Free groups, RAAGs, and their automorphisms

Schriftliche Habilitationsleistung

vorgelegt von

Dawid Kielak

aus Gduńsk

Juni 2017

Universität Bielefeld

Contents

Introduction

The research presented in this thesis revolves around the structure of free groups, their close relatives the right-angled Artin groups, and the automorphisms of groups in both classes. As such, it positions itself inside the areas of combinatorial and geometric group theory. The former originated in the early 20^{th} century in the work of Dehn, and was further developed by Higman, Magnus, Nielsen, and Whitehead, to name just a few. The theory evolved with time and in its later stages was heavily reshaped by the work of Stallings. In the 1980s Mikhail Gromov revolutionised the subject, and geometric group theory was born. This latter branch of mathematics is concerned with studying groups via their actions on metric spaces; this point of view proved to be extremely fruitful.

As in the case of every theory, there are objects that geometric group theory is particularly focused on. Initially, these were the classes of hyperbolic and CAT(0) groups, as well as mapping class groups of surfaces, which connect geometric group theory to Thurston-style geometric topology. More recently, we observe a shift of interest from mapping class groups to the groups of automorphisms of free groups (most notably $\mathrm{Out}(F_n)$, the group of outer such automorphisms); it is worth pointing out that automorphism groups of free groups as well as mapping class groups were an active area of research in the earlier days of combinatorial group theory.

The solution of the Virtually Fibred Conjecture of Thurston due to Agol [**Ago**] (based on the work of Wise) brought right-angled Artin groups (RAAGs) into focus. These groups are given by presentations in which the only relators are commutators of some generators. Hence RAAGs interpolate between free and free-abelian groups, and so their outer automorphism groups form a natural class of groups interpolating between $\mathrm{Out}(F_n)$ and $\mathrm{GL}_n(\mathbb{Z})$, the latter being an important group of more classical interest.

Free groups

Free groups form a cornerstone of group theory, and are objects of central importance in combinatorial and geometric group theory. The structure of a free group is extremely simple (indeed, as simple as possible for a group), and hence most questions one can ask about a group have immediate answers for free groups. This is however not the case for all questions – there are non-trivial results about the structure of free groups. One such result (of significance for the current discussion) is the fact (essentially due to Magnus) that free groups are *orderable*, that is one can endow a free group F with a total order invariant under left multiplication. In fact, free groups are *biorderable*, that is the total order can be taken to be invariant under left and right multiplication simultaneously.

The orderability of F_n, the free group on n generators, features crucially in the proof of the Hanna Neumann Conjecture due to Mineyev [**Min**].

When a group G is orderable, the set of all possible orderings can be topologised in the following way: to each ordering $\leqslant$ we associate the *positive cone*

$$P_{\leqslant} = \{x \in G \mid x > 1\}$$

This allows us to view the set of orderings as a subset of the power set of F_n given the product topology, thus becoming the *space of orderings* (note that the topology of this space is referred to as the Sikora topology). For F_n this space has been shown to be homeomorphic to a Cantor set – this follows from the work of McCleary [**McC**], but appears in this form for the first time in an article of Navas [**Nav**]. Finding the homeomorphism type of the corresponding space of biorderings of F_n is still an open problem.

Observe that the positive cone $P_{\leqslant}$ is a subsemigroup of G satisfying

$$G = P_{\leqslant} \sqcup {P_{\leqslant}}^{-1} \sqcup \{1\}$$

where ${P_{\leqslant}}^{-1} = \{x^{-1} \mid x \in P_{\leqslant}\}$, and $\sqcup$ denotes the disjoint union. In particular, $G = P_{\leqslant}{P_{\leqslant}}^{-1}$, that is every element in G can be written as a *fraction* of two elements in $P_{\leqslant}$ (we say that G is the *group of fractions* of its subsemigroup $P_{\leqslant}$).

One important feature of the Sikora topology is that any positive cone in G finitely generated as a semigroup is isolated in the space of orderings. Thus, observing that the space of left orderings of F_n has no isolated points (as it is homeomorphic to the Cantor set), we deduce that $P_{\leqslant}$ is not finitely generated as a semigroup for any left-invariant ordering $\leqslant$ on F_n.

One can investigate the general structure of subsemigroups $P \subseteq F_n$ of which F_n is a fraction group. Navas conjectured that such a P cannot at the same time be finitely generated (as a semigroup) and satisfy $P \cap P^{-1} = \emptyset$. Chapter I contains the proof of this conjecture; in fact much more is proved.

Theorem I.2.3. *Let P be a finitely generated subsemigroup of a finitely generated group F with infinitely many ends. If $PP^{-1} = F$ then $P = F$.*

Note that F_n is a particular example of a group with infinitely many ends.

The proof uses the topology of a group with infinitely many ends directly, but let us mention here a celebrated result of Stallings [**Sta1, Sta2**]: a finitely generated group has at least 2 ends if and only if it *splits over a finite subgroup*, that is acts on a tree without a global fixed point and with a single edge-orbit with a finite stabiliser. In the language of Bass–Serre theory, such a group is the fundamental group of a non-trivial graph of groups with exactly one edge and a finite edge group. Moreover, the group in question has infinitely many ends if and only if it is additionally not virtually cyclic. We will return to Stallings' theorem when discussing automorphisms of free groups.

Coming back to Theorem I.2.3, as an application we obtain a new proof of the fact that the space of orderings of F_n is homeomorphic to the Cantor set. It is worth noting that the proof offered in Chapter I is the first geometric one.

We also deduce that the orderings of finitely generated groups with infinitely many ends do not have finitely generated positive cones. This was already known for free products of orderable groups by the work of Rivas [**Riv**].

Free groups, interesting as they intrinsically are, perform also an extrinsic function in group theory. It is often useful to establish whether a given group G contains a free subgroup; this is especially important when studying amenability, as being amenable and possessing F_2 as a subgroup are mutually exclusive properties.

The von Neumann–Day Conjecture claimed that in fact possessing F_2 was the only obstruction to amenability. The conjecture has been proven wrong by Olshanskiy in 1980, but still groups which are not amenable and do not contain F_2 are considered to be somewhat exotic.

Amenability is an extremely interesting property of a group. One of the definitions of amenability of a (discrete) group G is the statement that G admits a left-multiplication invariant mean, but amenability has many other equivalent definitions which are more natural in specific contexts.

Tamari [**Tam**] showed that if G is amenable, then $\mathbb{K}G$, the group algebra of G with coefficients in a field $\mathbb{K}$, satisfies the *Ore condition*, that is if we take $p, q \in \mathbb{K}G$ where q is not a zero-divisor, then there exist $r, s \in \mathbb{K}G$ such that r is not a zero-divisor and

$$pr = qs$$

Is is easy to see that any group algebra of F_2 does not satisfy the Ore condition. Since the condition passes to subgroups, this implies that any group containing F_2 does not have a group algebra satisfying the Ore condition. In view of the above, the question of establishing the Ore condition becomes interesting for all counterexamples to the von Neumann–Day Conjecture. Prior to the appearance of the work contained in Chapter II, the question was not resolved even for a single such group.

A folklore conjecture (sometimes attributed to Guba) said that the converse to Tamari's result holds true, that is a group G with $\mathbb{K}G$ satisfying the Ore condition is amenable. If one assumes that $\mathbb{K}G$ has no non-trivial zero-divisors, this conjecture has been confirmed.

Theorem II.2.2. *Let G be a group, and let $\mathbb{K}$ be a field such that $\mathbb{K}G$ has no non-trivial zero divisors. Then G is amenable if and only if $\mathbb{K}G$ satisfies the Ore condition.*

Note that the assumption on $\mathbb{K}G$ not containing non-trivial zero-divisors is related to the notorious Zero-divisor Conjecture of Kaplansky [**Kap**], which states that $\mathbb{K}G$ contains non-trivial zero-divisors if and only if G contains torsion (independently of $\mathbb{K}$). Kaplansky's conjecture has been verified for many classes of groups, including elementary amenable and orderable groups.

Automorphisms of free groups

We start by investigating a single automorphism $g\colon F_n \to F_n$. To be able to use geometric methods, we need to associate a space to g. A natural such space is formed by the mapping torus T_g: we start with a classifying space Γ for F_n (this will be a graph), take a direct product with the unit interval $[0,1]$, and identify $\Gamma \times \{0\}$ with $\Gamma \times \{1\}$ using the homotopy equivalence associated to g (the classifying map of g). It is immediate that the fundamental group of T_g is isomorphic to $G = F_n *_g$, the HNN extension of F_n associated to g, or equivalently to the semi-direct product $F_n \rtimes_g \mathbb{Z}$.

Now, there are various properties of T_g that one can investigate, and we will be interested in L^2-invariants. The most basic L^2-invariants are the L^2-Betti numbers, that is the (von Neumann) dimensions of the L^2-homology spaces. But the L^2-Betti numbers are all trivial for any mapping torus, and so in particular for T_g. Thus, we need a secondary invariant.

To introduce the secondary invariant we will first look at an analogy: Consider a square matrix M over a field. The primary invariant, which corresponds to an L^2-Betti number, is the dimension of the kernel of M. When this dimension is 0, then M is an invertible linear map. Therefore M has a non-trivial determinant, which is the secondary invariant. Similarly for T_g: since we know that T_g is L^2-acyclic (i.e. its L^2-Betti numbers vanish), we need a version of the determinant. This is provided by the *universal L^2-torsion* $\rho_u^{(2)}$ of Friedl–Lück, introduced in a series of papers [**FL1, FL2, FL3**].

We investigate the universal L^2-torsion for a mapping torus T_g not only of an automorphism of F_n, but also of an injective endomorphism. On the level of the group, we look at $\rho_u^{(2)}$ not only for F_n-by-$\mathbb{Z}$ groups, but also for the (more general) descending HNN extensions of finitely generated free groups.

The universal L^2-torsion induces a semi-norm on the first real cohomology of the group $G = F_n{*}_g$ – the way in which this happens is somewhat indirect and involved, and explained in detail in Sections 2.6, 2.7 and 2.9 of Chapter III. The semi-norm induced by the universal L^2-torsion coincides with the Thurston norm when G is a 3-manifold group, and therefore, by analogy, the semi-norm induced by $\rho_u^{(2)}$ is also referred to as the *Thurston norm*, and denoted by $\| \,.\, \|_T$. Note that the Thurston norm is a well-established tool in the realm of 3-manifolds. It is a semi-norm on the first real cohomology of the manifold, intimately related to the question of whether and how the manifold fibres over the circle. Its definition is of a strictly topological character, and it is very interesting to see that it can be also defined using the L^2-invariants, which are analytic and algebraic in flavour.

For descending HNN extensions of F_2, we see in Chapter III that this Thurston norm is an upper bound for the Alexander semi-norm δ_0 defined by McMullen, as well as for the higher Alexander semi-norms δ_n defined by Harvey.

Theorem III.4.7. *Let $G = F_2{*}_g$ be a descending HNN extension of F_2 such that the first Betti number satisfies $b_1(G) \geqslant 2$. Then the Thurston and higher Alexander semi-norms satisfy for all $n \geqslant 0$ and $\varphi \in H^1(G;\mathbb{R})$ the inequality*

$$\delta_n(\varphi) \leqslant \|\varphi\|_T$$

The same inequalities are known to hold for 3-manifold groups, and are of particular importance for knots (and the associated complementary 3-manifolds).

We extend this result to higher rank free groups when the extension is taken over a *unipotent polynomially growing* (UPG) automorphism.

Corollary III.6.6. *Let $G = F_n \rtimes_g \mathbb{Z}$ with $n \geqslant 2$ and g a UPG automorphism. Let $\varphi \in H^1(G;\mathbb{R})$. Then for all $k \geqslant 0$ we have*

$$\delta_k(\varphi) = \|\varphi\|_T$$

When studying an infinite cyclic extension G of a group by an endomorphism, it is often interesting to find out in how many ways can G be decomposed as such an extension. Some information can be gained by looking at the *Bieri–Neumann–Strebel (BNS) invariants.* The (first) BNS invariant is a subset $\Sigma \subseteq H^1(G;\mathbb{R}) \smallsetminus \{0\}$, closed under positive homothety, such that $\varphi\colon G \to \mathbb{R}$ lies in Σ if and only if G admits a Cayley graph (with respect to a finite generating set) in which the vertices mapped by φ to $[0,\infty)$ induce a connected subgraph.

