$, v and w are adjacent which contradicts the fold-domination. □

Corollary 2. *The fold-equivalence class of a twist-dominant element only contains the element itself.*

2.2 Whitehead partitions

Definition 3. Let $\Gamma = (V, E)$ be a finite simple graph. A Γ-*Whitehead partition* $\mathfrak{P} = (P, P^*, L)$ is the data of three subsets P, P^* and L of $V^\pm$ satisfying the following conditions :

- $P \sqcup P^* \sqcup L = V^\pm$
- P and P^* have at least two elements

[1] These notations can get confusing when two graphs share the same vertices. The name of the graph will appear in subscript in case of ambiguïty.

- There exists $v \in V$ (not necessarily unique), called *basepoint*, such that $L = lk_{\pm}(v)$ (in particular, L is inversion-stable) and either $v \in P$, $v^{-1} \in P^*$ or $v^{-1} \in P$, $v \in P^*$.

- For every $a \in P$ such that $a^- \in P^*$, $lk_{\pm}(a) \subseteq L$

- For every $a \in P$ and $b \in P^*$, a and b are not adjacent in $\Gamma^{\pm}$.

P and P^* are called the *sides* and L is called the *link* of the partition.

A *based* Γ*-Whitehead partition* is the data $(\mathfrak{P}, v)$ of a partition and a chosen basepoint. It is determined by the data of P and v.

Notation. Given $\mathfrak{P} = (P, P^*, L)$ a Γ-Whitehead partition, we will denote :

- $cis(\mathfrak{P}) = \{v \in V \mid v, v^- \in P \text{ or } v, v^- \in P^*\}$

- $trans(\mathfrak{P}) = \{v \in V \mid v \in P,\ v^- \in P^* \text{ or } v^- \in P,\ v \in P^*\}$

- $max(\mathfrak{P}) = \{v \in trans(\mathfrak{P}) \mid lk_{\pm}(v) = L\}$ which is precisely the (nonempty by definition) set of possible basepoints for $\mathfrak{P}$.

Finally, we say that $\mathfrak{P}$ *splits* a vertex $v \in V$:if $v \in trans(\mathfrak{P})$.

Remark 2. The fourth assumption in the definition can be rephrased as : every vertex in $trans(\mathfrak{P})$ is isolated in $\Gamma \setminus (L \cap V)$, or even, if the partition is based at b : for every $v \in trans(P)$, $v \lhd_f b$.

One has a partition $cis(\mathfrak{P}) \sqcup trans(\mathfrak{P}) \sqcup (L \cap V) = V$.

Finally, any two elements of $max(\mathfrak{P})$ are fold-equivalent, but this set is not necessarily an entire fold-equivalence class. That way, we can extend the orderings $\lhd_f$ and $\lhd_t$ to the set of generators and Γ-Whitehead partitions by setting for example $v \lhd_t \mathfrak{P}$:if for any $w \in max(\mathfrak{P})$, $v \lhd_t w$. This doesn't depend on the choice of w, and applies to both orders, even in the case of comparing two partitions.

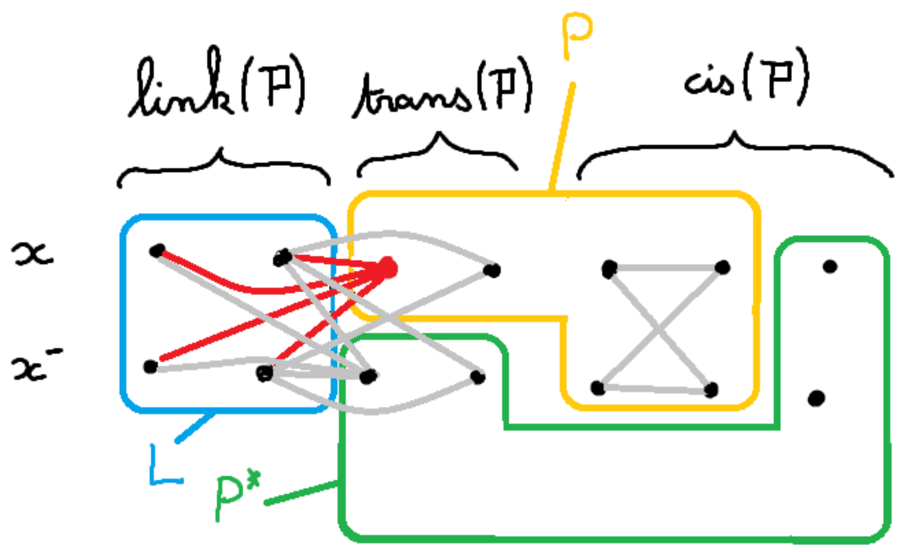

Figure 1: A Whitehead partition.
The two series of dots corrrespond to generators and their inverses. The red dot is the basepoint, whose link is the entire blue zone. There cannot be any edge joining the green and yellow zones.

Combined with Lemma 1, this yields :

Corollary 3. *If a Whitehead partition $\mathfrak{P}$ over a graph $\Gamma = (V, E)$ splits a twist-dominant vertex $v \in V$, then $max(\mathfrak{P}) = \{v\}$.*

Definition 4. Let $\mathfrak{P} = (P, P^*, L)$ and $\mathfrak{Q} = (Q, Q^*, M)$ be two distinct Whitehead partitions over the same graph Γ.

$\mathfrak{P}$ and $\mathfrak{Q}$ are called *square-compatible*[2] :if for every $v \in max(\mathfrak{P})$, $w \in max(\mathfrak{Q})$, v and w are distinct and adjacent in Γ.

Likewise, a vertex $v \in V$ and a partition $\mathfrak{P}$ are called *square-compatible* :if $v \notin max(\mathfrak{P})$ and for every $w \in max(\mathfrak{P})$, v and w are adjacent in Γ. For two vertices $v, w \in V$, *square-compatible* means adjacent in Γ.

Finally, $\mathfrak{P}$ and $\mathfrak{Q}$ are called *vertex-compatible* if they are not square-compatible and one side of $\mathfrak{P}$ is contained in one side of $\mathfrak{Q}$.

Two partitions are called *compatible* if they are either square-compatible or vertex-compatible.

Remark 3. In the definitions of square-compatibility, every occurrence of "for every" can actually be replaced by "there exists", which can be more amenable.

One can prove (using Lemma 3.4 of [2]) that if $\mathfrak{P}$ and $\mathfrak{Q}$ are not square-compatible and, say, $P \subseteq Q$, then necessarily $Q^* \subseteq P^*$, which ensures that vertex-compatibility is a symmetric relation. It can also be rephrased asking for a side of $\mathfrak{P}$ and a side of $\mathfrak{Q}$ to be disjoint (in our case P and Q^*).

If $\mathfrak{P}$ and $\mathfrak{Q}$ are vertex-compatible, exactly one of these four sets is empty : $P \cap Q$, $P \cap Q^*$, $P^* \cap Q$, $P^* \cap Q^*$.

2.3 Right-angled Artin groups

This is a class of finitely generated groups, which is designed to interpolate between two very important families of finitely generated groups : free groups $(F_n)_{n \geq 0}$ and free abelian groups $(\mathbb{Z}^n)_{n \geq 0}$.

Definition 5. Let $\Gamma = (V, E)$ be a finite simple graph. The *right-angled Artin group* (*RAAG*) A_Γ associated to Γ is given by the following group presentation :

$$A_\Gamma := \langle v \in V \mid \forall \{v, w\} \in E,\, vw = wv \rangle$$

The group elements induced by elements of V will be called *generators.*

Example. If Γ is discrete, A_Γ is isomorphic to the free group $F_{|V|}$. If Γ is complete A_Γ is isomorphic to $\mathbb{Z}^{|V|}$.

The class of RAAGs is stable under finite free products (by considering the disjoint union of graphs), and finite direct products (by adding to the disjoint union all possible edges from one term to another).

Each RAAG can be seen conveniently as the fundamental group of a locally CAT(0) cube complex[3] defined as follows :

Definition 6. The *Salvetti complex* $\mathbb{S}_\Gamma$ associated to a finite simple graph $\Gamma = (V, E)$ is the subcomplex of the torus $T^{|V|} = (\mathbb{R}/\mathbb{Z})^{|V|}$, with its standard cubical cell structure and with edges labeled by V, where are kept only the cubes whose edge labels form a clique in Γ (or equivalently pairwise commute in A_Γ).

Note that $\mathbb{S}_\Gamma$ always includes the vertex and the $|V|$ edges of the torus, because any set with less than two elements is a clique.

Example. If Γ is complete with n vertices, no cube is removed and $\mathbb{S}_\Gamma$ is the entire n-torus. If Γ is discrete, $\mathbb{S}_\Gamma$ is the 1-skeleton of T^n, which is a graph with one vertex and n loops.

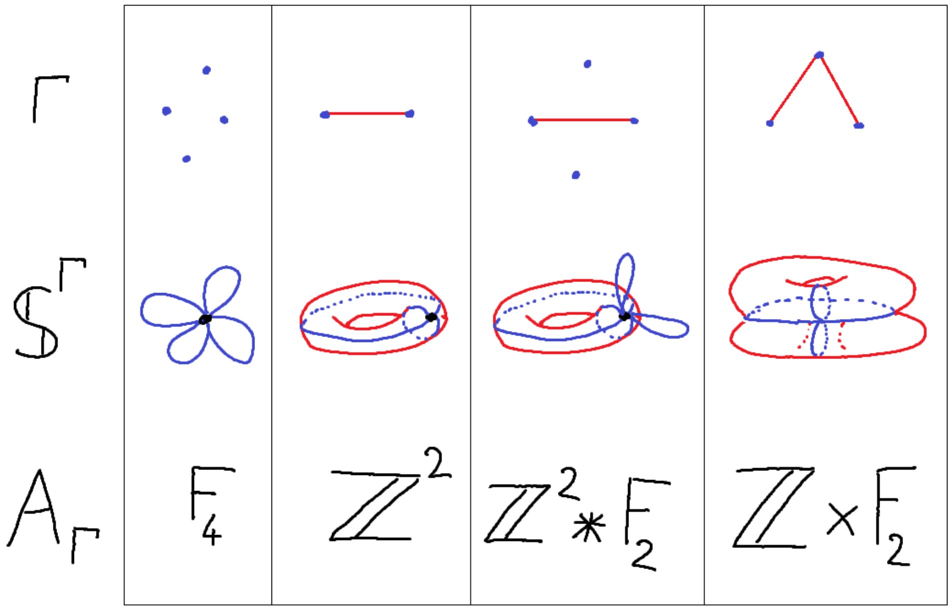

Figure 2: Exemple of right-angled Artin groups,
with corresponding Salvetti complexes.