When $G = F_n{*}_g$ is a descending HNN extension, then Σ becomes very meaningful – if φ and $-\varphi$ belong to Σ, then $\ker \varphi$ is a finitely generated free group (by the work of Geoghegan–Mihalik–Sapir–Wise [**GMSW**]), and so we can write G as a semi-direct product of $\ker\varphi$ by $\mathbb{Z}$.

In the case of two-generator one-relator groups G with $b_1(G) = 2$, Friedl–Tillmann [**FT**] established a close connection between $\rho_u^{(2)}$ and the Bieri–Neumann–Strebel invariant $\Sigma(G)$. Theorem III.5.13 and Corollary III.6.4 exhibit a similar connection in the realm of descending HNN extensions of free groups.

So far we have focused on a single automorphism g of F_n. Another way of studying automorphisms of F_n is to study the group they form, $\mathrm{Aut}(F_n)$. Since F_n has no centre (we are implicitly assuming that $n \geqslant 2$), it embeds into $\mathrm{Aut}(F_n)$ as the group of inner automorphisms. This way we have F_n contained in $\mathrm{Aut}(F_n)$ as a normal subgroup, and taking the quotient gives us $\mathrm{Out}(F_n)$, the group of *outer* automorphisms of F_n.

We will start by discussing finite subgroups of $\mathrm{Out}(F_n)$ – they are classified by the *Nielsen realisation* theorem for $\mathrm{Out}(F_n)$, due to, independently, Culler [**Cul**], Khramtsov [**Khr1**], and Zimmermann [**Zim1**] (compare the paper of Hensel–Osajda–Przytycki [**HOP**] for a different approach). The theorem says that any finite subgroup of $\mathrm{Out}(F_n)$ can be realised as a group acting on a finite graph with fundamental group F_n. Hence, whenever we are looking at a homomorphism $G \to \mathrm{Out}(F_n)$ for some group G, we can pin down the image of torsion subgroups of G, which in many situations leads to the classification of all possible homomorphisms $G \to \mathrm{Out}(F_n)$; for examples of this technique see the results of Bridson–Vogtman [**BV1, BV2, BV3**] and the author [**Kie1, Kie2**].

In its original form, the *Nielsen realisation problem* asks whether torsion subgroups of the mapping class group of a surface can be realised as groups of homeomorphisms of the surface. A celebrated result of Kerckhoff [**Ker1, Ker2**] answers this in the positive, and even allows for realisations by isometries of a suitable hyperbolic metric.

Coming back to free groups, suppose that we write F_n as a free product

$$F_n = A_1 * \cdots * A_k * B$$

and suppose that we have a finite subgroup $H < \mathrm{Out}(F_n)$ such that for every $h \in H$ and every j there exists j' with $h(A_j)$ conjugate to $A_{j'}$. Vogtmann asked whether in such a situation one can find a graph Γ realising the action of H so that for each j the graph Γ contains a subgraph Γ_j carrying the free factor A_j of F_n in such a way that h sends Γ_j to $\Gamma_{j'}$ and induces the given automorphism $A_j \to A_{j'}$ (up to conjugation) on the fundamental groups of Γ_j and $\Gamma_{j'}$. Chapter IV contains a positive answer to this question; in fact it contains a much more general result (Theorem IV.7.5). The statement of the theorem is somewhat technical, so let us unravel it here: we start with a free product

$$A = A_1 * \cdots * A_n * B$$

where the groups A_i are finitely generated, and B is a finitely generated (possibly trivial) free group. We are given a finite group H of outer automorphism of A such that each $h \in H$ takes each A_i to some conjugate of A_j (that is, H *preserves the free product decomposition*). Additionally, for each i, we are given a complete non-positively curved space X_i with fundamental group A_i, such that the subgroup of H fixing A_i up to conjugation acts on X_i realising its (outer) action on A_i. Then we can construct a complete non-positively curved space X with fundamental group A, with an action of H realising the outer action of H on A, and such that X contains the spaces X_i in a way equivariant for the appropriate subgroup of H.

The proof of the above result involves going back to the proof of Stallings' theorem on groups with at least two ends, and investigating what can be said when the group in question is a finite extension of a free product, where the finite group acts in a way preserving the free product decomposition. Under such an assumption the following was obtained.

Theorem IV.2.9. *Let $\varphi\colon H \to \mathrm{Out}(A)$ be a monomorphism with a finite domain. Let $A = A_1 * \cdots * A_n * B$ be a free product decomposition with each A_i and B finitely generated, and suppose that it is preserved by H. Let $\overline{A}$ be the preimage of $H = \mathrm{im}\ \varphi$ in $\mathrm{Aut}(A)$. Then $\overline{A}$ splits over a finite group in such a way that each A_i fixes a vertex in the associated action on a tree.*

Stallings' theorem was an important tool used by Karrass–Pietrowski–Solitar [**KPS**] to prove that every finitely generated virtually free group is the fundamental group of a finite graph of finite groups. Using the relative version of Stallings' result above, we prove

Theorem IV.4.1. *Let*

$$\varphi\colon H \to \mathrm{Out}(A)$$

be a monomorphism with a finite domain, and let

$$A = A_1 * \cdots * A_n * B$$

be a decomposition preserved by H*, with each* A_i *finitely generated and non-trivial, and* B *a (possibly trivial) finitely generated free group. Let* $A_1, \dots, A_m$ *be the minimal factors. Then the associated extension* $\overline{A}$ *of* A *by* H *is isomorphic to the fundamental group of a finite graph of groups with finite edge groups, with* m *distinguished vertices* $v_1, \dots, v_m$*, such that the vertex group associated to* v_i *is a conjugate of the extension* $\overline{A_i}$ *of* A_i *by* $\mathrm{Stab}_H(i)$*, and vertex groups associated to other vertices are finite.*

There are two corollaries of note:

Corollary IV.5.1. *Let* $H \leqslant \mathrm{Out}(A, \{A_1, \dots, A_n\})$ *be a finite subgroup, and suppose that the factors* A_i *are finitely generated. Then* H *fixes a point in the free-splitting graph* $\mathcal{FS}(A, \{A_1, \dots, A_n\})$*.*

Here, $\mathcal{FS}(A, \{A_1, \dots, A_n\})$ denotes the free-splitting graph of Handel–Mosher [**HM**] – this is an analogue of the usual free splitting graph of $\mathrm{Out}(F_n)$, adapted to the setting of more general free products. This free-splitting graph comes with an action of $\mathrm{Out}(A, \{A_1, \dots, A_n\})$, the group of outer automorphisms of A preserving the free product decomposition.

Corollary IV.6.1. *Let* A *be a finitely generated group, and let* $H \leqslant \mathrm{Out}(A)$ *be a finite subgroup. Then* H *fixes a vertex in* $P\mathcal{O}$*.*

Here, $P\mathcal{O}$ denotes the Outer Space for free products of Guirardel–Levitt [**GL**] – again this is analogue of the Culler–Vogtmann Outer Space of $\mathrm{Out}(F_n)$, again adapted to the setting of more general free products. This time however we do not specify the free factors A_i – they are the one-ended factors coming from the Grushko decomposition.

Both of the above corollaries extend results known in the setting of $\mathrm{Out}(F_n)$ to the setting of free products, and so allow us to better understand the latter in a way that already proved fruitful in the study of free groups.

Chapter IV also contains a proof of Nielsen realisation for limit groups, that is finitely generated fully residually free groups.

Theorem IV.8.11. *Let* A *be a limit group, and let*

$$A \to \overline{H} \to H$$

be an extension of A *by a finite group* H*. Then there exists a complete locally CAT(*κ*) space* X *realising the extension* $\overline{H}$*, where* $\kappa = -1$ *when* A *is hyperbolic, and* $\kappa = 0$ *otherwise.*

This theorem is obtained by combining the classical Nielsen realisation theorems (for free, free-abelian and surface groups – see Theorems 8.1 to 8.3 in Chapter IV) with the existence of an invariant JSJ decomposition shown by Bumagin–Kharlampovich–Myasnikov [**BKM**].

Observe that the curvature bounds for the space X are optimal – it was proved by Alibegović–Bestvina [**AB**] that limit groups are CAT(0), and by Sam Brown [**Bro2**] that a limit group is CAT(-1) if and only if it is hyperbolic.

Limit groups have been intensely studied ever since they appeared as crucial objects in Sela's solution to the Tarski problem on elementary theory of free groups. Limit groups combine the study of free groups and surface groups and are quite general; at the same time they have enough structure to allow for interesting results.

After studying finite subgroups of $\mathrm{Out}(F_n)$, we turn our attention to its finite quotients. Investigating such quotients of mapping class groups and $\mathrm{Out}(F_n)$ has a long history. The first fundamental result here is that groups in both classes are residually finite – this is due to Grossman [**Gro2**].

Once we know that the groups admit many finite quotients, we can start asking questions about the structure or size of such quotients. This is of course equivalent to studying normal subgroups of finite index in mapping class groups and $\mathrm{Out}(F_n)$.

When $n \geqslant 3$, the groups $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$ have unique subgroups of index 2, denoted respectively by $\mathrm{SOut}(F_n)$ and $\mathrm{SAut}(F_n)$. Both of these subgroups are perfect, and so the abelian quotients of $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$ are well understood – they are precisely the trivial group and $\mathbb{Z}/2\mathbb{Z}$. The situation for mapping class groups is very similar.

The simplest way of obtaining a non-abelian quotient of $\mathrm{Out}(F_n)$ or $\mathrm{Aut}(F_n)$ comes from observing that $\mathrm{Out}(F_n)$ acts on the abelianisation of F_n, that is $\mathbb{Z}^n$. In this way we obtain (surjective) maps

$$\mathrm{Aut}(F_n) \to \mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$$

The finite quotients of $\mathrm{GL}_n(\mathbb{Z})$ are controlled by the congruence subgroup property and are well understood. In particular, the smallest (in terms of cardinality) such quotient is $\mathrm{PSL}_n(\mathbb{Z}/2\mathbb{Z}) = \mathrm{L}_n(2)$, obtained by reducing $\mathbb{Z}$ modulo 2. According to a conjecture of Mecchia–Zimmermann [**MZ**], the group $\mathrm{L}_n(2)$ is the smallest non-abelian quotient of $\mathrm{Out}(F_n)$.

In [**MZ**] Mecchia and Zimmermann confirmed their conjecture for $n \in \{3, 4\}$. In Chapter V the conjecture is shown for all $n \geqslant 3$. In fact more is proven:

Theorem V.9.1. *Let $n \geqslant 3$. Every non-trivial finite quotient of $\mathrm{SAut}(F_n)$ is either greater in cardinality than $\mathrm{L}_n(2)$, or isomorphic to $\mathrm{L}_n(2)$. Moreover, if the quotient is $\mathrm{L}_n(2)$, then the quotient map is the natural map postcomposed with an automorphism of $\mathrm{L}_n(2)$.*

The natural map $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ is obtained by acting on $\mathrm{H}_1(F_n; \mathbb{Z}/2\mathbb{Z})$.

Zimmermann [**Zim2**] also obtained a partial solution to the corresponding conjecture for mapping class groups, which states that the smallest non-trivial quotient of a mapping class group of a compact connected oriented surface of genus g without punctures is $\mathrm{Sp}_{2g}(2)$. This conjecture has been confirmed by Pierro and the author in [**KP**].

In order to determine the smallest non-trivial quotient of $\mathrm{SAut}(F_n)$, we restrict our attention to the finite simple groups which, by the Classification of Finite Simple Groups (CFSG), fall into one of the following four families:

(1) the cyclic groups of prime order;
(2) the alternating groups A_n, for $n \geqslant 5$;
(3) the finite groups of Lie type; and
(4) the 26 sporadic groups.

For the purpose of the current discussion, the finite groups of Lie type are divided into the following two families:

(3C) the "classical groups": $\mathtt{A}_n$, ${}^2\mathtt{A}_n$, $\mathtt{B}_n$, $\mathtt{C}_n$, $\mathtt{D}_n$ and ${}^2\mathtt{D}_n$, and;
(3E) the "exceptional groups": ${}^2\mathtt{B}_2$, ${}^2\mathtt{G}_2$, ${}^2\mathtt{F}_4$, ${}^3\mathtt{D}_4$, ${}^2\mathtt{E}_6$, $\mathtt{G}_2$, $\mathtt{F}_4$, $\mathtt{E}_6$, $\mathtt{E}_7$ and $\mathtt{E}_8$.

In the context of the alternating groups the following is shown.

Theorem V.3.16. *Let $n \geqslant 3$. Any action of $\mathrm{SAut}(F_n)$ or $\mathrm{SOut}(F_n)$ on a set with fewer than $k(n)$ elements is trivial, where*

$$k(n) = \begin{cases} 7 & & n = 3 \\ 8 & & n = 4 \\ 12 & \text{if} & n = 5 \\ 14 & & n = 6 \end{cases}$$

and $k(n) = \max_{r \leqslant \frac{n}{2}-3} \min\{2^{n-r-p(n)}, \binom{n}{r}\}$ *for* $n \geqslant 7$*, where* $p(n)$ *equals* 0 *when* n *is odd and* 1 *when* n *is even.*

In particular, we conclude that the smallest alternating quotient of $\mathrm{SAut}(F_n)$ has to have rank at least $k(n)$, and so this alternating group is much bigger than $\mathrm{L}_n(2)$.

The bound given above for $n \geqslant 7$ is somewhat mysterious; one can however easily see that $k(n)$ is bounded below by $2^{\frac{n}{2}}$ for large n.

Note that, so far, no such result was available for $\mathrm{SAut}(F_n)$ (one could extract a bound of $2n$ from the work of Bridson–Vogtmann [**BV1**]).

The question of the smallest set on which $\mathrm{SAut}(F_n)$ or $\mathrm{SOut}(F_n)$ can act non-trivially remains open, but the result above does answer the question on the growth of the size of such a set with n – it is exponential. Note that the corresponding question for mapping class groups has been answered by Berrick–Gebhardt–Paris [**BGP**].

Note that $\mathrm{Out}(F_n)$ (and hence also $\mathrm{SAut}(F_n)$) has plenty of alternating quotients – indeed, it was shown by Gilman [**Gil**] that $\mathrm{Out}(F_n)$ is residually alternating.

Following the alternating groups, the sporadic groups are ruled out. It was observed by Bridson–Vogtman [**BV1**] that any quotient of $\mathrm{SAut}(F_n)$ which does not factor through $\mathrm{SL}_n(\mathbb{Z})$ must contain a subgroup isomorphic to

$$(\mathbb{Z}/2\mathbb{Z})^{n-1} \rtimes A_n = 2^{n-1} \rtimes A_n$$

(one can easily see this subgroup inside of $\mathrm{SAut}(F_n)$, as it acts on the n-rose, that is the bouquet of n circles). Thus, for large enough n, sporadic groups are never quotients of $\mathrm{SAut}(F_n)$.

When dealing with the finite groups of Lie type, the strategy differs depending on whether it is the classical or exceptional groups that are being considered. The exceptional groups are handled in a similar fashion to the sporadic groups – this time an alternating subgroup A_{n+1} inside $\mathrm{SAut}(F_n)$ is used; this subgroup rigidifies $\mathrm{SAut}(F_n)$ in a similar way as the subgroup $2^{n-1} \rtimes A_n$ did. The degrees of the largest alternating subgroups of exceptional groups of Lie type are known (and listed for example in [**LS**]); in particular this degree is bounded above by 17 across all such groups.

The most involved part of Chapter V deals with the classical groups. In characteristic 2 an inductive strategy is used, and results in

Theorem V.6.9. *Let* $n \geqslant 3$*. Let* K *be a finite group of Lie type in characteristic* 2 *of twisted rank less than* $n-1$*, and let* $\overline{K}$ *be a reductive algebraic group over an algebraically closed field of characteristic* 2 *of rank less than* $n-1$*. Then any homomorphism* $\mathrm{Aut}(F_n) \to K$ *or* $\mathrm{Aut}(F_n) \to \overline{K}$ *has abelian image, and any homomorphism* $\mathrm{SAut}(F_{n+1}) \to K$ *or* $\mathrm{SAut}(F_{n+1}) \to \overline{K}$ *is trivial.*

In odd characteristic the proof goes through the representation theory of $\mathrm{SAut}(F_n)$. The following is obtained:

Theorem V.7.12. *Let* $n \geqslant 8$*. Every irreducible projective representation of* $\mathrm{SAut}(F_n)$ *of dimension less than* $2n-4$ *over a field of characteristic greater than* 2 *which does not factor through the natural map* $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ *has dimension* $n+1$*.*

Over a field of characteristic greater than $n+1$, every linear representation of $\mathrm{Out}(F_n)$ of dimension less than $\binom{n+1}{2}$ factors through the map

$$\mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$$

mentioned above (see [**Kie1**, 3.13]). Representations of $\mathrm{SAut}(F_n)$ over characteristic other than 2 have also been studied by Varghese [**Var**].

The group $\mathrm{Out}(F_n)$ contains many groups of classical interest as subgroups. In particular, the braid group

$$B_n = \langle \sigma_1, \ldots, \sigma_{n-1} \mid \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}, \sigma_i\sigma_j = \sigma_j\sigma_i \text{ when } j \neq i+1 \rangle$$

acts on $F_n = \langle a_1, \ldots, a_n \rangle$ via the Hurwitz action:

$$\sigma_i.a_j = \begin{cases} a_j & j \notin \{i, i+1\} \\ a_{i+1} & \text{if} \quad j = i \\ {a_{i+1}}^{-1}a_i a_{i+1} & j = i+1 \end{cases}$$

and it is easy to see that this action gives an embedding $B_n < \mathrm{Out}(F_n)$.

In [**Cha**] Charney asked whether braid groups are CAT(0), that is whether they act geometrically (properly discontinuously and cocompactly) on a CAT(0) space. The property of being CAT(0) has far reaching consequences for a group. Algorithmically, such groups have quadratic Dehn functions and hence soluble world problem; geometrically, all free-abelian subgroups are undistorted; algebraically, the centralisers of infinite cyclic subgroups split.

Brady and McCammond showed in [**BM2**] that the n-strand braid groups are CAT(0) if $n = 4$ or 5. However, their proof for $n = 5$ relies heavily on a computer program. They also conjectured that the same statement should hold for arbitrary n [**BM2**, Conjecture 8.4].

In Chapter VI we exploit the close relationship between braid groups, non-crossing partitions of a regular n-gon, and the geometry of spherical buildings (the latter relationship was discovered by Brady and McCammond [**BM2**]). We prove

Theorem VI.4.17. *For every $n \leqslant 6$ the diagonal link in the orthoscheme complex of non-crossing partitions NCP_n is* CAT(1)*.*

The *orthoscheme complex* is a metric simplicial complex associated to a bounded graded poset (and NCP_n is such a poset). On the combinatorial level, it is just the simplicial realisation of the poset. Each maximal simplex, which corresponds to a chain of a fixed length, say m, is given the metric of the standard m-orthoscheme, that is the convex hull in $\mathbb{R}^n$ of the points $v_1, v_1 + v_2, \ldots, v_1 + \cdots + v_m$, where $v_1, \ldots, v_m$ is a standard basis of $\mathbb{R}^n$.

As a consequence we obtain

Corollary VI.5.6. *For every $n \leqslant 6$, the n-strand braid group is* CAT(0)*.*

Brady and McCammond conjectured further that the orthoscheme complex of any bounded graded modular lattice is CAT(0) [**BM2**, Conjecture 6.10]. Chapter VI contains a partial result towards the solution of this problem.

Theorem VI.4.18. *The orthoscheme complex of any bounded graded modular complemented lattice is* CAT(0)*.*

Since the article [**HKS**] was published, [**BM2**, Conjecture 6.10] was proven in full generality by Chalopin–Chepoi–Hirai–Osajda [**CCHO**].

Automorphisms of RAAGs

Consider a finite simplicial graph Γ. We associate to it a right-angled Artin group (RAAG) A_Γ in the following way. We let

$$A_\Gamma = \big\langle V(\Gamma) \mid \{[v, w] : (v, w) \in E(\Gamma)\} \big\rangle$$

where $V(\Gamma)$ and $E(\Gamma)$ are, respectively, the vertex and edge set of Γ. (We will refer to Γ as the *defining graph*.)

It is immediate that when Γ has no edges then A_Γ is a finitely generated free group, and when Γ is complete (equivalently, is a clique), then A_Γ is a finitely generated free-abelian group. Thus, RAAGs form a (rich) family of groups interpolating between free and free-abelian groups. From this point of view, it is natural to also think of their (outer) automorphism groups as interpolating between $\mathrm{Aut}(F_n)$ (or $\mathrm{Out}(F_n)$) and $\mathrm{GL}_n(\mathbb{Z})$.

Within the family formed by groups $\mathrm{Out}(A_\Gamma)$, it is interesting to study rigidity phenomena occurring for maps between various members of the family. Homomorphisms

$$\mathrm{Out}(F_n) \to \mathrm{Out}(F_m)$$

have been investigated by Bogopolski–Puga [**BP**], Khramtsov [**Khr2**], Bridson–Vogtmann [**BV3**], and the author [**Kie1, Kie3**]. Homomorphisms

$$\mathrm{Out}(F_n) \to \mathrm{GL}_m(\mathbb{Z})$$

or more general representation theory of $\mathrm{Out}(F_n)$ have been studied by Grunewald–Lubotzky [**GL**], Potapchik–Rapinchuk [**PR**], Turchin–Wilwacher [**TW**], and the author [**Kie1, Kie3**].

Let us look at homomorphisms $\mathrm{Out}(F_n) \to \mathrm{Out}(A_\Gamma)$ for more general Γ, or, equivalently, at the ways in which $\mathrm{Out}(F_n)$ can act via outer automorphisms on a RAAG A_Γ. There are two natural ways of constructing non-trivial homomorphisms

$$\varphi\colon \mathrm{Out}(F_n) \to \mathrm{Out}(A_\Gamma)$$

When Γ is a join of two graphs, Δ and Σ say, then $\mathrm{Out}(A_\Gamma)$ contains

$$\mathrm{Out}(A_\Delta) \times \mathrm{Out}(A_\Sigma)$$

as a finite index subgroup. When additionally Δ is isomorphic to the discrete graph with n vertices, then $\mathrm{Out}(A_\Delta) = \mathrm{Out}(F_n)$, and so we have an obvious embedding φ. In fact this method works for a discrete Δ with a very large number of vertices as well, since there are injective maps $\mathrm{Out}(F_n) \to \mathrm{Out}(F_m)$ constructed by Bridson–Vogtmann [**BV3**] for specific values of m growing exponentially with n.

The other way of constructing non-trivial homomorphisms φ becomes possible when Γ contains n vertices with identical stars. In this case, it is immediate that these vertices form a clique Θ, and we have a map

$$\mathrm{GL}_n(\mathbb{Z}) = \mathrm{Aut}(A_\Theta) \to \mathrm{Aut}(A_\Gamma) \to \mathrm{Out}(A_\Gamma)$$

We also have the projection

$$\mathrm{Out}(F_n) \to \mathrm{Out}(H_1(F_n)) = \mathrm{GL}_n(\mathbb{Z})$$

and combining these two maps gives us a non-trivial (though also non-injective) φ. This second method does not work in other situations, due to a result of Wade [**Wad**].

Chapter VII contains the following.

Theorem VII.3.7. *Let $n \geqslant 6$. Suppose that Γ is a simplicial graph with fewer than $\frac{1}{2}\binom{n}{2}$ vertices, which does not contain n distinct vertices with equal stars, and is not a join of the discrete graph with n vertices and another (possibly empty) graph. Then every homomorphism $\mathrm{SOut}(F_n) \to \mathrm{Out}(A_\Gamma)$ is trivial.*

Here $\mathrm{SOut}(F_n)$ denotes the unique index two subgroup of $\mathrm{Out}(F_n)$.