Proposition 4. *For every finite simple* Γ*,* $\mathbb{S}_\Gamma$ *is connected, and its universal cover is a CAT(0) space, hence is contractible. The fundamental group of* $\mathbb{S}_\Gamma$ *is* A_Γ *and all its higher homotopy groups are trivial.*

As the relators in a RAAG are quite elementary, one can expect that commutation between elements is well-characterized, namely coming only from the commutation of generators, and the fact that two powers of any element in a group always commute.

Proposition 5 (Centralizer theorem, special case, [6], III)**.** *Let* $u \in A_\Gamma$ *be represented by a cyclically reduced word* w*. There exists a writing* $w = w_1 w_2 \dots w_n$ *such that the* $u_i \in A_\Gamma$ *representing the* w_i *pairwise commute and such that the centralizer of* u *is generated by the* u_i *and the generators* $v \in V$ *which commute with every letter of* w*.*

The theorem is in fact more precise and gives a characterization of the w_i.

Corollary 6. *The center of* A_Γ *is generated by the set* Z *of generators which commute with every generator.*

Proof. By the theorem, the centralizer of a generator $v \in V$ is the subgroup generated by the star of v. The center of A_Γ, being the intersection of all element centralizers, is contained in the intersection of all such subgroups, which is generated by the intersection of all stars, *i.e.* the vertices of Z. Conversely every generator in Z commutes with every generator, hence centralizes the whole group. □

[2]This is the notion of *commuting partitions* in [2] and [3] we renamed for the sake of clarity
[3]We will not recall the basics of CAT(0) geometry, [4] is an extensive reference

3 Automorphisms of RAAGs

From now on, we will give ourselves a fixed finite simple graph $\Gamma = (V, E)$ and study the properties of the associated RAAG A_Γ, starting with its automorphism group.

3.1 A generating set for the automorphism group

Definition 7. The *automorphism group* $Aut(A_\Gamma)$ has a normal subgroup whose elements are called *inner automorphisms*, $Inn(A_\Gamma) := \{g \mapsto hgh^{-1} \mid h \in G\}$. The *outer automorphism group*, $Out(A_\Gamma)$ of A_Γ is the quotient of the former by the latter, which fits in the exact sequence :

$$1 \to A_Z \simeq Z(A_\Gamma) \to A_\Gamma \to Aut(A_\Gamma) \to Out(A_\Gamma) \to 1$$
$$h \mapsto (g \mapsto hgh^{-1})$$

Example. The case of a complete graph boils down to the general linear group : $Aut(\mathbb{Z}^n) \simeq Out(\mathbb{Z}^n) \simeq GL_n(\mathbb{Z})$

Now we highlight some specific elements of $Aut(A_\Gamma)$:

Definition 8. The following automorphisms are described by the mapping of generators, the preservation of relations being easily verified.

- For $v \in V$, the automorphism sending $v \mapsto v^{-1}$ and fixing all the generators $w \in V \setminus \{v\}$ is called an *inversion*.
- For ϕ an automorphism of Γ (sending bijectively vertices to vertices and adjacency-preserving), the automorphism of A_Γ sending each generator $v \in V$ to $\phi(v)$ is an extension of ϕ still called a *graph automorphism*.
- For $v, w \in V$ such that $v \lhd w$, the automorphism sending $v \mapsto vw$ and fixing all generators different from v is called a *dominated (right-)transvection*. More precisely, it is called a *fold* when $v \lhd_f w$ and a *twist* when $v \lhd_t w$.
- For $v \in V$ and $C \subset V$ a connected component of $\Gamma \setminus st(v)$, the automorphism sending every generator w belonging to C to its conjugate vwv^{-1} and fixing all other generators is called a *partial conjugation*.

Remark 4. In the abelian case, transvections (which are all dominated) coincide with their original meaning, and are known to generate $SL_n(\mathbb{Z})$. Adding inversions, we get a generating set for $GL_n(\mathbb{Z}) = Aut(\mathbb{Z}^n)$. There are no partial conjugations other than the identity, and graph automorphisms are simply permutation matrices.

In the free case, any connected component has only one vertex, hence a partial conjugation is a product of a right-transvection and a left-transvection. The latter can be obtained by conjugating by an inversion the group inverse of a right-transvection. The three other types are called *Nielsen elementary transformations.* Nielsen found them to generate $Aut(F_n)$ in [5] and even gave a (rather complicated) presentation for this generating set.

The following theorem will help us reduce the study of actions of the infinite group $Out(A_\Gamma)$ to that of a finite number of generators.

Theorem 7 (Laurence, [7], after Servatius, [6])**.** *The finite set of inversions, graph automorphisms, dominated transvections and partial conjugations is generating the group $Aut(A_\Gamma)$.*

Obviously, their classes *modulo* inner automorphisms generate $Out(A_\Gamma)$.

Definition 9. The group of *untwisted outer automorphisms* (or *long-range automorphisms*) $U(A_\Gamma)$ is the subgroup of $Out(A_\Gamma)$ generated by (the classes of) inversions, graph automorphisms, folds and partial conjugations. The group of *short-range automorphism* is the subgroup of $Out(A_\Gamma)$ generated by (the classes of) twists.

3.2 Whitehead automorphisms

We now turn to some specific generators of $U(A_\Gamma)$ called *Whitehead automorphisms*, that appear notably in a finite presentation of $Aut(A_\Gamma)$ found by Day in [8].

Definition 10. Let $(\mathfrak{P} = (P, P^*, L), b)$ be a based Whitehead partition. The *Whitehead automorphism* associated to it is defined by[4] :

$$\varphi \colon v \in V \mapsto \begin{cases} v^{-1} & \text{if } v = b \\ vb^{-1} & \text{if } v \in trans(\mathfrak{P}) \cap P,\ v \neq b \\ bv & \text{if } v \in trans(\mathfrak{P}) \cap P^* \\ bvb^{-1} & \text{if } v \in cis(\mathfrak{P}) \cap P \\ v & \text{otherwise (if } v \in (cis(\mathfrak{P} \cap P^*) \sqcup L) \end{cases}$$

Of course, one has to verify that this map sends adjacent vertices to commuting words, which is straightforward. Moreover, as $\varphi(b) = b^{-1}$, $\varphi \circ \varphi$ fixes every generator hence φ is involutive and a well-defined automorphism of A_Γ.

Lemma 8. *$U(A_\Gamma)$ is generated by (the images in the quotient of) graph automorphisms, inversions and Whitehead automorphisms.*

Proof. The Whitehead automorphism associated to $(\mathfrak{P}, b)$ can be written as a product of an inversion, a fold per element of $trans(\mathfrak{P})$ and a partial conjugation per connected component of $cis(\mathfrak{P}) \cap P$ in $\Gamma \setminus (L \cap V)$, hence belongs to $U(A_\Gamma)$.

Conversely, if $v \lhd_f w \in V$, choosing $P = \{v, w\}$ and w as a basepoint gives a Whitehead automorphism sending $w \mapsto w^{-1}$, $v \mapsto vw^{-1}$ and fixing every other generator. Composing that automorphism with the inversion of w yields a generic fold.

Moreover, if $w \in V$ and $C \subset V$ is a connected component of $\Gamma \setminus st(w)$, choosing $P = C \cup C^- \cup \{w\}$ and w as a basepoint and composing by an inversion generates a generic partial conjugation, except in the case where $C = \Gamma \setminus st(w)$, where P^* would have only one element. In that case, the partial conjugation is actually a conjugation, thus corresponds to the identity in $Out(A_\Gamma)$. □

3.3 Topological realization

It is a straightforward algebraic topology fact that pointed-homotopic pointed maps induce the same map on fundamental groups. However, the following statement is both more fitting to our setup (by removing basepoints) and more precise, using the fact stated in Proposition 4 that $\mathbb{S}_\Gamma$ is a classifying space for A_Γ : it provides a topological intepretation of $Out(A_\Gamma)$

[4]This is the convention used by [2] and [3] for a more geometric interpretation. Usually the basepoint is mapped to itself without inversion.

Proposition 9. *Let x be the vertex of $\mathbb{S}_\Gamma$, so that $A_\Gamma = \pi_1(\mathbb{S}_\Gamma, x)$. The group $Out(A_\Gamma)$ identifies with the group $hHeq(\mathbb{S}_\Gamma)$ of (free) homotopy equivalences up to (free) homotopy. This identification sends the homotopy class of f to the coset modulo inner automorphisms of $\phi(f) = i_f \circ f_*\colon \pi_1(\mathbb{S}_\Gamma, x) \to \pi_1(\mathbb{S}_\Gamma, x)$ where $i_f\colon \pi_1(\mathbb{S}_\Gamma, f(x)) \to \pi_1(\mathbb{S}_\Gamma, x)$ is a fixed automorphism given by the (non-canonical) choice of a path from x to $f(x)$ in A_Γ. However the choice of such a path doesn't affect the resulting coset in $Out(A_\Gamma)$.*

Proof. If p and q are two paths from x to $f(x)$, let $l = pq^{-1}$, a loop based at x. For every loop m based at x, let $i_p\colon \pi_1(\mathbb{S}_\Gamma, f(x)) \to \pi_1(\mathbb{S}_\Gamma, x)$ map $[m]$ to $[pmp^{-1}]$ and i_q map $[m]$ to $[qmq^{-1}]$ likewise. Hence, $i_p = c_{[l]} \circ i_q$ where c_l is the group conjugacy by the element $[l]$. This ensures that $\phi(f)$ is canonically defined up to conjugacy.

One has now to prove that such $\phi(f)$ is indeed an automorphism. As f is a homotopy equivalence, let g be its homotopy inverse, and let H be a (free) homotopy from $id_{\mathbb{S}_\Gamma}$ to $g \circ f$. Let p_f be a path from x to $f(x)$, p_g a path from x to $g(x)$ and p the path from x to $g \circ f(x)$ given by H. Let $l = p_g(g_*(p_f))p^{-1}$, a loop based at x. For every loop m based at x the following equality holds :

$$\begin{aligned}\phi(g) \circ \phi(f)([m]) &= i_g \circ g_* \circ i_f \circ f_*([m]) \\ &= [p_g](g_*([p_f](f_*([m]))[p_f^{-1}]))[p_g^{-1}] \\ &= [l][p](g \circ f)_*([m])[p^{-1}][l^{-1}] \\ &= c_{[l]}([p](g \circ f)_*([m])[p^{-1}])\end{aligned}$$

However by definition of p, H gives a pointed homotopy between m and $p(g \circ f \circ m)p^{-1}$, hence those two loops based at x have the same class and $\phi(g) \circ \phi(f) = c_{[l]}$ which is trivial up to conjugacy. So is $\phi(f) \circ \phi(g)$ by the same argument, hence the image of f is indeed an element of $Out(A_\Gamma)$, with inverse the image of g. An argument of the same nature proves that the mapping is a group homomorphism and only depends on the homotopy class of f hence is well defined.