In the course of the proof it becomes important to control actions of $\mathrm{Out}(F_n)$ on small sets. This is done in a somewhat indirect fashion, using the following representation theoretic statement (which is of independent interest). Note that now the representation theoretic approach can be be avoided by using Theorem V.3.16 instead.

Theorem VII.2.27. *Let V be a non-trivial, irreducible $\mathbb{K}$-linear representation of*

$$\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$$

where $n \geqslant 3$, q is a power of a prime p, and where $\mathbb{K}$ is an algebraically closed field of characteristic 0. Then

$$\dim V \geqslant \begin{cases} 2 & \text{if} \quad (n,p) = (3,2) \\ p^{n-1}-1 & \text{otherwise} \end{cases}$$

This result seems not to be present in the literature; it extends a theorem of Landazuri–Seitz [**LS**] yielding a very similar statement for $q = p$ (see Theorem 2.26 in Chapter VII).

Chapter VII also contains

Theorem VII.4.1. *There are no injective homomorphisms $\mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_{\Gamma'})$ when Γ' has fewer vertices than Γ.*

Since the article underlying Chapter VII was written, Imamura [**Ima**] produced new examples of embeddings $\mathrm{Out}(A_\Gamma) \hookrightarrow \mathrm{Out}(A_{\Gamma'})$, and new obstructions prohibiting such embeddings.

We have already seen the Nielsen realisation result for free groups: every finite subgroup of $\mathrm{Out}(F_n)$ can be realised by an action on a finite graph. A similar result holds for torsion subgroups of $\mathrm{Out}(\mathbb{Z}^n) = \mathrm{GL}_n(\mathbb{Z})$: each such subgroup can be realised by an action on a metric torus of dimension n.

These two examples provide motivation for the investigation of Nielsen realisation for general RAAGs, which is the content of Chapter VIII. RAAGs exhibit a natural cubical geometry: each RAAG A_Γ has a classifying space, the *Salvetti complex*, which is a finite cube complex. Moreover, the fundamental group of a cube complex embeds in a RAAG provided that the cube complex is *special* (in the sense of Wise).

The close relationship between RAAGs and cube complexes tempts one to try to achieve Nielsen realisation using cube complexes. This is however bound to fail, since already for finite subgroups of $\mathrm{GL}_n(\mathbb{Z})$ the action on the torus described above cannot (in general) be made cubical and cocompact simultaneously. Thus we look at a subgroup $\mathcal{U}^0(A_\Gamma) \leqslant \mathrm{Out}(A_\Gamma)$, and prove

Corollary VIII.8.4. *Let $\varphi\colon H \to \mathcal{U}^0(A_\Gamma)$ be a homomorphism with a finite domain. Then there exists a metric NPC cube complex X realising φ.*

Here, 'realising' means that H acts on X, and the induced outer action on $\pi_1(X) = A_\Gamma$ is the given one.

The group $\mathcal{U}^0(A_\Gamma) < \mathrm{Out}(A_\Gamma)$ is the intersection of the group $\mathcal{U}(A_\Gamma)$ of untwisted outer automorphisms (introduced by Charney–Stambaugh–Vogtmann [**CSV**]) with the finite index subgroup $\mathrm{Out}^0(A_\Gamma)$ of pure outer automorphisms. Note that for many graphs Γ the groups $\mathcal{U}^0(A_\Gamma)$ and $\mathrm{Out}(A_\Gamma)$ coincide. Note also that the above is still the most general Nielsen realisation result available for automorphisms of RAAGs.

Structure of the thesis. The material contained in this thesis has already appeared in form of several articles; specifically

- Chapter I coincides with [**Kie2**], published in Groups, Geometry and Dynamics in 2015.
- Chapter II coincides with the appendix to [**Bar**] by Laurent Bartholdi, to appear in the Journal of the European Mathematical Society.
- Chapter III coincides with [**FK**] written with Florian Funke.

- Chapter IV coincides with [**HK2**] written with Sebastian Hensel, to appear in the Michigan Mathematical Journal.
- Chapter V coincides with [**BKP**] written with Barbara Baumeister and Emilio Pierro.
- Chapter VI coincides with [**HKS**] written with Thomas Haettel and Petra Schwer, published in Geometriae Dedicata in 2016.
- Chapter VII coincides with [**Kie4**].
- Chapter VIII coincides with [**HK1**] written with Sebastian Hensel.

The mathematical content of each chapter is identical to the article version, however the presentation has been altered in order to make the thesis more coherent.

Acknowledgments. The author would like to thank first and foremost his collaborators: Barbara Baumeister, Florian Funke, Thomas Haettel, Sebastian Hensel, Emilio Pierro, and Petra Schwer. He would also like to thank Kai-Uwe Bux and Ursula Hamenstädt for scientific guidance. Finally, he would like to thank the many people with whom he had mathematical discussions during the research resulting in the work presented here: Brian Bowditch, Martin Bridson, Stefan Friedl, Wolfgang Lück, Andrés Navas, Beatrice Pozzetti, Piotr Przytycki, Yuri Santos Rego, Saul Schleimer, Karen Vogtmann, and Stefan Witzel.

The author was supported by the ERC Grant 'Moduli' of Ursula Hamenstädt, and by the SFB 701 of Bielefeld University.

Free groups

CHAPTER I

Groups with infinitely many ends are not fraction groups

Abstract. We show that any finitely generated group F with infinitely many ends is not a group of fractions of any finitely generated proper subsemigroup P, that is F cannot be expressed as a product PP^{-1}. In particular this solves a conjecture of Navas in the positive. As a corollary we obtain a new proof of the fact that finitely generated free groups do not admit isolated left-invariant orderings.

1. Introduction

The existence of a left-invariant order on a group G is equivalent to the existence of a positive cone $P \subset G$, that is a subsemigroup such that G can be written as a disjoint union $G = \{1\} \sqcup P \sqcup P^{-1}$. In fact there is a one-to-one correspondence between left-invariant orderings and such positive cones.

In this note we prove that whenever a finitely generated group F with infinitely many ends can be written as $F = PP^{-1}$, where P is a finitely generated subsemigroup of F, then $P = F$. Our result answers a question of Navas, who conjectured that finitely generated free groups are not groups of fractions of finitely generated subsemigroups P with $P \cap P^{-1} = \emptyset$.

As an application we obtain a new proof of the fact that the space of left-invariant orderings of a finitely generated free group (endowed with the Chabauty topology) does not have isolated points. This result follows from the work of McCleary [**McC**], but appears in this form for the first time in the work of Navas [**Nav**]. It is worth noting that our proof is the first geometric one.

We also deduce that the left-orderings of finitely generated groups with infinitely many ends do not have finitely generated positive cones. This was already known for free products of left-orderable groups by the work of Rivas [**Riv**].

Our theorem complements a folklore result stating that whenever $\mathcal{S}$ is a finite generating set for a group G, and G does not contain a free subsemigroup, then G is a group of fractions of P, the semigroup generated by $\mathcal{S}$.

We should note here that finitely generated groups with infinitely many ends have been classified by Stallings [**Sta1, Sta2**]. They are precisely those fundamental groups of non-trivial graphs of groups with exactly one edge and a finite edge group, which are finitely generated and not virtually cyclic.

Acknowledgments: the author wishes to thank Andrés Navas for introducing him to this problem, and for his many comments. He also wishes to thank Thomas Haettel and Ursula Hamenstädt for many helpful conversations, and the referee for pointing out ways of significantly improving the presentation of this note.

2. The result

In the following we will use X to denote the (right) Cayley graph of a finitely generated group F with respect to some finite generating set. We will identify F

with vertices of X, and use d to denote the standard metric on the Cayley graph X. The isometric left-action of F on X and its subsets will be denoted by left multiplication. The notation $B(x, \xi)$ will stand for the closed ball centred at x of radius ξ.

We will assume that F has infinitely many ends, and so there will exist a constant κ such that the ball $B = B(1, \kappa)$ disconnects X into a space with at least 3 infinite components. We will use S to denote the set of vertices of B.

Definition 2.1. We say that $A \subset X$ is a *shoot* if and only if there exists $w \in F$ such that A is a connected component of $X \smallsetminus wB$. We say that wB *bounds* the shoot.

Lemma 2.2. *Let A be an infinite shoot bounded by B. Then there exists $w \in F$ such that $w(X \smallsetminus A) \subseteq A$ and $w^{-1}(X \smallsetminus A) \subseteq A$.*

Proof. Note that the ball $B(1, 2\kappa)$ is finite, since X is locally finite. Since F has infinitely many ends and A is infinite, there exists λ such that

$$L = \{x \in F \mid d(1, x) = \lambda\} \cap A$$

has more than $|B(1, 2\kappa)|$ elements. Take $l \in L$. The cardinality of L guarantees that there exists $l' \in L$ such that $l'l^{-1} \notin B(1, 2\kappa)$.

Let $w = l'l^{-1}$. Observe that

$$d(w, l') = d(l'l^{-1}, l') = d(1, l) = \lambda$$

Consider a shortest path between w and l'. If it lies entirely in A, then in particular so does w. If not, then it must contain some point $b \in B$, since B bounds A. Now we have

$$\lambda = d(w, l') = d(w, b) + d(b, l') \geqslant d(w, b) + \lambda - \kappa$$

which implies that $d(w, b) \leqslant \kappa$, and hence that $w \in B(1, 2\kappa)$, which is a contradiction. We have thus established that $w \in A \smallsetminus B(1, 2\kappa)$, and therefore that $wB \subset A$.

Note that $w^{-1} = ll'^{-1} \notin B(1, 2\kappa)$ enjoys the same properties as w, and so we immediately conclude that $w^{-1}B \subset A$, or equivalently that $B \subset wA$.

Since wB and B are disjoint, every shoot bounded by wB either contains B or is disjoint from it. Clearly, there is a unique shoot bounded by wB containing B, and we have already shown that it is wA. Each of the other shoots bounded by wB lies in a single shoot bounded by B, namely in the shoot bounded by B which contains wB. But we have already seen that this is A. We are left with the conclusion that $w(X \smallsetminus A) \subset A$, and our proof is finished by making the analogous observations for w^{-1}. □

We are now ready for the main result.

Theorem 2.3. *Let P be a finitely generated subsemigroup of a finitely generated group F with infinitely many ends. If $PP^{-1} = F$ then $P = F$.*

Proof. For ease of notation we will refer to the elements of P as positive, and to the elements of P^{-1} as negative.

We first note that any finite generating set of P is a generating set for F. Let X be the Cayley graph of F with respect to some such generating set. Note that this allows us to view generators of P as positive edges of X, and hence any positive element $p \in P$ is realised by a positive path between 1 and the vertex p in X.

We will use the notation κ, B and S as defined above.

Step 1: We claim that $S(P^{-1} \cup \{1\}) = F$.

If P intersects each ball $B(x, \kappa) = xB$ then each $x \in F$ is a concatenation of an element of P (namely any positive path from 1 to xB) with an element in S

(connecting the end of the positive path to the centre of the ball). Thus we have $x \in PS$, and our claim follows by taking inverses.

Let us now suppose that there exists an $x \in F$ such that

$$P \cap xB = \emptyset$$

Let A_0 denote an infinite shoot bounded by xB such that $1 \notin A_0$.

Let $z \in F \smallsetminus S$ be any element, and let A be the shoot bounded by B containing z. We claim that there exists $y \in F$ such that $yA \subseteq A_0$.

There are two cases we need to consider. The first one occurs when

$$xA \subseteq A_0$$

in which case we take $y = x$. The other one (illustarted in Figure 2.1) occurs when $xA \not\subseteq A_0$, that is when xA is a shoot bounded by xB other than A_0. Lemma 2.2 applied to $x^{-1}A_0$ gives us an element $w \in F$ such that $w(X \smallsetminus x^{-1}A_0) \subseteq x^{-1}A_0$. So $y = xw$ satisfies

$$yA = xwA \subseteq xx^{-1}A_0 = A_0$$

and so we have proven the claim.

Now, since $yz \in F = PP^{-1}$, we can write $yz = pq$, where p is positive and q is negative. Since there are no positive elements in xB by assumption, we see that $p \notin A_0$, and therefore q is a negative path connecting a vertex $p \in X \smallsetminus A_0$ to $yz \in yA \subseteq A_0$. The shoot yA is bounded by yB and contained in A_0, hence any path from $X \smallsetminus A_0$ to yA has to cross yB. This is in particular true for q, so there is a negative path (a terminal subpath of q) from some vertex of yB to yz, and hence from a vertex of B to z (after translating by y^{-1}). In the group language we have thus shown that $z \in SP^{-1}$, and so

$$F \smallsetminus S \subseteq SP^{-1}$$

But clearly $S \subset S(P^{-1} \cup \{1\})$, and so we have proven the claim of step 1.

Step 2: We claim that $P = F$.

We have established above that $S(P^{-1} \cup \{1\}) = F$, with S being finite. Let Q be a minimal (with respect to cardinality) finite subset of F such that $Q(P^{-1} \cup \{1\}) = F$. Suppose that there exist distinct $q, q' \in Q$. Then $q^{-1}q' \in F = PP^{-1}$, and so $q^{-1}q' = ab^{-1}$ with $a, b \in P$. Hence

$$q, q' \in qaP^{-1}$$

and therefore we could replace Q by $(Q \cup \{qa\}) \smallsetminus \{q, q'\}$ of smaller cardinality. This shows that $|Q| = 1$. Without loss of generality we can take $Q = \{1\}$, and thence get

$$P^{-1} \cup \{1\} = F$$

Now let $f \in F \smallsetminus \{1\}$. We have $f, f^{-1} \in P^{-1}$, and since P^{-1} is a semigroup, also $1 = ff^{-1} \in P^{-1}$. So $P^{-1} = F$. Taking an inverse concludes the theorem. ☐

We now easily deduce the following.

Corollary 2.4. *Let F be a finitely generated group with infinitely many ends. Then F does not allow a left-invariant ordering with a finitely generated positive cone.*

Proof. Let P be the positive cone of a left-invariant ordering of F. Then

$$F = P \cup P^{-1} \cup \{1\}$$

and so in particular $F = PP^{-1}$. But also $P \cap P^{-1} = \emptyset$, and so $P \neq F$. Now the contrapositive of Theorem 2.3 tells us that P is not finitely generated. ☐

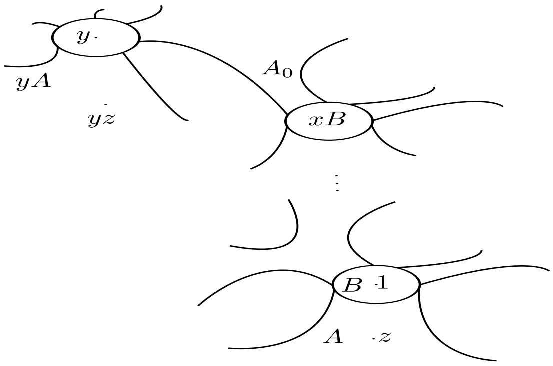

FIGURE 2.1. Step 1 of the theorem.

The statement above follows from the work of Rivas [**Riv**], since left-orderable groups are torsion-free, and so they have infinitely many ends only when they are free products.

We also get the following corollary.

Corollary 2.5. *The space of left-invariant orderings on any finitely generated free group has no isolated points.*

Proof. Let P be the positive cone of an isolated ordering of F, a finitely generated free group. By above, P is not finitely generated.

The order defined by P is isolated, and so there exists a finite set $S \subset F$ such that whenever we have another positive cone of an ordering P' such that $P \cap S = P' \cap S$, then $P = P'$. However the work of Smith and Clay [**CS**, Theorem E] allows us to construct an order (in fact infinitely many such orders) whose positive cone P' satisfies $P \cap S = P' \cap S$, but such that $P \neq P'$. This is a contradiction. □

Corollary 2.6. *Let G be a group admitting a finite Garside structure. Then G has at most two ends.*

Proof. A finite Garside structure gives us a finitely generated subsemigroup P such that $G = PP^{-1}$, and $P \neq G$. Thus G cannot have infinitely many ends. □

CHAPTER II

A characterisation of amenability via Ore domains

This appeared as an appendix to an article of Laurent Bartholdi.

Abstract. We prove that a group algebra $\mathbb{K}G$ without non-trivial zero-divisors is an Ore domain if and only if G is amenable, thus proving a folklore conjecture attributed to Guba.

1. Introduction

An *Ore domain* is a ring R which satisfies the Ore condition, that is for any $p, q \in R$ where q is not a zero-divisor, there exist $r, s \in R$ such that r is not a zero-divisor and

$$pr = qs$$

It was shown by Tamari [**Tam**] that if G is amenable, then $\mathbb{K}G$, the group algebra of G with coefficients in a field $\mathbb{K}$, satisfies the Ore condition. On the other hand, it is easy to see that no group algebra of F_2, the free group on two letters, satisfies the Ore condition. This implies that any group containing F_2 does not have a group algebra satisfying the Ore condition.

In view of the above, the question of establishing the Ore condition becomes interesting for the group algebras of the counterexamples to the von Neumann–Day Conjecture, that is for non-amenable groups which do not contain non-abelian free groups. So far, the condition has not been verified for a group algebra of a single such group.

A folklore conjecture (sometimes attributed to Guba) states that the converse to Tamari's result holds true, that is any group G with $\mathbb{K}G$ satisfying the Ore condition is amenable. We confirm this conjecture under the additional assumption that $\mathbb{K}G$ has no non-trivial zero-divisors.

Theorem 2.2. *Let G be a group, and let $\mathbb{K}$ be a field such that $\mathbb{K}G$ has no non-trivial zero divisors. Then G is amenable if and only if $\mathbb{K}G$ satisfies the Ore condition.*

2. The result

We restrict ourselves to group rings without zero divisors – conjecturally, these coincide with group rings of torsion-free groups (this is the content of the notorious conjecture of Kaplansky). Note that if G has torsion, then the question of when $\mathbb{K}G$ satisfies the Ore condition with respect to its regular elements is more complicated; for example, Linnell, Lück and Schick prove in [**LLS**] that $H \wr \mathbb{Z}$ is never an Ore ring, for H a finite, non-trivial group.

A crucial fact we will use is the following theorem of Bartholdi.

Theorem 2.1 ([**Bar**])**.** *Let $\mathbb{K}$ be a field and G a non-amenable group. Then there exists a finite extension $\mathbb{L}$ of $\mathbb{K}$ and an $n \times n$ matrix M over $\mathbb{L}G$ such that the induced $\mathbb{L}G$-linear map*

$$M\colon\ \mathbb{L}G^n \to \mathbb{L}G^n$$

is injective, but the last row of M consists only of zeros.

We prove

Theorem 2.2. *Let G be a group, and let $\mathbb{K}$ be a field such that $\mathbb{K}G$ has no zero divisors. Then G is amenable if and only if $\mathbb{K}G$ satisfies the Ore condition.*

Proof. When G is amenable, the result is due to Tamari [**Tam**].

Assume that G is non-amenable. Theorem 2.1 yields a finite extension $\mathbb{L}$ of $\mathbb{K}$ and an $n \times n$ matrix M over $\mathbb{L}G$. Forgetting that last row of M and using the isomorphism of $\mathbb{K}G$-modules $\mathbb{L} =\cong \mathbb{K}^d$ for some d, we obtain an exact sequence of free $\mathbb{K}G$-modules

$$0 \longrightarrow (\mathbb{K}G)^{dn} \longrightarrow (\mathbb{K}G)^{d(n-1)}. \tag{1}$$

Suppose now that $\mathbb{K}G$ is an Ore domain; then $\mathbb{K}G$ embeds into its classical field of fractions $\mathbb{F}$. Crucially, $\mathbb{F}$ is a flat $\mathbb{K}G$ module, that is the functor $- \otimes_{\mathbb{K}G} \mathbb{F}$ preserves exactness of sequences (see e.g. [**MR**, Proposition 2.1.16]). Also, $\mathbb{F}$ is a skew field, and upon tensoring (1) with $\mathbb{F}$ we obtain an exact sequence

$$0 \longrightarrow \mathbb{F}^{dn} \longrightarrow \mathbb{F}^{d(n-1)}$$

which is impossible for reasons of dimension. □

Automorphisms of free groups

CHAPTER III

Alexander and Thurston norms, and the Bieri–Neumann–Strebel invariants for free-by-cyclic groups

This is joint work with Florian Funke.

Abstract. We investigate Friedl–Lück's universal L^2-torsion for descending HNN extensions of finitely generated free groups, and so in particular for F_n-by-$\mathbb{Z}$ groups. This invariant induces a semi-norm on the first cohomology of the group which is an analogue of the Thurston norm for 3-manifold groups.
For descending HNN extensions of F_2, we prove that this Thurston semi-norm is an upper bound for the Alexander semi-norm defined by McMullen, as well as for the higher Alexander semi-norms defined by Harvey. The same inequalities are known to hold for 3-manifold groups.
We also prove that the Newton polytopes of the universal L^2-torsion of a descending HNN extension of F_2 locally determine the Bieri–Neumann–Strebel invariant of the group. We give an explicit means of computing the BNS invariant for such groups.
When the HNN extension is taken over F_n along a polynomially growing automorphism with unipotent image in $\mathrm{GL}(n,\mathbb{Z})$, we show that the Newton polytope of the universal L^2-torsion and the BNS invariant completely determine one another. We also show that in this case the Alexander norm, its higher incarnations, and the Thurston norm all coincide.

1. Introduction

Whenever a free finite G-CW-complex X is L^2-acyclic, i.e. its L^2-Betti numbers vanish, a secondary invariant called the L^2-torsion $\rho^{(2)}(X;\mathcal{N}(G))$ enters the stage [**Lüc1**, Chapter 3]. It takes values in $\mathbb{R}$ and captures in many cases geometric data associated to X: If X is a closed hyperbolic 3-manifold, then it was shown by Lück and Schick [**LS**] that

$$\rho^{(2)}(\widetilde{X};\mathcal{N}(\pi_1(X))) = -\frac{1}{6\pi}\cdot \mathrm{vol}(X)$$

and if X is the classifying space of a free-by-cyclic group $F_n \rtimes_g \mathbb{Z}$, with $g \in \mathrm{Aut}(F_n)$, then $-\rho^{(2)}(\widetilde{X}; F_n \rtimes_g \mathbb{Z})$ gives a lower bound on the growth rates of g, as shown by Clay [**Cla**, Theorem 5.2].

Many generalisations of the L^2-torsion have been constructed, e.g. the L^2-Alexander torsion (by Dubois–Friedl–Lück [**DFL**]) and L^2-torsion function, or more generally L^2-torsion twisted with finite-dimensional representations (by Lück [**Lüc2**]).

In a series of papers, Friedl and Lück [**FL2, FL1, FL3**] constructed the *universal L^2-torsion* $\rho_u^{(2)}(X;\mathcal{N}(G))$ for any free finite L^2-acyclic G-CW-complex. It takes values in $\mathrm{Wh}^w(G)$, a weak version of the Whitehead group of G which is adapted to the setting of L^2-invariants. The Fuglede–Kadison determinant induces a map $\mathrm{Wh}^w(G) \to \mathbb{R}$ taking $\rho_u^{(2)}(X;\mathcal{N}(G))$ to $\rho^{(2)}(X;\mathcal{N}(G))$, and similar maps with $\mathrm{Wh}^w(G)$ as their domain take the universal L^2-torsion to the aforementioned generalisations of L^2-torsion.