Finally, one has to prove the mapping one-to-one. To prove surjectivity, given an automorphism $\psi \in Aut(A_\Gamma)$, one can create f sending x to itself and sending each 1-cell in $\mathbb{S}_\Gamma$ (*i.e.* each generator of A_Γ) to a representative of the image by ψ of the generator, effectively inducing ψ. One has now to extend the mapping to higher-dimensional cells, which requires some slight work for dimension 2 using that $\psi(1) = 1$ and is immediate in higher dimensions as $\pi_k(\mathbb{S}_\Gamma)$ is trivial for $k > 1$.

To prove injectivity, considering some f in the kernel, $\phi(f)$ is a conjugacy by the class of some loop l, one has to define a homotopy H between f and $id_{\mathbb{S}_\Gamma}$, moving x along l, which makes H defined at least on the 1-skeleton of $\mathbb{S}_\Gamma \times [0, 1]$. The extension to higher-dimensional cells is as straightforward as before, which concludes the proof. □

Remark 5. This result (valid for all classifying spaces) underlines the interest of considering outer automorphism groups instead of general automorphism groups. The latter would fit into such a correspondence but with pointed topological spaces, which would create some trouble with the following construction. However [10] deals with this difficulty of tracking basepoints and defines some concept of *Auter Space* with an action of $Aut(G)$, in the special case of free groups.

4 Γ-complexes

We now characterize a class of cube complexes with fundamental group A_Γ which will play a central role in the construction of Outer space.

4.1 Blow-ups

We saw earlier in section 2.3 the definition of the Salvetti complex $\mathbb{S}_\Gamma$, as well as the notion of *compatible partitions*. This section will give a purpose to that compatibility, using partitions to alter $\mathbb{S}_\Gamma$ into several homotopically equivalent spaces called *blow-ups*.

Definition 11. Let $n \geq 0$ be an integer and $\Pi = (\mathfrak{P}_1 \dots \mathfrak{P}_n)$ a family of distinct pairwise-compatible Γ-Whitehead partitions. A *region* is a family of sides $(R_1, \dots, R_n) \in \prod_{i=1}^{n} \{P_i^*, P_i\}$ such that for all $1 \leq i \neq j \leq n$, if $\mathfrak{P}_i$ and $\mathfrak{P}_j$ are vertex-compatible, $R_i \cap R_j \neq \emptyset$. We impose nothing if $\mathfrak{P}_i$ and $\mathfrak{P}_j$ are square-compatible.

The *blow-up core* $\mathbb{C}^\Pi$ of $\mathbb{S}_\Gamma$ relative to Π is the subcomplex of $[0,1]^n$ spanned by the vertices in $\{0,1\}^n \simeq \prod_{i=1}^{n} \{P_i^*, P_i\}$ corresponding to regions. Edges corresponding to a change in the k^{th} coordinate are all labeled $\mathfrak{P}_k$ and oriented from P_k^* to P_k.

Let $v \in V$ be a vertex. A region $R = (R_1, \dots R_n)$ is called *terminal* for v :when for every index i such that $\mathfrak{P}_i$ splits v, R_i contains v. For such an R, let R_v^- be the family of sides obtained from R by choosing the opposite side (containing v^-) at the indices that split v (one may have $R_v^- = R$ if no partition splits v.)

Lemma 10 (3.9 in [2])**.** *Any incomplete family of sides that verify the definition of a region can be completed into a region.*

Corollary 11 (3.10 in [2])**.** *With the setup of the definition, for every $v \in V$ there exists at least one terminal region for v, and for every such region R, R_v^- is still a region. Moreover, if R is terminal for two commuting vertices $v, w \in V$, R_v^- is also terminal for w.*

Definition 12. With the same setup, the *blow-up* $\mathbb{S}^\Pi$ of $\mathbb{S}_\Gamma$ relative to Π is the labeled oriented cube complex obtained by :

1. Starting with $\mathbb{C}^\Pi$
2. For every $v \in V$, for every region R terminal for v, gluing on $\mathbb{C}^\Pi$ a directed edge labeled v from R_v^- to R.
3. For every $k \geq 2$, gluing a k-cube over each possible 1-skeleton where parallel edges have identical label and orientation, and labels are pairwise square-compatible (except when all the labels are partitions, because this k-cube already exists in $\mathbb{C}^\Pi$)

Proposition 12. *Let $\Pi = (\mathfrak{P}_1, \dots, \mathfrak{P}_n)$ be a family of distinct pairwise-compatible Γ-Whitehead partitions.*

1. *Changing the order of $\mathfrak{P}_i$'s doesn't change the isomorphism type of $\mathbb{S}^\Pi$ or $\mathbb{C}^\Pi$*
2. *$\mathbb{S}^\Pi$ and $\mathbb{C}^\Pi$ are connected*
3. *$\mathbb{C}^\Pi$ and $\mathbb{S}^\Pi$ are locally CAT(0) cube complexes*
4. *The set of edges having a given label l is exactly the set of edges dual to some hyperplane H_l*

5. *Every maximal set of pairwise square-compatible labels (which are either partitions or vertices) is the label set of a unique cube in* $\mathbb{S}^\Pi$

Here we merge the notation for a hyperplane (as a formal equivalence class of edges) and the notation for its carrier (as the full subcomplex spanned by those edges) into a single notation H_l, l being the common label of all such edges.

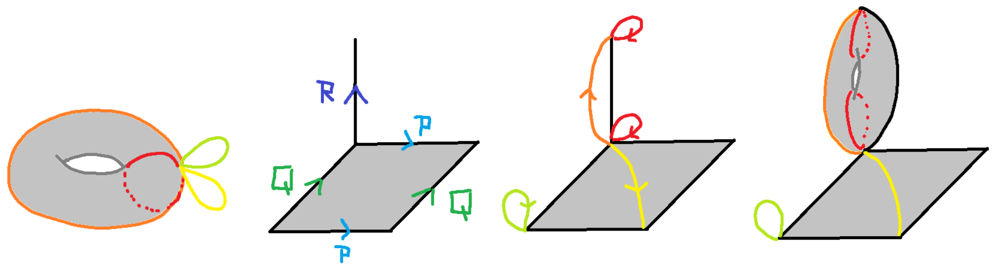

Figure 3: The construction of a blow-up.
From left to right : $\mathbb{S}_\Gamma$, $\mathbb{C}^\Pi$ with three partitions, two of which are square-compatible, the complex obtained after step 2 and the full blow-up $\mathbb{S}^\Pi$.
The two partitions $\mathfrak{P}$ and $\mathfrak{Q}$ are distinct, as they differ by the orientation of the yellow edge.

Proof. 1. The isomorphism is clearly deduced from the permutation of the $\mathfrak{P}_i$'s

2. Take two vertices of $\mathbb{C}^\Pi$, which correspond to regions, and induct over the number of sides on which those regions differ. If that number is 1, both vertices of some edge in $[0,1]^n$ are present in $\mathbb{C}^\Pi$ and so is the edge. If it is more than 1, say up two permutation that the regions write : $R = (C_1, \dots C_{k-1}, D_k, \dots D_n)$ and $R' = (C_1, \dots C_{k-1}, D'_k, \dots D'_n)$ where D_i and D'_i are opposite sides. If $\mathfrak{P}_k$ and $\mathfrak{P}_{k+1}$ are vertex-compatible, by Remark 3, one of the intersections $D_k \cap D'_{k+1}$, $D'_k \cap D_{k+1}$ is non-empty, say the first one. Then $(C_1, \dots C_{k-1}, D_k, D'_{k+1})$ is an incomplete region and can be completed to a full region R'', by Lemma 10, which is connected to both R and R' by induction hypothesis. If they are square-compatible, the given R'' works whatsoever. Hence $\mathbb{C}^\Pi$ is connected, and so is $\mathbb{S}^\Pi$, having the same vertex set.

3. Let R be a vertex of $\mathbb{S}^\Pi$, and l_1, l_2 two square-compatible labels of edges which meet at R. To simplify assume both edges point towards R, and start from R_1, R_2. If the labels are of two square-compatible partitions, the existence of the three regions R, R_1, R_2, corresponding to three choices of sides of those partitions guarantees the existence of the fourth, and hence of the square in $\mathbb{C}^\Pi$. If the labels are of two square-compatible (commuting) vertices, the last part of Corollary 11 guarantees the same conclusion. Finally for a partition $\mathfrak{P}$ and a vertex v which are square-compatible, v must belong to the link of $\mathfrak{P}$, so that R_1 is still terminal for v which again provides the fourth vertex.

   For more than two pairwise-compatible labels, one can apply the previous argument several times to build progressively more and more of the 2-skeleton of the cube (with accurately labeled edges), which is then filled-in by definition of the construction itself, proving that $\mathbb{S}^\Pi$ is CAT(0). The argument for $\mathbb{C}^\Pi$ is strictly contained in this one.

4. By construction, two edges dual to the same hyperplane must have the same label. Conversely, if two edges have the same label l, connect their targets by an edge path in $\mathbb{C}^\Pi$ given by Statement 2. If $l \in V$ is a vertex, the edge path corresponds to switching one side for its opposite successively for partitions, each partition being switched at most once. In particular, since both endpoints are terminal for v, partitions which split v are never switched, and a hyperplane associated to v runs all along the path. A similar argument holds if l is a partition, guaranteeing the converse.

5. This is Proposition 3.2 and Corollary 3.5 in [3], which mostly follow the same proof strategies as above.

□

Definition 13. A Γ-*complex* is a cube complex X isomorphic to $\mathbb{S}^\Pi$ for some (possibly empty) family Π of pairwise-compatible distinct partitions. Neither Π nor the isomorphism have to be unique or fixed in the definition. A choice of such Π and isomorphism will be called a *blow-up structure* for X.