Assuming that G satisfies the Atiyah Conjecture, Friedl–Lück [**FL3**] construct a *polytope homomorphism*

$$\mathbb{P}\colon \mathrm{Wh}^w(G) \to \mathcal{P}_T(H_1(G)_f)$$

where $H_1(G)_f$ denotes the free part of the first integral homology of G, and $\mathcal{P}_T(H_1(G)_f)$ denotes the Grothendieck group of the commutative monoid whose elements are polytopes in $H_1(G)_f \otimes \mathbb{R}$ (up to translation) with pointwise addition (also called *Minkowski sum*). The image of $-\rho_u^{(2)}(X;\mathcal{N}(G))$ under $\mathbb{P}$ is the *L^2-torsion polytope of X*, denoted by $P_{L^2}(X;G)$. If $M \neq S^1 \times D^2$ is a compact connected aspherical 3-manifold with empty or toroidal boundary such that $\pi_1(M)$ satisfies the Atiyah Conjecture, then it is shown in [**FL3**, Theorem 3.27] that $P_{L^2}(\widetilde{M};\pi_1(M))$ induces another well-known invariant of M, the *Thurston norm*

$$\|\cdot\|_T\colon H^1(M;\mathbb{R}) \to \mathbb{R}$$

This semi-norm was defined by Thurston [**Thu**] and is intimately related to the question of the manifold fibering over the circle.

McMullen [**McM**] constructed an *Alexander semi-norm* from the Alexander polynomial and showed that it provides a lower bound for the Thurston semi-norm. This was later generalised by Harvey [**Har2**] to higher Alexander semi-norms

$$\delta_n\colon H^1(M;\mathbb{R}) \to \mathbb{R}$$

Friedl–Lück's theory can also be applied to free-by-cyclic groups, or more generally to descending HNN extensions $G = F_n*_g$, with g an injective endomorphism of F_n, and yields in this context a semi-norm

$$\|\cdot\|_T\colon H^1(G;\mathbb{R}) \to \mathbb{R}$$

which we also call *Thurston norm* due to the analogy with the 3-manifold setting. In the case $n = 2$, we build a similar picture as for 3-manifolds and prove that this semi-norm is an upper bound for McMullen–Harvey's Alexander semi-norms:

Theorem 4.7. *Let $G = F_2*_g$ be a descending HNN extension of F_2 such that the first Betti number satisfies $b_1(G) \geqslant 2$. Then the Thurston and higher Alexander semi-norms satisfy for all $n \geqslant 0$ and $\varphi \in H^1(G;\mathbb{R})$ the inequality*

$$\delta_n(\varphi) \leqslant \|\varphi\|_T$$

We extend this result to higher rank free groups for a particular type of automorphism called *UPG* (see Definition 6.1) where we even obtain an equality:

Corollary 6.6. *Let $G = F_n \rtimes_g \mathbb{Z}$ with $n \geqslant 2$ and g a UPG automorphism. Let $\varphi \in H^1(G;\mathbb{R})$. Then for all $k \geqslant 0$ we have*

$$\delta_k(\varphi) = \|\varphi\|_T.$$

In the case of two-generator one-relator groups G with $b_1(G) = 2$, the L^2-torsion polytope has been studied by Friedl–Tillmann [**FT**]. They established a close connection between $P_{L^2}(G) := P_{L^2}(EG;G)$ and the Bieri–Neumann–Strebel invariant $\Sigma(G)$. We prove similar results in our setting:

Theorem 5.13. *Let $g\colon F_2 \to F_2$ be a monomorphism and let $G = F_2*_g$ be the associated descending HNN extension. Given $\varphi \in \mathrm{Hom}(G,\mathbb{R}) \smallsetminus \{0\}$ such that $-\varphi$ is not the epimorphism induced by F_2*_g, there exists an element $d \in \mathcal{D}(G)^\times$ such that:*

(1) *The image of* d *under the quotient maps*

$$\mathcal{D}(G)^{\times} \to \mathcal{D}(G)^{\times}/[\mathcal{D}(G)^{\times}, \mathcal{D}(G)^{\times}] \cong K_1^w(\mathbb{Z}G) \to \mathrm{Wh}^w(G)$$

is $-\rho_u^{(2)}(G)$. *In particular* $P_{L^2}(G) = P(d)$ *in* $\mathcal{P}_T(H_1(G)_f)$.

(2) *There exists an open neighbourhood* U *of* $[\varphi]$ *in* $S(G)$ *such that for every* $\psi \in \mathrm{Hom}(G,\mathbb{R}) \smallsetminus \{0\}$ *which satisfies* $[\psi] \in U$ *and is* d*-equivalent to* φ, *we have* $[-\varphi] \in \Sigma(G)$ *if and only if* $[-\psi] \in \Sigma(G)$.

The d-equivalence is induced by the Newton polytopes associated to d in a simple way (see Definition 5.11). Over arbitrary rank we can strengthen this result again for UPG automorphisms:

Corollary 6.4. *Let* $G = F_n \rtimes_g \mathbb{Z}$ *with* $n \geqslant 2$ *and* g *a UPG automorphism. Let* $\varphi \in H^1(G;\mathbb{R})$. *Then* $[\varphi] \in \Sigma(G)$ *if and only if* $F_\varphi(P_{L^2}(G)) = 0$ *in* $\mathcal{P}_T(H_1(G)_f)$.

The face map F_φ is defined in Definition 5.10. This theorem is motivated by Cashen-Levitt's computation of the BNS invariant of such groups [**CL**, Theorem 1.1].

Finally, we formulate a question about the Newton polytopes of two different notions of a determinant for certain square matrices over $\mathbb{Z}G$ (Question 3.6). This purely algebraic statement would immediately yield the inequality of semi-norms $\delta_n(\cdot) \leqslant \|\cdot\|_T$ also for descending HNN extensions of higher rank free groups.

Acknowledgments: The first-named author was supported by GRK 1150 'Homotopy and Cohomology' funded by the DFG, the Max Planck Institute for Mathematics, and the Deutsche Telekom Stiftung.

The second-named author was supported by the ERC grant 'Moduli' of Ursula Hamenstädt, and the SFB 701 'Spectral Structures and Topological Methods in Mathematics' of the Bielefeld University.

The authors would like to thank Stefan Friedl and Wolfgang Lück for helpful discussions, and the organisers of the 'Manifolds and Groups' conference on Ventotene, where some of this work was conducted.

2. Preliminaries

2.1. Descending HNN extensions.

Definition 2.1. Let G be a group, $H \leqslant G$ a subgroup, and $g\colon H \to H$ a monomorphism. The *HNN extension associated to* g is the quotient of the free product of G with $\langle t \rangle \cong \mathbb{Z}$ by

$$\langle\langle \{t^{-1}xtg(x)^{-1} \mid x \in H\} \rangle\rangle$$

The element t is called the *stable letter* of the HNN extension. The HNN extension is called *descending* if $H = G$. The natural epimorphism $G*_g \to \mathbb{Z}$, sending t to 1 with G in its kernel, is called the *induced epimorphism.*

Remark 2.2. Note that when $g\colon G \to G$ is an isomorphism, then $G*_g = G \rtimes_g \mathbb{Z}$ is a semi-direct product, or a G-by-$\mathbb{Z}$ group (since extensions with a free quotient always split).

In the final sections of this paper we will focus on descending HNN extensions $G = F_2*_g$. The following (well-known) result illustrates that this is somewhat less restrictive than it might seem.

Proposition 2.3. *Let* $g\colon F_2 \to F_2$ *be a monomorphism which is not onto. There exists* $N \in \mathbb{N}$ *such that for every* $n \geqslant N$ *there exists a monomorphism* $g_n\colon F_n \to F_n$ *such that*

$$F_2*_g \cong F_n*_{g_n}$$

Proof. We start by observing that Marshall Hall's theorem [**Hal**] tells us that there exists $N \in \mathbb{N}$ such that $g(F_2)$ is a free factor of a finite index subgroup F_N of F_2. In fact it is easy to see (using the proof of Stallings [**Sta3**]) that this statement holds for any $n \geqslant N$ (here we are using the fact that g is not onto; otherwise $N = 2$ and we cannot take larger values of n).

Now g factors as

$$F_2 \xrightarrow{a} F_n \xrightarrow{b} F_2$$

where a embeds F_2 as a free factor, and b is an embedding with image of finite index. We let $g_n = a \circ b \colon F_n \to F_n$.

Next we construct the desired isomorphism. Let t (resp. s) denote the stable letter of $F_2 *_g$ (resp. $F_n *_{g_n}$). Let $F_2 = \langle x_1, x_2 \rangle$ and $F_n = \langle x_1, \ldots, x_n \rangle$; with this choice of generators the map a becomes the identity.

Consider $h \colon F_2 *_g \to F_n *_{g_n}$ defined by

$$h(x_i) = x_i \text{ and } h(t) = s$$

It is a homomorphism since

$$t^{-1} x_i t = b(x_i)$$

and

$$h(t^{-1})h(x_i)h(t) = s^{-1} x_i s = b(x_i) = h(b(x_i))$$

Now consider $h' \colon F_n *_{g_n} \to F_2 *_g$ induced by

$$h'(x_i) = t b(x_i) t^{-1} \text{ and } h'(s) = t$$

It is clear that h' is the inverse of h. □

Remark 2.4. Of course there is nothing special about F_2 in the above result. The proof works verbatim when F_2 is replaced by F_m with $m \geqslant 2$.

2.2. Dieudonné determinant. While working with the universal L^2-torsion, the Dieudonné determinant for matrices over skew-fields is of fundamental importance. We review here its definition and fix a so-called *canonical representative*.

Definition 2.5. Given a ring R, we will denote its group of units by $R^\times$.

Definition 2.6 (Dieudonné determinant)**.** Given a skew field $\mathcal{D}$ and an integer n, let $M_n(\mathcal{D})$ denote the ring of $n \times n$ matrices over $\mathcal{D}$. The *Dieudonné determinant* is a multiplicative map

$$\det_{\mathcal{D}} \colon M_n(\mathcal{D}) \to \mathcal{D}^\times / [\mathcal{D}^\times, \mathcal{D}^\times] \cup \{0\}$$

defined as follows: First we construct its *canonical representative*

$$\det^{\mathrm{c}}_{\mathcal{D}} \colon M_n(\mathcal{D}) \to \mathcal{D}$$

and then set $\det_{\mathcal{D}}(A)$ to be image of $\det^{\mathrm{c}}_{\mathcal{D}}(A)$ under the obvious map

$$\mathcal{D} \to \mathcal{D}^\times / [\mathcal{D}^\times, \mathcal{D}^\times] \cup \{0\}$$

The canonical representative is defined inductively:

- for $n = 1$ we have $\det^{\mathrm{c}}_{\mathcal{D}}((a_{11})) = a_{11}$;
- if the last column of A contains only zeros we set $\det^{\mathrm{c}}_{\mathcal{D}}(A) = 0$;
- for general n (and a matrix A with non-trivial last column) we first identify the bottommost non-trivial element in the last column of A. If this is a_{nn} we take $P = \mathrm{id}$; otherwise, if the element is a_{in}, we take P to be the permutation matrix which swaps the i^{th} and n^{th} rows of A; in either case we have $PA = A' = (a'_{ij})$ with $a'_{nn} \neq 0$. Now we define $B = (b_{ij})$ by

$$b_{ij} = \begin{cases} 1 & \text{if} \quad i = j \\ 0 & \text{if} \quad i \neq j \text{ and } j < n \\ -a'_{in} {a'_{nn}}^{-1} & \text{if} \quad i \neq j = n \end{cases}$$

This way we have

$$BPA = A'' = (a''_{ij})$$

with $a''_{in} = 0$ for all $i \neq n$. Let us set C to be the $(n-1) \times (n-1)$ matrix $C = (a''_{ij})_{i,j<n}$. We define

$$\det{}^{c}_{\mathcal{D}}(A) = \det P \cdot \det{}^{c}_{\mathcal{D}}(C) \cdot a''_{nn}$$

Note that the canonical representative $\det^{c}_{\mathcal{D}}$ is not multiplicative, but the determinant itself is, as shown by Dieudonné [**Die**].

It is immediate from the definition that when $\mathcal{D}$ is a commutative field, then the Dieudonné determinant agrees with the usual determinant.

Proposition 2.7 (Formula for square matrices)**.**

$$\det{}^{c}_{\mathcal{D}} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{cases} ad - bd^{-1}cd & \text{if } d \neq 0 \\ -bc & \text{if } d = 0 \end{cases}$$

2.3. Crossed products.

Definition 2.8 (Crossed product group ring)**.** Let R be a ring and G a group together with maps of sets $\varphi\colon G \to \mathrm{Aut}(R)$ and $\mu\colon G \times G \to R^{\times}$ such that

$$\varphi(g) \circ \varphi(g') = c(\mu(g, g')) \circ \varphi(gg')$$
$$\mu(g, g') \cdot \mu(gg', g'') = \varphi(g)(\mu(g', g'')) \cdot \mu(g, g'g'')$$

where $c\colon R^{\times} \to \mathrm{Aut}(R)$ maps an invertible element r to the conjugation by r on the left. Then the *crossed product group ring* $R * G$ is the free left R-module with basis G and multiplication induced by the rule

$$(\kappa g) \cdot (\lambda h) = \kappa\varphi(g)(\lambda)\mu(g, h)gh \tag{2}$$

for any $g, h \in G$ and $\kappa, \lambda \in R$. The conditions on μ and φ ensure the associativity of the multiplication, so that $R * G$ is indeed a ring.

Note that when φ and μ are trivial, we obtain the usual group ring RG.

Example 2.9. Crossed product group rings appear naturally: Given an extension of groups

$$1 \to K \to G \to Q \to 1$$

we can identify $RG \cong (RK) * Q$, where the structure maps φ and μ are defined as follows: Let $s\colon Q \to G$ be a set-theoretic section of the given epimorphism $G \to Q$. Define

$$\varphi(q) \left(\sum_{k \in K} a_k \cdot k \right) = \sum_{k \in K} a_k \cdot s(q)ks(q)^{-1}$$

and

$$\mu(q, q') = s(q)s(q')s(qq')^{-1} \in K$$

The isomorphism $(RK) * Q \to RG$ is given by

$$\sum_{q \in Q} \lambda_q \cdot q \mapsto \sum_{q \in Q} \lambda_q \cdot s(q)$$

Definition 2.10. Given an element $x = \sum_{h \in G} \lambda_h \cdot h \in R * G$ we define its *support* to be

$$\mathrm{supp}(x) = \{h \in G \mid \lambda_h \neq 0\}$$

Note that the support is a finite subset of G.

2.4. Ore localisation. We briefly review non-commutative localisation.

Definition 2.11. Let R be a unital ring without zero-divisors, and let $T \subseteq R$ be a subset containing 1 such that for every $s, t \in T$ we also have $st \in T$. Then T satisfies the *(left) Ore condition* if for every $r \in R$, $t \in T$ there are $r' \in R$, $t' \in T$ such that $t'r = r't$.

One can then define a ring $T^{-1}R$, called the *Ore localisation*, whose elements are fractions $t^{-1}r$ with $r \in R, t \in T$, subject to the usual equivalence relation. There is an obvious ring monomorphism $R \to T^{-1}R$.

One instance of the Ore localisation will be of particular interest in this paper. If G is an amenable group, $\mathcal{D}$ a skew field and $\mathcal{D} * G$ a crossed product which is a domain, then a result of Tamari [**Tam**] shows that $\mathcal{D} * G$ satisfies the left (and right) Ore condition with respect to the non-zero elements in $\mathcal{D} * G$. This applies in particular to the case where G is finitely generated free-abelian.

Throughout the paper, we will only take the Ore localisation with respect to all non-zero elements of a ring.

2.5. The Atiyah Conjecture and $\mathcal{D}(G)$. In this section we review techniques which were originally developed for proving the Atiyah Conjecture, but have meanwhile been shown to be fruitful on many other occasions.

Given a group G, let $L^2(G)$ to denote the complex Hilbert space with Hilbert basis G on which G acts by translation. We use $\mathcal{N}(G)$ to denote the *group von Neumann algebra of* G, i.e. the algebra of bounded G-equivariant operators on $L^2(G)$. Associated to any $\mathcal{N}(G)$-module M (in the purely ring-theoretic sense), there is a *von Neumann dimension* $\dim_{\mathcal{N}(G)}(M) \in [0, \infty]$ (see [**Lüc1**, Chapter 6]).

Conjecture 2.12 (Atiyah Conjecture)**.** *Let G be a torsion-free group. Given a matrix $A \in \mathbb{Q}G^{m \times n}$, we denote by $r_A : \mathcal{N}(G)^m \to \mathcal{N}(G)^n$ the $\mathcal{N}(G)$-homomorphism given by right multiplication with A. Then G satisfies the Atiyah Conjecture if for every such matrix the number $\dim_{\mathcal{N}(G)}(\ker(r_A))$ is an integer.*

The class of groups for which the Atiyah Conjecture is known to be true is large. It includes all free groups, is closed under taking directed unions, as well as extensions with elementary amenable quotients. Infinite fundamental groups of compact connected orientable irreducible 3-manifolds with empty or toroidal boundary which are not closed graph manifolds are also known to satisfy the Atiyah Conjecture. For these statements and more information we refer to [**FL1**, Chapter 3].

Definition 2.13. Let $R \subseteq S$ be a ring extension. Then the *division closure* of R inside S is the smallest subring D of S which contains R, such that every element in D which is invertible in S is already invertible in D. We denote it by $\mathcal{D}(R \subseteq S)$.

Let $\mathcal{U}(G)$ denote the algebra of affiliated operators of $\mathcal{N}(G)$. This algebra is carefully defined and examined in [**Lüc1**, Chapter 8]. Note that $\mathbb{Q}G$ embeds into $\mathcal{N}(G)$, and therefore $\mathcal{U}(G)$, as right multiplication operators. Let $\mathcal{D}(G)$ denote the division closure of $\mathbb{Q}G$ inside $\mathcal{U}(G)$.

The following theorem appears in [**Lüc1**, Lemma 10.39] for the case where $\mathbb{Q}G$ is replaced by $\mathbb{C}G$ in the above definitions, but the proof also carries over to rational coefficients.

Theorem 2.14. *A torsion-free group satisfies the Atiyah Conjecture if and only if $\mathcal{D}(G)$ is a skew field.*

It is known that if $H \subseteq G$ is a subgroup, then there is a canonical inclusion $\mathcal{D}(H) \subseteq \mathcal{D}(G)$.

Recall from Example 2.9 that for an extension of groups

$$1 \to K \to G \to Q \to 1$$

the group ring $\mathbb{Z}G$ is isomorphic to the crossed product $\mathbb{Z}K * Q$, where Q acts on $\mathbb{Z}K$ by conjugation. When G satisfies the Atiyah Conjecture, this action extends to an action on $\mathcal{D}(K)$ and one can identify the crossed product $\mathcal{D}(K) * Q$ with a subring of $\mathcal{D}(G)$ (see [**Lüc1**, Lemma 10.58]). If Q is finitely generated free-abelian, then $\mathcal{D}(K) * Q$ satisfies the Ore condition with respect to the non-zero elements T and the Ore localisation admits by [**Lüc1**, Lemma 10.69] an isomorphism

$$T^{-1}(\mathcal{D}(K) * Q) \xrightarrow{\cong} \mathcal{D}(G) \tag{3}$$

2.6. Universal L^2-torsion. Let G be a group satisfying the Atiyah Conjecture. In [**FL3**, Definition 1.1], Friedl and Lück define the *weak K_1-group* $K_1^w(\mathbb{Z}G)$ as the abelian group generated by $\mathbb{Z}G$-endomorphisms $f\colon \mathbb{Z}G^n \to \mathbb{Z}G^n$ that become a weak isomorphism (a bounded injective operator with dense image) upon applying $-\otimes_{\mathbb{Z}G} L^2(G)$, subject to the usual relations in K_1. The above condition is equivalent to f becoming invertible after applying $-\otimes_{\mathbb{Z}G} \mathcal{D}(G)$ (see [**FL3**, Lemma 1.21]). The *weak Whitehead group* $\mathrm{Wh}^w(G)$ *of* G is defined as the quotient of $K_1^w(\mathbb{Z}G)$ by $\{\pm g \mid g \in G\}$ considered as endomorphisms of $\mathbb{Z}G$ via right multiplication. An injective group homomorphism $i\colon G \to H$ induces maps

$$i_*\colon K_1^w(\mathbb{Z}G) \to K_1^w(\mathbb{Z}H)$$
$$i_*\colon \mathrm{Wh}^w(G) \to \mathrm{Wh}^w(H)$$

Example 2.15. For H a finitely generated free-abelian group, we have isomorphisms

$$K_1^w(\mathbb{Z}H) \cong K_1(T^{-1}(\mathbb{Z}H)) \cong T^{-1}(\mathbb{Z}H)^\times$$

where T denotes the set of non-trivial elements of $\mathbb{Z}H$. The first isomorphism is a special case of the main result of [**LL**] by Linnell–Lück, and the second one is well-known and induced by the Dieudonné determinant over the field $T^{-1}(\mathbb{Z}H)$.

A $\mathbb{Z}G$-chain complex is called *based free* if every chain module is free and has a preferred basis. Given an L^2-acyclic finite based free $\mathbb{Z}G$-chain complex C_*, Friedl-Lück [**FL3**, Definition 1.7] define the *universal L^2-torsion of* C_*

$$\rho_u^{(2)}(C_*; \mathcal{N}(G)) \in K_1^w(\mathbb{Z}G)$$

in a similar fashion as the Whitehead torsion.

If X is an L^2-acyclic finite free G-CW-complex, then its cellular chain complex $C_*(X)$ is finite and free, and we equip it with some choice of bases coming from the CW-structure. Since this is only well-defined up to multiplication by elements in G, the *universal L^2-torsion* $\rho_u^{(2)}(X; \mathcal{N}(G)) \in \mathrm{Wh}^w(G)$ *of* X is defined as the image of $\rho_u^{(2)}(C_*(X); \mathcal{N}(G))$ under the projection $K_1^w(\mathbb{Z}G) \to \mathrm{Wh}^w(G)$.

A finite connected CW-complex X is *L^2-acyclic* if its universal cover $\widetilde{X}$ is an L^2-acyclic $\pi_1(X)$-CW-complex. If this is the case, then the *universal L^2-torsion of* X is

$$\rho_u^{(2)}(\widetilde{X}) := \rho_u^{(2)}(\widetilde{X}; \mathcal{N}(\pi_1(X))) \in \mathrm{Wh}^w(\pi_1(X))$$

If X is a (possible disconnected) finite CW-complex, then it is *L^2-acyclic* if each path component is L^2-acyclic in the above sense. In this case, its *universal L^2-torsion* is defined by

$$\rho_u^{(2)}(\widetilde{X}) := (\rho_u^{(2)}(\widetilde{C}))_{C \in \pi_0(X)} \in \mathrm{Wh}^w(\Pi(X)) := \bigoplus_{C \in \pi_0(X)} \mathrm{Wh}^w(\pi_1(C))$$

A map $f\colon X \to Y$ of finite CW-complexes such that

$$\pi_1(f, x)\colon \pi_1(X, x) \to \pi_1(Y, f(x))$$

is injective for all $x \in X$ induces a homomorphism

$$f_*\colon \operatorname{Wh}^w(\Pi(X)) \to \operatorname{Wh}^w(\Pi(Y))$$

by

$$f_* := \big((f|_C)_*\colon \operatorname{Wh}^w(\pi_1(C)) \to \operatorname{Wh}^w(\pi_1(D))\big)_{C\in\pi_0(X)}$$

where $f(C) \subseteq D$.

The main properties of the universal L^2-torsion are collected in [**FL3**, Theorem 2.5], respectively [**FL3**, Theorem 2.11], of which we recall here the parts needed in this paper.

Lemma 2.16. *(1) Let $f\colon X \to Y$ be a G-homotopy equivalence of finite free G-CW-complexes. Suppose that X or Y is L^2-acyclic. Then both X and Y are L^2-acyclic and we get*

$$\rho_u^{(2)}(X;\mathcal{N}(G)) - \rho_u^{(2)}(Y;\mathcal{N}(G)) = \zeta(\tau(f))$$

where $\tau(f) \in \operatorname{Wh}(G)$ is the Whitehead torsion of f and

$$\zeta\colon \operatorname{Wh}(G) \to \operatorname{Wh}^w(G)$$

is the obvious homomorphism.