This notion, which relies on forgetting part of the structure, will prove crucial for the construction of Outer Space. In counterpart, it means that we have to make definitions as independant from labellings as possible from now on, in order to transfer them from blow-ups to Γ-complexes. For example the following lemma will prove useful :

Lemma 13. *In a blow-up* $\mathbb{S}^\Pi$*, two hyperplanes, seen as unions of midplanes of cubes, intersect if and only if their labels are square-compatible. That is, square-compatibility can be detected with the cubical structure only, without the labels.*

Moreover, if H_1 *and* H_2 *are two hyperplanes in a* Γ*-complex* X*, the relation* $l_1 \lhd_f l_2$ *for* l_1 *and* l_2 *the labels of* H_1 *and* H_2 *in some blow-up structure for* X *is verified if and only if it is verified for* every *possible blow-up structure. The same is true for the relation* $l_1 \lhd_t l_2$*.*

Proof. An intersection of hyperplanes guarantees that two of their dual edges meet at a vertex and span a square there, hence must be square-compatible. Conversely, two square-compatible labels are part of a maximal family of square-compatible labels, hence by Proposition 12 appear in a cube, hence their dual hyperplanes intersect.

Moreover, it can be checked that the relation $l_1 \lhd_f l_2$ is equivalent to the statement that every label l_3 square-compatible with l_1 is also square-compatible with l_2, and the previous result allows transferring that statement to the hyperplane structure, making it independant of the chosen blow-up structure.

The proof for $l_1 \lhd_t l_2$ is more complex and will be found as Corollary 4.4 in [3]. □

4.2 Collapses

Now we want to define an operation in a Γ-complex somewhat inverse to blowing partitions up, namely collapsing along hyperplanes. Let us first consider the fundamental example of a one-partition blow-up.

Example. Let $\mathfrak{P}$ be a Γ-Whitehead partition and consider the hyperplane $H_\mathfrak{P}$ in $\mathbb{S}^\mathfrak{P}$. By construction of the blow-up, there is a single edge $e_\mathfrak{P}$ which joins the only two vertices of $\mathbb{S}^\mathfrak{P}$, in particular which isn't a loop. The edges of the carrier of $H_\mathfrak{P}$ are $e_\mathfrak{P}$ and all the edges labeled by an element v square-compatible with $\mathfrak{P}$. By Proposition 12.5, they form squares with $e_\mathfrak{P}$ twice on the boundary, hence there are exactly two of each label, a loop at each vertex. Let A be the subcomplex spanned by those loops at one vertex : $H_{\mathfrak{P}}$ then decomposes as a product $A \times e_{\mathfrak{P}}$.

Besides, labels in A are exactly the labels which appear twice by construction of the blow-up, all other labels appear once. Hence, with $e_{\mathfrak{P}}$ being identified with $[0,1]$, by deleting $A \times (0,1)$ and identifying $A \times \{0\}$ and $A \times \{1\}$, one recovers $\mathbb{S}_\Gamma$. Seen differently, there is a map $c_{\mathfrak{P}} \colon \mathbb{S}^{\mathfrak{P}} \to \mathbb{S}_\Gamma$ projecting $A \times e_{\mathfrak{P}}$ onto A and leaving the rest unchanged. This map is called *collapse of* $H_{\mathfrak{P}}$ and is a homotopy equivalence.

Other homotopy equivalences $\mathbb{S}^{\mathfrak{P}} \to \mathbb{S}_\Gamma$ exist, for example $c_{\mathfrak{P}} \circ i$ where i is an automorphism of $\mathbb{S}^{\mathfrak{P}}$. However, they are somewhat trivial if i fixes the edge $e_{\mathfrak{P}}$ as they can be rewritten $\tilde{i} \circ c_{\mathfrak{P}}$ with $\tilde{i}$ an automorphism of $\mathbb{S}_\Gamma$: these other collapses will be considered *equivalent* in some sense that will be made explicit in Section 4.3. The only remaining collapses of interest are those where i exchanges $e_{\mathfrak{P}}$ with some edge e_v and fixes every other edge. Then, the labels which are square compatible to $\mathfrak{P}$ and those compatible to v coincide, that is $lk(v) = lk(b)$ where b is any basepoint of $\mathfrak{P}$, or also : $v \in max(\mathfrak{P})$. We denote the associated collapse map by c_v.

Moreover, the homotopy inverse $c_{\mathfrak{P}}^{-1}$ is well-defined up to homotopy, so $c_v \circ c_{\mathfrak{P}}^{-1}$ lies in a single homotopy class of homotopy equivalences $\mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$, that is, according to Proposition 9, induces a unique element of $Out(A_\Gamma)$. A simple tracking of the image of edges shows that this element is exactly the class of the *Whitehead automorphism corresponding to the based partition* (P, v). This geometric realization of Whitehead automorphisms highlight the central role they play in the theory, as "elementary moves".

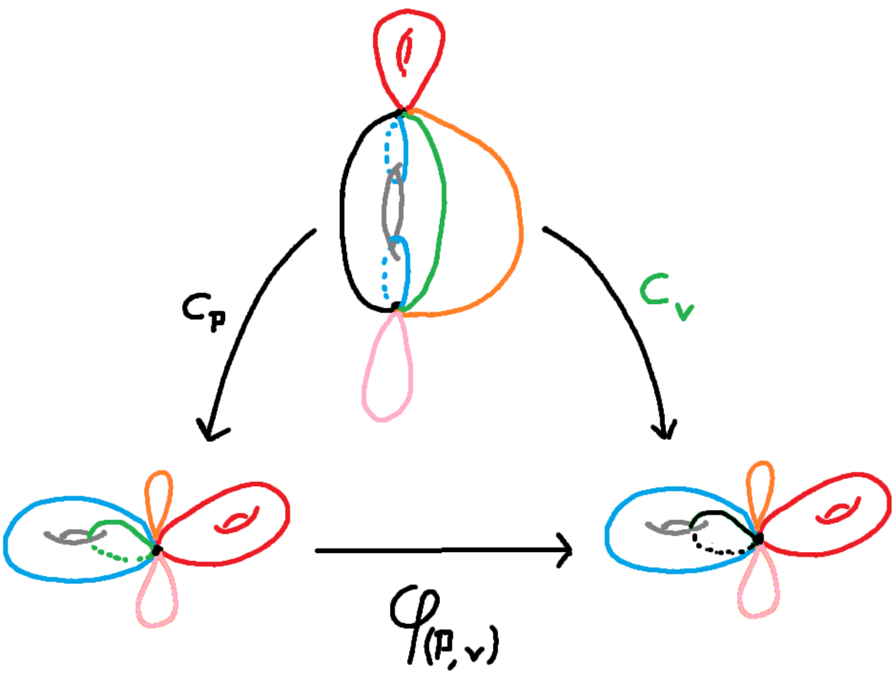

Figure 4: Realizing a Whitehead automorphism using two collapse maps. The green edge corresponds to the basepoint v, the black edge to the partition $\mathfrak{P}$, blue edges represent L, orange edges represent $trans(\mathfrak{P})$, pink edges represent $cis(\mathfrak{P}) \cap P$ and red edges represent $cis(\mathfrak{P}) \cap P^*$

.

Now we deal with the general case.

Definition 14. A family of hyperplanes in a cube complex X is called *collapsible*

:if the two following conditions hold :

- Every carrier of a hyperplane in the family decomposes as an embedded product subcomplex $e \times A$, where e is its dual edge
- Every loop path going through edges dual to the family only is nullhomotopic.

The first condition is always verified in a Γ-complex (because two edges with identical labels cannot meet at a vertex), but we will need this more general setup, as we cannot ensure that the result of a collapse will be a Γ-complex itself[5]. For example, in a Γ-complex, a single hyperplane is collapsible if and only if its dual edge isn't a loop.

Lemma 14. *Let X be a cube complex, H a collapsible hyperplane of X, and write the carrier of H as $e \times A$.*

1. *Contracting the subcomplex $e \times A$ to A without changing any other cell gives a new cube complex X_H and a cellular homotopy equivalence $c_H : X \to X_H$, which induces a bijection between hyperplanes in X_H and hyperplanes in X different from H*
2. *c_H establishes a bijection between collapsible families in X_H and collapsible families in X containing H*
3. *If H was labeled $\mathfrak{P}$ in some blow-up structure $\mathbb{S}^{\Pi}$ for X, X_H is isomorphic to $\mathbb{S}^{\Pi \setminus \{\mathfrak{P}\}}$, in a coherent way with respect to the action of $c_{\mathfrak{P}}$ on labels.*

Proof.

1. Like in the previous example, identifying e with $[0,1]$, deleting $A \times (0,1)$ and gluing the remaining copies of A in X gives a new subcomplex X_H with one hyperplane less, and a homotopy equivalence c_H as wanted (the homotopy being obtained by varying the width of H to 0). The construction makes obvious that cells outside the carrier of H are untouched.

   The only edges in X that are contracted by c_H are dual to H, and the only edges that are identified are parallel. Hence every other class of parallel edges (that is, every hyperplane different from H) is preserved, and every edge in X_H comes from an edge in X (not dual to H) hence no new hyperplane is created.

2. A collapsible family in X_H lifts by the bijection of 1. to a family of hyperplanes of X, to which we can add H, and conversely, a family of hyperplanes in X with H removed descends to X_H. c_H being a homotopy equivalence, it must send nullhomotopic paths to nullhomotopic paths, which guarantees one of the conditions. For the other, in one direction, adding cubes cannot destroy the product structure on a hyperplane. The other condition comes in play to prove that collapsing doesn't destroy the structure either (which it clearly true already for edges which lie in the carrier of the hyperplane).

3. Collapsing every edge $e_{\mathfrak{P}}$ amounts to identifying regions which differ only by the side of $\mathfrak{P}$ considered. The identification then holds by definition of the blow-up construction.

□

[5]There probably exists a weak condition on X that guarantees the collapse to remain a Γ-complex, but it hasn't been made explicit yet.

Definition 15. If $\mathcal{H} = (H_1, \dots H_n)$ is a collapsible family of hyperplanes of X, let $c_{\mathcal{H}}$ be the successive composition of the collapse of H_1 in X, then of $c_{H_1}(H_2)$ in X_{H_1}, carrying on until every hyperplane has been collapsed. $c_{\mathcal{H}}$ is called the *collapse map* of $\mathcal{H}$ in X and its image $X_{\mathcal{H}}$ is the *collapse* of X along $\mathcal{H}$.

Lemma 15. *1. The collapse map of a family $\mathcal{H}$ and its target $X_{\mathcal{H}}$ are independent from the order of the hyperplanes in H.*

2. If X has a blow-up structure $\mathbb{S}^{\Pi}$, the family of hyperplanes labeled by Π is collapsible, the associated collapse is isomorphic to $\mathbb{S}_{\Gamma}$, and the labellings are compatible with c_{Π}. c_{Π} is called the standard collapse *of $\mathbb{S}^{\Pi}$.*

Proof. 1. This comes from the preservation of cells which don't feature an edge dual to H.

2. This is simply an induction over Lemma 14.3

□

Proposition 16 (Theorem 4.12 in [2])**.** *Let $\mathcal{H}$ and $\mathcal{K}$ be two collapsible families of hyperplanes over the same Γ-complex X such that $X_{\mathcal{H}}$ and $X_{\mathcal{K}}$ are both isomorphic to $\mathbb{S}_{\Gamma}$. Any element of $\mathcal{H}$ can be replaced by some element of $\mathcal{K}$ so that the resulting set $\mathcal{H}'$ is still collapsible and $X_{\mathcal{H}'}$ is still isomorphic to $\mathbb{S}_{\Gamma}$.*

Remark 6. This result classifies as an exchange lemma, which derives from a characterization of collapsible families of hyperplanes by the fact that their dual edges form a forest in the 1-skeleton of the Γ-complex. The fact that the collapse is isomorphic to $\mathbb{S}_{\Gamma}$, guarantees a form of maximality, or spanning, of the forest, hence the name of a *tree-like* family given by [2].

However, there is a weakness here, as [2] doesn't specify how the different isomorphisms to $\mathbb{S}_{\Gamma}$ play out with each other, in particular, why the difference between them would be untwisted. Provided this obstruction be overcome, collapses can play a symmetric role to blow-ups.

Lemma 17. *If X admits two hyperplanes H_1 and H_2 such that both collapses X_{H_1} and X_{H_2} are isomorphic to $\mathbb{S}_{\Gamma}$, the outer automorphism induced by $c_{H_1} \circ c_{H_2}^{-1}$ is either trivial or a Whitehead automorphism.*

Proof. If X can be proven to be a blow-up with respect to the edges dual to H_1, this boils down to the example at the start of this section (*modulo* the obstruction of exactly identifying what the isomorphisms to $\mathbb{S}_{\Gamma}$ are). The Whitehead partition $\mathfrak{P}$ that would correspond to the hyperplane H_1 is almost entirely determined : whether edges are loops or not describes which generators lie in $cis(\mathfrak{P})$ or $trans(\mathfrak{P})$. Almost all the conditions in the definition are automatically verified (an edge dual to H_2 can play the role of a basepoint.) The only problem would be that one side of $\mathfrak{P}$ can have only one element. But in that case, necessarily one vertex of X is of degree 2 and the two collapses are homotopy equivalent, hence the induced automorphism is the identity. □

Corollary 18. *If X, a Γ-complex, has two blow-up structures $\mathbb{S}^{\Pi}$ and $\mathbb{S}^{\Pi'}$ with associated standard collapses c_{Π}, $c_{\Pi'} : X \to \mathbb{S}_{\Gamma}$, the outer automorphism of $\pi_1(\mathbb{S}_{\Gamma}) = A_{\Gamma}$ given by $(c_{\Pi'})_* \circ (c_{\Pi})_*^{-1}$ belongs to $U(A_{\Gamma})$.*

Proof. Let $\mathcal{H}$ the family of hyperplanes of X coming from the $(H_{\mathfrak{P}})_{\mathfrak{P}\in\Pi}$ and $\mathcal{K}$ the family coming from the $(H_{\mathfrak{P}})_{\mathfrak{P}\in\Pi'}$. Upon applying the proposition iteratively, one can let $\mathcal{H} = \mathcal{H}_1, \dots, \mathcal{H}_k = \mathcal{K}$ be collapsible hyperplane families so that two consecutive families differ only by one hyperplane.

Let $\mathcal{H}_1 = (H, H_2, \ldots, H_n)$ and $\mathcal{H}_2 = (H', H_2, \ldots H_n)$. Upon writing the collapses as compositions of single hyperplane collapses, one obtains a homotopy $c_{\mathcal{H}_2} \circ c_{\mathcal{H}_1}^{-1} \sim c_{H'} \circ c_H^{-1}$. Both maps c_H and $c_{H'}$ are from a common cube complex $Y = X_{H_2,\ldots H_n}$ in which H and H' are collapsible with collapse isomorphic to $\mathbb{S}_\Gamma$. The Lemma 17 then concludes. □

4.3 Markings

Definition 16. A *marked* Γ-complex is a Γ-complex X together with a homotopy equivalence $c\colon X \to \mathbb{S}_\Gamma$ such that, given a blow-up structure $\mathbb{S}^\Pi \simeq X$ with its associated standard collapse c_Π, the automorphism induced by $c_* \circ (c_\Pi)_*^{-1}$ belongs to $U(A_\Gamma)$. By corollary 18, this property doesn't depend on the blow-up structure.

Two marked Γ-complexes (X, c) and (Y, d) are *equivalent* if there exists $i\colon X \to Y$ an isomorphism of Γ-complexes such that $d \circ i$ and c are homotopic maps $X \to \mathbb{S}_\Gamma$.

The class of marked Γ-complexes modulo equivalence is a set $\mathcal{X}_\Gamma$, as there are only finitely many Whitehead partitions over Γ. Under the identification of $Out(A_\Gamma)$ with the group of homotopy equivalences $\mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ modulo homotopy, the subgroup $U(A_\Gamma)$ acts on $\mathcal{X}_\Gamma$ by the formula $h \cdot (X, c) = (X, h \circ c)$.

A *(marked) salvetti* is an element of $\mathcal{X}_\Gamma$ which is represented by marked Γ-complexes isomorphic to $\mathbb{S}_\Gamma$.

Given a marked Γ-complex (X, c) and a non-empty collapsible family of hyperplanes $\mathcal{H}$ of X, the image $X_\mathcal{H}$ under this collapse is endowed with the marking $c \circ d$ where $d\colon X_\mathcal{H} \to X$ is a homotopy inverse of the collapse map. We write $(X_\mathcal{H}, c \circ d) < (X, c)$.

Lemma 19.
- *The relation $<$ is compatible with equivalence, and quotients into a strict partial order over $\mathcal{X}_\Gamma$, still denoted $<$. This order is preserved by the action of $U(A_\Gamma)$*
- *Salvettis are exactly the minimal elements of $(\mathcal{X}_\Gamma, <)$ and form a single infinite orbit under $U(A_\Gamma)$*
- *$U(A_\Gamma)$ acts on $\mathcal{X}_\Gamma$ with finite stabilizers*

Proof.
- This stems from the fact that collapses behave well with cube complex isomorphisms (once again irrespective of the labelling). The action composes the marking on the left while the order relies on composition on the right, hence they are compatible.
- In a salvetti, every edge is a loop, hence no hyperplane is collapsible. Moreover every Γ-complex which isn't a salvetti contains a collapsible hyperplane (namely any one dual to an edge labeled by a partition in some isomorphism with a blow-up), hence isn't minimal. Besides, given a salvetti represented by (X, c), where $i\colon \mathbb{S}_\Gamma \to X$ is an isomorphism : $c \circ i\colon \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ is a homotopy equivalence defined up to homotopy, which corresponds to an element h of $U(A_\Gamma)$ by definition, and it is clear that $h \cdot (\mathbb{S}_\Gamma, id)$ and (X, c) are equivalent, thus every salvetti is in the orbit of $[\mathbb{S}_\Gamma, id]$, and the set of salvettis is obviously closed under the action. By properness of the action, $U(A_\Gamma)$ being infinite, the orbit must be infinite as well.
- For every element h of $Stab([\mathbb{S}_\Gamma, id])$ there exists an isomorphism $i\colon \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ such that $h \circ i$ is homotopic to $id_{\mathbb{S}_\Gamma}$. As homotopic elements in $U(A_\Gamma)$ are equal, h lies in the finite set of classes of combinatorial isomorphisms, seen as homotopy equivalences. To be more precise, the elements of this set preserve

edges up to orientation, so they are exactly generated by inversions and graph automorphisms. The stabilizer of any salvetti is conjugate to that of $[\mathbb{S}_\Gamma, id]$ hence finite as well.

For a general $\sigma \in \mathcal{X}_\Gamma$, there is a finite non-empty set S of salvettis σ' which verify $\sigma' < \sigma$. A group element which fixes σ must send S to itself, inducing a permutation in $\mathfrak{S}_S$ which is finite. Moreover, two elements inducing the same permutation differ by an element which has to fix every salvetti in S, hence lie in a finite subgroup. Hence the stabilizer of σ has to be finite. □

Definition 17. The *spine* K_Γ is the simplicial complex[6] whose n-faces are the chains $\sigma_1 < \cdots < \sigma_n$ of elements of $\mathcal{X}_\Gamma$, together with the action of $U(A_\Gamma)$ given by the action on the vertex set $\mathcal{X}_\Gamma$.

Lemma 20. *K_Γ is connected and locally finite. Moreover, the action of $U(A_\Gamma)$ over K_Γ is proper and cocompact.*

Proof. Every vertex of K_Γ is connected to a minimal vertex, that is a salvetti, so it suffices to prove that any two salvettis are connected, or rather that any salvetti $[\mathbb{S}_\Gamma, c]$ is connected to $[\mathbb{S}_\Gamma, id]$, where c is an untwisted homotopy equivalence. Thanks to Lemma 8, we can decompose the automorphism induced by c into a product of graph automorphisms, inversions and Whitehead automorphisms. Conjugating a Whitehead automorphism by a graph automorphism or inversion gives a Whitehead automorphism or its inverse, hence we can write $c = w \circ f$ where f acts as a product of automorphisms and inversions and w acts as a composition of Whitehead automorphisms. Besides, f can be realized as a cube complex automorphism, hence making $(\mathbb{S}_\Gamma, id)$ and $(\mathbb{S}_\Gamma, f)$ equivalent. Using the action, $[\mathbb{S}_\Gamma, c] = [\mathbb{S}_\Gamma, w]$. Up to gluing paths after each other, what is left to study is the case where w is a single Whitehead automorphism. In this case, we saw in the fundamental example starting Section 4.2 that w is induced by $c_b \circ c_\Pi^{-1}$ where (Π, b) is the based partition associated to W. Thus, $[\mathbb{S}_\Gamma, id]$ is connected to $[\mathbb{S}_\Gamma, w] = [\mathbb{S}_\Gamma, c]$ which proves connectedness.

Moreover, the fact that any salvetti can be blown-up in only finitely many ways guarantees that the 1-neighborhood of any vertex in K_Γ is a finite subcomplex, hence the local finiteness.

By local finiteness and the fact that stabilizers of vertices are finite, the action on K_Γ is proper.

Finally, for cocompactness, notice that quotienting $\mathcal{X}_\Gamma$ by the action of $U(A_\Gamma)$ identifies complexes which are isomorphic but with different markings. Thus, there are only finitely many isomorphism types of unmarked Γ-complexes as there are finitely many Γ-Whitehead partitions and a pairwise-compatible set of them has a size bounded in terms of the number of vertices of Γ. As the action is order-preserving, the quotient of K_Γ is still a simplicial complex (built using the same process as K_Γ but with the quotient $\overline{\mathcal{X}_\Gamma}$), with finite vertex set hence finite. □

5 Untwisted Outer Space

The ordered set $\mathcal{X}_\Gamma$ previously constructed, and its simplicial realization K_Γ are very close to the notion of Outer Space we want to obtain, at least for untwisted automorphisms. We will first outline the exact notion, and the associated metric structure, but also prove that topological invariants of Outer Space are already contained in K_Γ, which we will then proceed to prove contractible.

[6] We won't consider the abstract complex here, only its geometric realization

5.1 Definition and reduction to K_Γ

Definition 18. A *rectilinear metric structure* on a Γ-complex X is the data of a positive real number for each hyperplane, called its *width*. It defines a metric d on X by fixing all edges dual to a hyperplane to have its width for length and all cubes to be orthotopes (isometric to an orthogonal product of segments).

A marking on a rectilinear Γ-complex (X, d) is, as before, a homotopy equivalence c between X and $\mathbb{S}_\Gamma$ which induces an untwisted automorphism when composed with any standard collapse of X.

Two rectilinear Γ-complexes (X, d, c) and (X', d', c') are called *equivalent* :when there exists an isometry $i\colon (X, d) \to (X', d')$ such that $c' \circ i$ is homotopic to c. The group $U(A_\Gamma)$ acts as before on equivalence classes on the left.

The *untwisted (unreduced) Outer Space* for A_Γ is the set Σ_Γ of rectilinear Γ-complexes up to equivalence together with the action $U(A_\Gamma)$. The *untwisted reduced Outer Space* is its projectivization $P\Sigma_\Gamma$, the subset where the sum of widths of hyperplanes equals 1, which is preserved under the action. There is a straightforward equivariant embedding of $\mathcal{X}_\Gamma$ into $P\Sigma_\Gamma$, with all hyperplanes given the same width, and this embedding extends to $K_\Gamma \hookrightarrow P\Sigma_\Gamma$.

Given a fixed element σ of $\mathcal{X}_\Gamma$, the subset Σ_σ of Σ_Γ made of the rectilinear Γ-complexes for which forgetting the metric structure (and quotienting appropriately) gives either σ or an element $\sigma' < \sigma$, is identified to a subset of $(\mathbb{R}_+)^n$ containing $(\mathbb{R}_+^*)^n$, where n is the number of hyperplanes in σ. This is done by ordering the hyperplanes in a cube complex representing σ and identifying σ' with the point whose coordinates are the widths of these hyperplanes in order (and 0 if the hyperplane is collapsed). Moreover, Σ'_σ then embeds in Σ_σ in a nice way, which allows for gluing all the thus metrized $(\Sigma_\sigma)_{\sigma\in\mathcal{X}_\Gamma}$ into a *metric structure for* Σ_Γ.

Remark 7. This metric decomposition of Σ_Γ splits $P\Sigma_\Gamma$ into *faces* which are all standard open simplices. However, it is not exactly a simplicial complex, as some faces are missing, namely those that should result of the collapse of a hyperplane which is not collapsible (*e.g.* because it is dual to a loop). For example, the origin doesn't exist in any of the aforementioned $(\mathbb{R}_+)^n$ as it would correspond to a trivial metric complex.

The metric induced on K_Γ under the embedding is simplicial (euclidean on faces), and K_Γ is actually a subcomplex of the barycentric subdivision of $P\Sigma_\Gamma$ (the complex spanned by the barycenters of the faces which are actually present and not missing).

As K_Γ is connected, intersects every subset $\Sigma_\sigma \subseteq \Sigma_\Gamma$, and those cover the whole of Σ_Γ and are each path-connected, Σ_Γ is a path-connected metric space.

Lemma 21. *The (equivariant) map $\Sigma_\Gamma \to P\Sigma_\Gamma$ which scales the metric by a constant factor is a (equivariant) deformation retraction. There also exists a (equivariant) deformation retraction from $P\Sigma_\Gamma$ onto K_Γ. From this follows easily that Σ_Γ and K_Γ are homotopically equivalent.*

The proof of this result, of simplicial essence, will be omitted.

We now turn to the proof of a much harder result, namely the contractibility of the space we just built, which makes up most of [2].

5.2 Contractibility : outline of the proof

Definition 19. Given a salvetti $\sigma \in K_\Gamma$, its *star*, $st(\sigma)$, is the subcomplex spanned by σ and all vertices σ' linked to σ by an edge in K_Γ, that is, that verify $\sigma < \sigma'$, *i.e.* which can collapse to σ. It is contractible via straight-line homotopies to σ.

The central idea of the proof is to add inductively more and more stars of salvettis, which relies on a well-ordering of them, while ensuring at each time that the intersection between the new star and the previous simplicial complex is contractible.

Definition 20. Let $\mathcal{G} = (g_1, g_2, \dots)$ be the family of conjugacy classes of A_Γ with a fixed arbitrary order.

For any salvetti $\sigma = [\mathbb{S}_\Gamma, c]$, given $g \in \mathcal{G}$, the *length of g with respect to σ*, $l_\sigma(g)$ is the minimum length of an edge path in $\mathbb{S}_\Gamma$ which induces the conjugacy class $c^{-1}(g)$ in A_Γ (such an edge path needn't be based as we are considering everything up to conjugacy).

Let $\mathcal{G}_0 = \{g \in \mathcal{G} \mid l_{[\mathbb{S}_\Gamma, id]}(g) \leq 2\}$ and define the *reduced norm* of σ as

$$|| \, \sigma \, ||_0 = \sum_{g \in \mathcal{G}_0} l_\sigma(g) \in \mathbb{Z}$$

and the *norm* of σ as

$$|| \, \sigma \, || = (|| \, \sigma \, ||_0 \, , l_\sigma(g_1), l_\sigma(g_2), \dots) \in \mathbb{Z}^{\mathbb{N}}$$

We endow the set of salvettis with the lexicographic ordering $<$ of their norms (the *weak order*) and with the ordering of their reduced norms (the *strong order*). A strong inequality implies the same weak inequality.

Lemma 22 (Lemma 6.2 in [2])**.** *$\sigma_{\min} := [\mathbb{S}_\Gamma, id]$ is the unique minimal salvetti for the strong order.*

The proof, here omitted, uses the fact that $\mathcal{G}_0$ contains all the classes of length 1 and 2.

Corollary 23 (Corollary 6.3 in [2])**.** *The norm $|| \cdot ||$ is injective on salvettis*

Proof. Given two salvettis $\sigma = [\mathbb{S}_\Gamma, c]$ and σ' with same norm, by the previous lemma, $c^{-1} \cdot \sigma = [\mathbb{S}_\Gamma, id]$ is the only minimum for $|| \cdot ||_0$. However, for every $g \in \mathcal{G}$, $l_\sigma(g) = l_{\sigma'}(g)$ hence $\left|\left| \, c^{-1} \cdot \sigma' \, \right|\right|_0 = \left|\left| \, c^{-1} \cdot \sigma \, \right|\right|_0 = || \, \sigma_{\min} \, ||_0$ which requires $\sigma = \sigma'$. □

Definition 21. A *Whitehead move* from a salvetti σ consists in blowing up a single partition and then collapsing a single hyperplane in σ, producing a new salvetti σ'. The Whitehead move is *weakly reductive* (resp. *strongly reductive*) :if $|| \, \sigma' \, || < || \, \sigma \, ||$ (resp. $|| \, \sigma' \, ||_0 < || \, \sigma \, ||_0$).

The following results make up the core of the proof.

Lemma 24 (Corollary 6.20 in [2])**.** *For every $\sigma \neq \sigma_{\min}$, there exists a strongly reductive Whitehead move from σ.*

Corollary 25 (Proposition 6.22 in [2])**.** *Sorting salvettis by increasing norm $|| \cdot ||$ gives a numbering $\sigma_{\min} = \sigma_0, \sigma_1 \dots$ with index set $\mathbb{N}$ (on which we are able to induct).*

Lemma 26. *For every $n \in \mathbb{N}$ the following intersection is a contractible subcomplex of K_Γ :*

$$st(\sigma_n) \cap \bigcup_{k=0}^{n-1} st(\sigma_k)$$

Corollary 27. *K_Γ and Σ_Γ are contractible.*

Among those results, the hardest are Lemmas 24 and 26. We will only outline their proofs. We start however by proving Corollaries 25 and 27.

5.3 Proof of the corollaries

Proof of Corollary 25 assuming Lemma 24. Let σ_1 be a salvetti. The existence, granted by Lemma 24, of a strongly reductive Whitehead move from σ_1 gives a salvetti σ_2 such that $||\,\sigma_2\,||_0 \leq ||\,\sigma_1\,||_0 - 1$. Following inductively during at most $N = ||\,\sigma_1\,||_0 - ||\,\sigma_{\min}\,||_0$ steps, one is guaranteed to reach the minimal salvetti. Hence σ_1 is at most N Whitehead moves away from $\sigma_{\min}$, and the same is true for every σ with $||\,\sigma\,||_0 \leq ||\,\sigma_1\,||_0$.

Moreover, there is only a finite number of possibilities for a Whitehead move from a given salvetti (bounded by the numbers of partitions and hyperplanes over $\mathbb{S}_\Gamma$). Hence there is a finite number of salvettis σ with $||\,\sigma\,||_0 \leq ||\,\sigma_1\,||_0$, *a fortiori* with $||\,\sigma\,|| < ||\,\sigma_1\,||$. In particular, the set of salvettis is well-ordered by $||\ ||$ and every proper initial segment is finite. It is however infinite by Lemma 19, hence must have the exact order type of $\mathbb{N}$, allowing us to number it and perform an induction.

□

Proof of Corollary 27 assuming the other results. Σ_Γ has already been seen as homotopy equivalent to K_Γ in Lemma 21.

For every $n \in \mathbb{N}$, set

$$K_n := \bigcup_{k=0}^{n-1} st(\sigma_k)$$

By Lemma 26, $K_n \cap st(\sigma_n)$ is contractible for every n. By extension of homotopies in a CW-complex, the quotient map $q_1 \colon K_n \to K_n/(K_n \cap st(\sigma_n))$ is a homotopy equivalence.

Likewise, as $st(\sigma_n)$ is contractible, the quotient map $q_2 \colon K_n \cup st(\sigma_n) \to (K_n \cup st(\sigma_n))/st(\sigma_n)$ is a homotopy equivalence as well.

However, for every $n > 0$ (so that K_n is not empty), $K_n/(K_n \cap st(\sigma_n))$ and $(K_n \cup st(\sigma_n))/st(\sigma_n)$ are actually the same space. In this case, the inclusion $\iota_n \colon K_n \to K_n \cup st(\sigma_n)$ verifies : $q_2 \circ \iota = q_1$. Composing by a homotopy inverse for q_2, ι is a homotopy equivalence. Besides one has the rewriting :

$$K_n \cup st(\sigma_n) = \bigcup_{k=0}^{n} st(\sigma_k) = K_{n+1}$$

.

Finally $d \in \mathbb{N}^*$ being fixed, we prove by induction over n that every continuous map $f \colon \mathbb{S}^d \to K_n$ is homotopic to a constant. There is a special case for $n = 0, 1$: K_0 is empty, so the property is vacuously verified, and $K_1 = st(\sigma_0)$ is contractible so the property is verified as well.

Assume the property verified up to rank $n-1 \geq 1$ and consider a map $f \colon \mathbb{S}^d \to K_n$. Its compact image is contained in a finite subcomplex of K_n. f is homotopic to $\iota_{n-1} \circ \iota_{n-1}^{-1} \circ f \colon \mathbb{S}^d \to K_n$, whose image is contained in K_{n-1} with $\iota_{n-1} \colon K_{n-1} \to K_n$ the inclusion, hence the induction hypothesis applies to $\iota_{n-1}^{-1} \circ f$.

As $K_\Gamma = \bigcup K_n$, any representative of $\pi_d(K_\Gamma)$ has (compact) image in one of the K_n hence is nullhomotopic. K_Γ is a connected CW-complex with trivial homotopy groups, hence is contractible by Whitehead's theorem. □

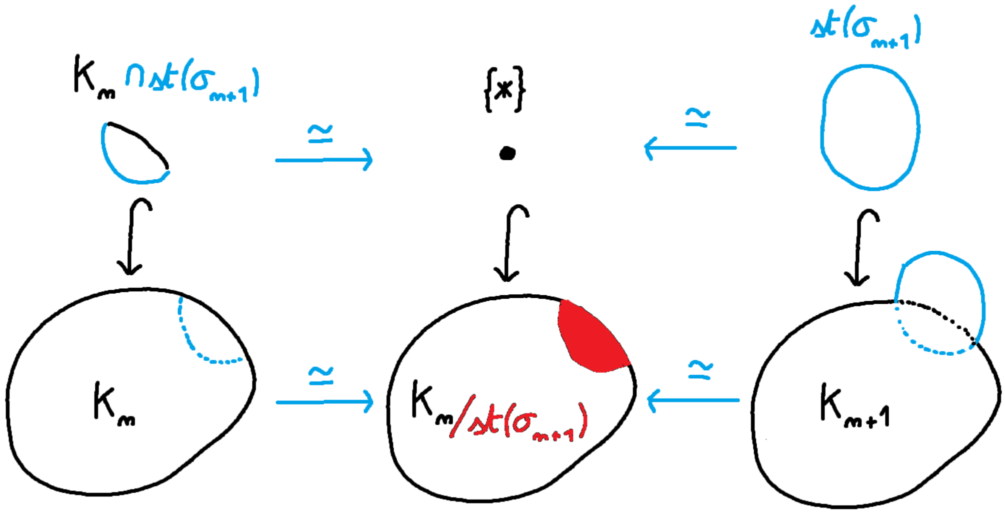

Figure 5: Inductive step of the contraction process
The blue edges on the top row are homotopy equivalences, which extend to homotopy equivalences on the bottom row by the vertical cofibrations.

.

5.4 Elements of proof of the lemmas

To prove Lemma 24, the argument takes the form of a peak reduction. Given a salvetti $\sigma \neq \sigma_{\min}$, there exists a sequence of Whitehead moves starting from σ and ultimately reaching $\sigma_{\min}$, as was proved in Lemma 20 for connectedness. We consider the integer sequence $||\,\sigma\,||_0 = n_0, n_1, \ldots, n_k = ||\,\sigma_{\min}\,||_0$ of reduced norms of the salvettis appearing through the Whitehead moves. As $n_0 > n_k$, two cases arise :

- Either $n_0 > n_1$: the first Whitehead move of the sequence was strongly reductive and we are done
- Or $n_0 \leq n_1$ and the maximum value (the *peak*) of the sequence is attained elsewhere, at index $0 < i < k$. Taking i maximal, we can assume $n_{i-1} \leq n_i > n_{i+1}$.

Now in that second case, granted we can perform a *peak reduction* and substitute n_i for possibly several terms lesser than n_i, after this step :

- Either the maximum was only attained at n_i, and we strictly reduced the value of the maximum
- Or we strictly reduced the number of terms which attained the maximum

Iterating this process, the first possibility can only be performed a finite number of times, and after that the second must also be performed a finite number of times, hence the algorithm always terminates in a situation where there is no peak and the first move was strongly reductive.

Remains the peak reduction argument itself : consider the central salvetti σ' and the two neighbors $\sigma'^{\mathfrak{P}}_v$ and $\sigma'^{\mathfrak{Q}}_w$ (the Whitehead moves must perform a collapse by an edge dual to a vertex, otherwise they would collapse the very partition they blew up and would be trivial.)

The case where $\mathfrak{P}$ and $\mathfrak{Q}$ are compatible is quite straightforward : if $v = w$, the two partitions were only a single Whitehead move away (blowing $\mathfrak{Q}$ up then collapsing $\mathfrak{P}$) and simply deleting the i^{th} term works. Else, $v \neq w$ and essentially going through $\sigma'^{\{\mathfrak{P},\mathfrak{Q}\}}_{v,w}$ instead of σ' works, because blowing up $\mathfrak{Q}$ then collapsing w has to decrease the norm in $\sigma'^{\mathfrak{P}}_v$ in the same way as in σ'.

For the case where $\mathfrak{P}$ and $\mathfrak{Q}$ aren't compatible, one needs a stronger result, namely the *Higgins-Lyndon lemma* (Lemma 6.17 in [2]) which exhibits a new partition $\mathfrak{P}'$ compatible with both $\mathfrak{P}$ and $\mathfrak{Q}$ having good reductiveness properties. Adding in between a Whitehead move blowing $\mathfrak{P}'$ up boils this case down to the previous one and concludes. This lemma involves a lot of combinatorial counting to ensure reductiveness and the remark that P' can be chosen among $\{P \cap Q,\ P \cap Q^*,\ P^* \cap Q,\ P^* \cap Q^*\}$.

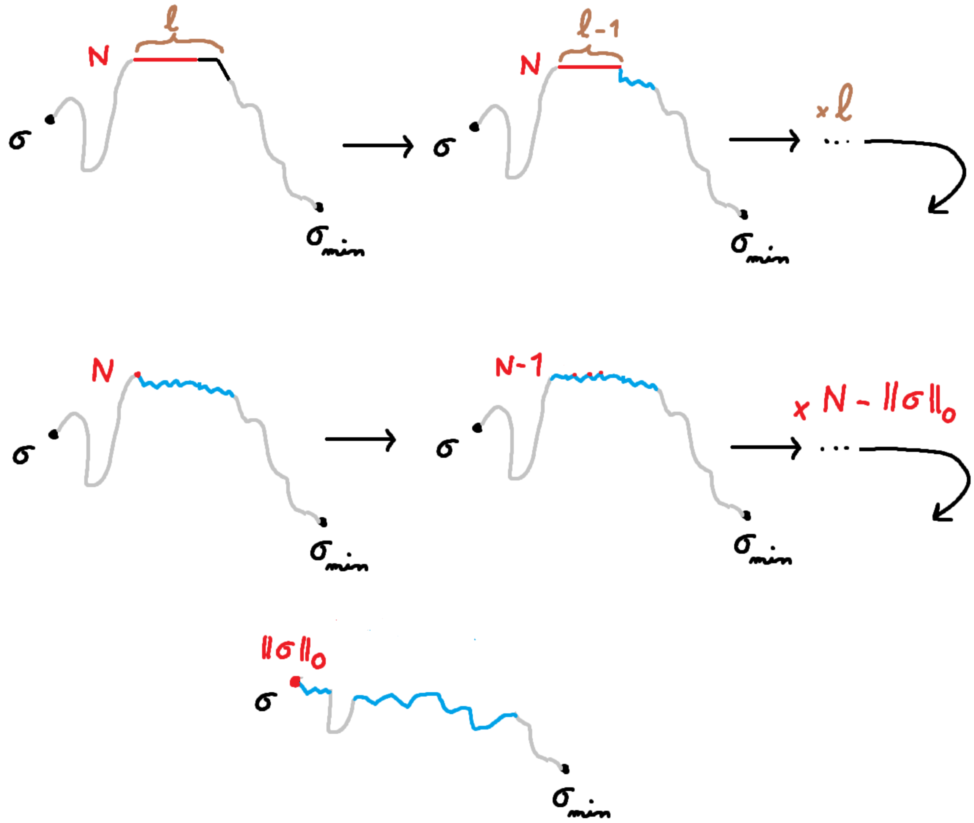

Figure 6: Peak reduction of the path from σ to σ_{min}.
The number of times the maximum value is attained is decreased, then the value itself is decreased and the process is repeated until the first move is reductive.

The proof of Lemma 26 relies heavily on a simplicial lemma due to Quillen ([9]) :

Lemma 28. *Given a poset $A \subseteq \mathcal{X}_\Gamma$ and S the induced subcomplex of K_Γ, any non-decreasing map $f\colon A \to A$ such that $f(x) \leq x$ for every x induces a deformation retraction of S onto $f(S)$.*

Its proof works essentially by completing the homotopy dimension by dimension, the initialization being given by the hypothesis $f(x) \leq x$, and for $x_1 \leq \cdots \leq x_n$ a simplex, using the intermediate simplices $f(x_1) \leq \cdots \leq f(x_i) \leq x_i \leq \cdots \leq x_n$.

Given $n \in \mathbb{N}$, the set A of vertices of $S = st(\sigma_n) \cap \bigcup_{k=0}^{n-1} st(\sigma_k)$ corresponds exactly to the blow-ups $(\sigma_n)^\Pi$ which can collapse to a weakly smaller simplex σ_k. These several blow-ups and collapses can be seen as a sequence of Whitehead moves, at least one of them having to be (weakly) reductive. Let Π' be the subset of Π containing only the partitions which can produce a reductive move. The map

$f\colon (\sigma_n)^\Pi \mapsto (\sigma_n)^{\Pi'}$ satisfies the hypotheses of Lemma 28, hence up to a homotopy equivalence, we can assume that every partition is part of a reductive Whitehead move.

Finally, the goal would be to find a partition $\mathfrak{Q}$ with which every Π arising in our setup is compatible. If such a partition exists, the maps $f'\colon (\sigma_n)^\Pi \mapsto (\sigma_n)^{\Pi\cup\{\mathfrak{Q}\}}$ and $f''\colon (\sigma_n)^{\Pi\cup\{\mathfrak{Q}\}} \to (\sigma_n)^{\mathfrak{Q}}$ both verify the hypotheses of Quillen's lemma (with reversed poset order for f'), and the whole of S is homotopy equivalent to the point $(\sigma_n)^{\mathfrak{Q}}$, hence its contractibility.

If such a $\mathfrak{Q}$ doesn't exist, the proof uses a variant of Higgins-Lyndon lemma to replace the partitions $\mathfrak{P}$ which aren't compatible with a candidate $\mathfrak{Q}$ by partitions $\mathfrak{P}'$ which are, and these replacements can be made in a non-decreasing way like above in order to preserve the homotopy equivalence class.

6 Towards general Outer Space

The final goal of the whole construction, reached in [3] is to build a contractible space $\mathcal{O}_\Gamma$ with a proper action of $Out(A_\Gamma)$, not only the untwisted subgroup $U(A_\Gamma)$. We already raised the problem twists create in terms of geometric realization : they usually cannot preserve cubes. We will now outline the process and constructions made to overcome that obstruction.

6.1 Skewing metrics

If it is not cubical, a twist is still a homotopy equivalence from $\mathbb{S}_\Gamma$ to itself. The image of a cube, or metrically an orthotope, is not an orthotope itself but a parallelotope spanned (linearly) by an edge and a face, the face belonging to the cubical structure of $\mathbb{S}_\Gamma$ and the edge somewhat diagonal to that structure.

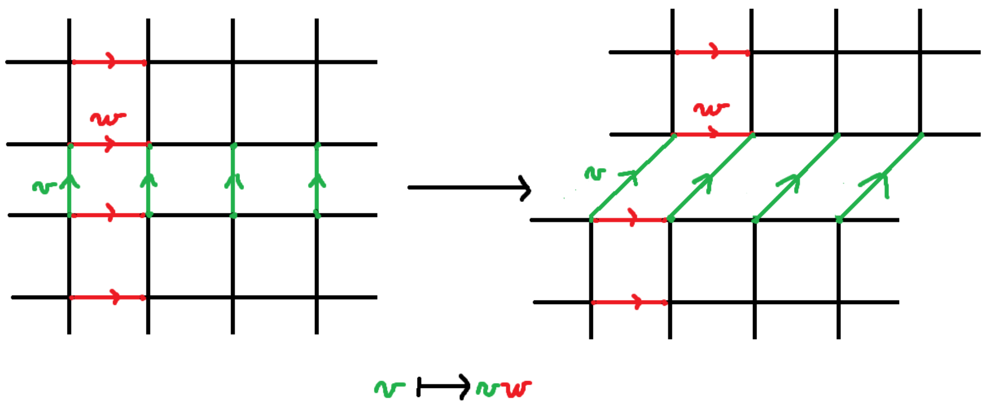

Figure 7: Homotopy equivalence in a part of a complex realizing a twist. The map isn't cubical but at least preserves vertices of the structure.

This enlightens the need to allow for new kinds of metrics on $\mathbb{S}_\Gamma$, and generally on Γ-complexes, called *skewed metrics*, where combinatorial cubes will be represented by parallelotopes. However those metrics still have to respect some allowability conditions, notably so that edges that could be exchanged by a cubical isomorphism still can be exchanged by a cubical isometry.

Definition 22. A parallelotope metric on a cube of a blow-up $\mathbb{S}^\Pi$ is *allowable* :if for any labels l_1 and l_2 which aren't related by $\lhd_t$, the face F_1 spanned by l_1 and all edges l with $l_1 \lhd_t l$ and the face F_2 spanned by l_2 and all edges l' with $l_2 \lhd_t l'$ form a dihedral right-angle[7].

A metric on the entire $\mathbb{S}^\Pi$ is *allowable* :if it restricts to an allowable metric on each maximal cube and besides, for $(\mathfrak{P}, b)$ a based partition with b is twist-dominant, for any label l square-compatible with $\mathfrak{P}$, the angle between $\mathfrak{P}$ and l is the same as the angle between b and l.

A *skewed* Γ*-complex* is a Γ-complex together with a (length space) metric which is allowable in any (hence every by Lemma 13) blow-up structure.

Such a metric can be entirely characterized by hyperplane widths, and the angle between twist-related edges ([3], Lemma 5.2). An issue arises in proving that the metric is locally CAT(0), as convexity of hyperplanes needn't be preserved. However it is true that any skewed Γ-complex is a locally CAT(0) space ([3], Corollary 5.7). The proof uses a *straightening process* where all angles are successively made right (from maximally twist-dominant edges to twist-minimal ones), effectively homotoping the skewed Γ-complex into a regular one, and proving that each step preserves the local CAT(0) nature.

Moreover, that process of straightening can be defined in a continuous way, so that applying it to all skewed Γ-complexes at once defines a deformation retraction of the space $\mathcal{T}_\Gamma$ of skewed Γ-complexes with untwisted marking modulo equivalence onto Σ_Γ. This new space $\mathcal{T}_\Gamma$, defined in a way very similar to Σ_Γ, is the last intermediate step before the full Outer-Space. This deformation retraction onto Σ_Γ guarantees its contractibility, but we cannot yet allow twisted markings.

6.2 Forgetting cubical structure

The last step, which enables finally twisting the markings, is to forget all combinatorial structure on points of $\mathcal{T}_\Gamma$, to keep only the underlying metric space. Of course, it requires the task of recovering the properties of Γ-complexes and blow-ups from this weakened version, which makes up most of [3].

Definition 23 (Full Outer Space)**.** An *Outer point*[8] is the data of a locally CAT(0) metric space (Y, d) which is isometric to some (unspecified) skewed Γ-complex, and a homotopy equivalence $m \colon Y \to \mathbb{S}_\Gamma$.

The *(full) Outer Space* for A_Γ, $\mathcal{O}_\Gamma$ is the set of all Outer points quotiented by the relation : $(Y, d, m) \sim (Y', d', m')$:when there exists an isometry $i \colon (Y, d) \to (Y', d')$ such that $m' \circ i$ is homotopic to m.

This set is endowed by an action of $Out(A_\Gamma)$, seen as homotopy equivalences of $\mathbb{S}_\Gamma$ modulo homotopy, by : $h \cdot (Y, d, m) := (Y, d, h \circ m)$, and with its equivariant Gromov-Hausdorff topology (as a set whose points are compact spaces)[9].

Keeping only the metric and marking from a marked skewed Γ-complex gives a natural map $\Theta \colon \mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$, equivariant with respect to $U(A_\Gamma)$, which can be proved continuous.

Note that this map Θ isn't *a priori* onto because the structure of $\mathcal{O}_\Gamma$ allows for more markings, as intended. However Proposition 7.3 in [3] states that it is actually surjective. That is, every Outer point is in fact equivalent to an Outer point whose marking is untwisted (for some skewed Γ-complex structure). This relies on, given the first Outer point, finding a new Γ-complex with a skewed metric that gives rise to the same metric space, but so that the identity of this

[7] That is : their orthogonal projections along their intersection are orthogonal.

[8] Non-standard terminology due to the author

[9] This metrizable topology, introduced by Gromov without equivariance, was defined in [11], Section 6

metric space amounts to a change of marking of exactly a twist. For a certain kind of twists (where the twisted element is twist-minimal), this can be achieved simply with the same Γ-complex and a different metric and marking (hence the reason for introducing skewed metrics). However for other twists, it requires entirely changing the Γ-complex.

One can then be interested in preimages of points by the map Θ. Corollary 7.18 in [3] identifies them with affine subspaces in some parameter space, and Theorem 7.22 in *loc. cit.* grants that those preimages are contractible.

The final result is that Θ is actually a fibration. Granted this, the associated long exact homotopy sequence has two trivial terms out of three, because of the contractibility of both $\mathcal{T}_\Gamma$ and the fibers. Hence every term is trivial and $\mathcal{O}$ has trivial homotopy[10].

All along this process, the labelling and even the cubical structure are forgotten and can only be partially recovered via results alike Lemma 13, using tools like the CAT(0) study of axes of elements of A_Γ and the concept of *branch loci*, which happen to determine some of the hyperplane widths back from only the metric structure.

References

[1] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. 84 (1986), no. 1, 91-119.

[2] R. Charney, N. Stambaugh, K. Vogtmann, *Outer space for untwisted automorphisms of right-angled Artin groups*, Geom. Topol. 21 (2017), no. 2, 1131-1178.

[3] C. Bregman, R. Charney, K. Vogtmann, *Outer space for RAAGs* (2020), arXiv:2007.09725.

[4] M. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Springer-Verlag, 1999.

[5] J. Nielsen, *Die Isomorphismengruppe der freien Gruppen*, Math. Ann. 91 (1924), 169–209.

[6] H. Servatius *Automorphisms of graph groups*, J. of Algebra 126 (1989), no. 1, 34-60.

[7] M. Laurence, *A generating set for the automorphism group of a graph group*, J. London Math. Soc. 52 (1995), no. 2, 318-334.

[8] M. Day, *Peak reduction and finite presentations for automorphism groups of right-angled Artin groups*, Geom. Topol. 13 (2009), 817-855

[9] D. Quillen, *Higher Algebraic K-Theory I*, Lect. Notes Math. 341 (1974), 85–147.

[10] A. Hatcher and K. Vogtmann, *Cerf theory for graphs*, J. London Math. Soc. 58 (1998), no. 3, 633–655.

[11] K. Fukaya, *Hausdorff convergence of Riemannian manifolds and its applications*, Adv. Stud. Pure Math. 18 (1990), 143–238.

[10] That is only a weak version of contractibility, as $\mathcal{O}_\Gamma$ has no obvious reason of being a CW-complex. More justification would be needed here than what is provided by the article.