(2) Let

$$\begin{array}{ccc} X_0 & \longrightarrow & X_1 \\ \big\downarrow & \overset{j_0}{\searrow} & \big\downarrow j_1 \\ X_2 & \xrightarrow{\ j_2\ } & X \end{array}$$

be a pushout of finite CW-complexes such that the top horizontal map is cellular, the left arrow is an inclusion of CW-complexes, and X carries the CW-structure coming from the ones on X_i, $i = 0,1,2$. Suppose that X_i for $i = 0,1,2$ is L^2-acyclic and that for any $x_i \in X_i$ the induced homomorphism $\pi_1(X_i,x_i) \to \pi_1(X,j_i(x_i))$ is injective. Then X is L^2-acyclic and we have

$$\rho_u^{(2)}(\widetilde{X}) = (j_1)_*(\rho_u^{(2)}(\widetilde{X}_1)) + (j_2)_*(\rho_u^{(2)}(\widetilde{X}_2)) - (j_0)_*(\rho_u^{(2)}(\widetilde{X}_0))$$

(3) Let $p\colon X \to Y$ be a finite covering of finite connected CW-complexes. Let $p^*\colon \operatorname{Wh}^w(\pi_1(Y)) \to \operatorname{Wh}^w(\pi_1(X))$ be the homomorphism induced by restriction with $\pi_1(p)\colon \pi_1(X) \to \pi_1(Y)$. Then X is L^2-acyclic if and only if Y is L^2-acyclic and in this case we have

$$\rho_u^{(2)}(\widetilde{X}) = p^*(\rho_u^{(2)}(\widetilde{Y}))$$

Next we apply this invariant to the groups we are interested in.

Definition 2.17. Let G be a group with a finite model for its classifying space BG, and let $g\colon G \to G$ be a monomorphism. Let T be the mapping torus of the realisation $Bg\colon BG \to BG$. Given a factorisation $G*_g \xrightarrow{p} \Gamma \xrightarrow{q} \mathbb{Z}$ of the induced epimorphism, denote by $\overline{T} \to T$ the Γ-covering corresponding to p. Suppose that the classical Whitehead group $\operatorname{Wh}(\Gamma)$ of Γ is trivial. Then $\overline{T}$ is L^2-acyclic [**Lüc1**, Theorem 1.39], and Lemma 2.16 (1) implies that we get a well-defined invariant

$$\rho_u^{(2)}(G*_g, p) := \rho_u^{(2)}(\overline{T};\mathcal{N}(\Gamma)) \in \operatorname{Wh}^w(\Gamma)$$

which only depends on G, g and p, but not on the realisations. If $p = \operatorname{id}_G$, then we write $\rho_u^{(2)}(G*_g) = \rho_u^{(2)}(G*_g, \operatorname{id}_G)$.

A classical theorem of Waldhausen [**Wal**, Theorem 19.4] says that $\operatorname{Wh}(F_n*_g) = 0$, so that we may apply this in particular to the special case where $\Gamma = G*_g = F_n*_g$, and $p = \operatorname{id}$.

2.7. The L^2-torsion polytope. Let H be a finitely generated free-abelian group. An *(integral) polytope* in $H \otimes_{\mathbb{Z}} \mathbb{R}$ is the convex hull of a non-empty finite set of points in H (considered as a lattice inside $H \otimes_{\mathbb{Z}} \mathbb{R}$).

Given two polytopes P_1 and P_2 in $H \otimes_{\mathbb{Z}} \mathbb{R}$, their *Minkowski sum* is defined as

$$P_1 + P_2 := \{x + y \in H \otimes_{\mathbb{Z}} \mathbb{R} \mid x \in P_1, y \in P_2\}$$

It is not hard to see that the Minkowski sum is *cancellative* in the sense that $P_1 + Q = P_2 + Q$ implies $P_1 = P_2$. It turns the set of polytopes in $H \otimes_{\mathbb{Z}} \mathbb{R}$ into a commutative monoid with the one-point polytope $\{0\}$ as the identity. The *(integral) polytope group of* H, denoted by $\mathcal{P}(H)$, is defined as the Grothendieck completion of this monoid, so elements are formal differences of polytopes $P - Q$, subject to the relation

$$P - Q = P' - Q' \Longleftrightarrow P + Q' = P' + Q$$

where on the right-hand side the symbol + denotes the Minkowski sum. With motivation originating in low-dimensional topology, integral polytope groups have recently received increased attention, see [**CFF, Fun**].

We define $\mathcal{P}_T(H)$ to be the cokernel of the homomorphism $H \to \mathcal{P}(H)$ which sends h to the one-point polytope $\{h\}$. In other words, two polytopes become identified in $\mathcal{P}_T(H)$ if and only if they are related by a translation with an element of H.

For a finite set $F \subseteq H$, we denote by $P(F)$ the convex hull of F inside $H \otimes_{\mathbb{Z}} \mathbb{R}$.

Let G be a torsion free group satisfying the Atiyah Conjecture. Then as before the integral group ring $\mathbb{Z}G$ embeds into the skew field $\mathcal{D}(G)$. Let $p\colon G \to H$ be an epimorphism onto a finitely generated free-abelian group H, and denote by K the kernel of the projection p. Friedl-Lück [**FL3**, Section 3.2] define a *polytope homomorphism*

$$\mathbb{P}\colon K_1^w(\mathbb{Z}G) \to \mathcal{P}(H) \tag{4}$$

as the composition of the following maps: Firstly, apply the obvious map

$$K_1^w(\mathbb{Z}G) \to K_1(\mathcal{D}(G)), \ [f] \mapsto [\mathrm{id}_{\mathcal{D}(G)} \otimes_{\mathbb{Z}G} f] \tag{5}$$

Since $\mathcal{D}(G)$ is a skew-field, the Dieudonné determinant constructed in Section 2.2 induces a map

$$\det_{D(G)}\colon K_1(\mathcal{D}(G)) \to \mathcal{D}(G)^\times/[\mathcal{D}(G)^\times, \mathcal{D}(G)^\times] \tag{6}$$

which is in fact an isomorphism (see Silvester [**Sil**, Corollary 4.3]). Finally, we use the isomorphism (3)

$$j\colon \mathcal{D}(G) \cong T^{-1}\left(\mathcal{D}(K) * H\right) \tag{7}$$

For $x \in \mathcal{D}(K) * H$ we define $P(x) := P(\mathrm{supp}(x)) \in \mathcal{P}(H)$. It is not hard to see that for two such elements x_1, x_2 we have $P(x_1x_2) = P(x_1) + P(x_2)$. We may therefore define a homomorphism

$$P\colon \left(T^{-1}\left(\mathcal{D}(K) * H\right)\right)^\times \to \mathcal{P}(H), \ t^{-1}s \mapsto P(s) - P(t) \tag{8}$$

Since the target of P is an abelian group, the composition $P \circ j|_{\mathcal{D}(G)^\times}$ factors through the abelianisation of $\mathcal{D}(G)^\times$. The polytope homomorphism announced in (4) is induced by the maps (5), (6), (7) and (8), and it does not depend on the choices used to construct the isomorphism (7). We get an induced polytope homomorphism

$$\mathbb{P}\colon \mathrm{Wh}^w(G) \to \mathcal{P}_T(H) \tag{9}$$

If x is an element in $\mathcal{D}(G)^\times$, we will henceforth use the isomorphism j without mention and therefore denote the image of x under $P \circ j|_{\mathcal{D}(G)^\times}$ simply by $P(x)$.

In the following definition we denote by $H_1(G)_f$ the free part of the abelianisation $H_1(G)$ of a group G.

Definition 2.18. Let X be a free finite G-CW-complex. We define the *L^2-torsion polytope* $P_{L^2}(X;\mathcal{N}(G))$ of X as the image of $-\rho_u^{(2)}(X;\mathcal{N}(G))$ under the polytope homomorphism (9).

Likewise, if $g\colon G \to G$ is a monomorphism of a group G with a finite classifying space, and the obvious epimorphism $G*_g \to H_1(G*_g)_f$ factors through some $p\colon G*_g \to \Gamma$ such that Γ satisfies the Atiyah Conjecture and $\mathrm{Wh}(\Gamma) = 0$, then the *L^2-torsion polytope of g relative to p*

$$P_{L^2}(G*_g, p) \in \mathcal{P}_T(H_1(\Gamma)_f) = \mathcal{P}_T(H_1(G*_g)_f)$$

is defined as the image of $-\rho_u^{(2)}(G*_g, p)$ under $\mathbb{P}\colon \mathrm{Wh}^w(\Gamma) \to \mathcal{P}_T(H_1(\Gamma)_f)$. If $p = \mathrm{id}_G$, then we just write $P_{L^2}(G*_g)$.

We expect the L^2-torsion polytope to carry interesting information about the monomorphism g. Even for free groups we get an interesting invariant, which is new also for their automorphisms.

2.8. The Alexander polytope. The Alexander polynomial was first introduced by Alexander in [**Ale**] as a knot invariant. Its definition was later extended by McMullen [**McM**] to all finitely generated groups in the following way.

Given a finite CW-complex X with a basepoint x and $\pi_1(X) = G$, consider the covering $\pi\colon \overline{X} \to X$ corresponding to the quotient map $p\colon G \to H_1(G)_f =: H$. The *Alexander module* of X is the $\mathbb{Z}H$-module

$$A(X) = H_1(\overline{X}, \overline{x}, \mathbb{Z})$$

where $\overline{x} = \pi^{-1}(x)$.

Now let A be any finitely generated $\mathbb{Z}H$-module. Since $\mathbb{Z}H$ is Noetherian, we may pick a presentation

$$\mathbb{Z}H^r \xrightarrow{M} \mathbb{Z}H^s \to A \to 0$$

The *elementary ideal* $I(A)$ of A is the ideal generated by all $(s-1)\times(s-1)$-minors of the matrix M. The *Alexander ideal* of X is $I(A(X))$, and the *Alexander polynomial* Δ_X is defined as the greatest common divisor of the elements in $I(A(X))$. This invariant is well-defined up to multiplication by units in $\mathbb{Z}H$ and we will view it as an element in $\mathrm{Wh}^w(H) \cong T^{-1}(\mathbb{Z}H)/\{\pm h \mid h \in H\}$, where this isomorphism comes from Example 2.15. Finally, the *Alexander polytope* $P_A(X)$ is defined as the image of Δ_X under the polytope homomorphism

$$\mathbb{P}\colon \mathrm{Wh}^w(H) \to \mathcal{P}_T(H)$$

The Alexander module and hence the Alexander polynomial depend only on the fundamental group, and we define $\Delta_G := \Delta_X$ and $P_A(G) := P_A(X)$ for any space with $\pi_1(X) = G$. This applies in particular to descending HNN extensions of finitely generated groups.

We emphasise that the Alexander polynomial is accessible from a finite presentation of G: We can take X to be the presentation complex, so that the $\mathbb{Z}H$-chain complex of the pair $(\overline{X}, \overline{x})$ looks like

$$0 \to \mathbb{Z}H^r \xrightarrow{F} \mathbb{Z}H^s \to C_0(\overline{X})/C_0(\overline{x}) = 0$$

where C_0 denotes the group of zero chains and F contains the Fox derivatives associated to the given presentation (see Section 2.11). Thus $A(X)$ is the cokernel of the map F, which immediately gives a finite presentation of $A(X)$ as desired.

2.9. Seminorms on the first cohomology. Given a polytope $P \subseteq H \otimes_{\mathbb{Z}} \mathbb{R}$, we obtain a seminorm $\| \cdot \|_P$ on $\operatorname{Hom}(H, \mathbb{R}) \cong \operatorname{Hom}_{\mathbb{R}}(H \otimes_{\mathbb{Z}} \mathbb{R}, \mathbb{R})$ by putting

$$\|\varphi\|_P := \sup\{\varphi(x) - \varphi(y) \mid x, y \in P\}$$

It is clear that $\| \cdot \|_P$ remains unchanged when P is translated within $H \otimes_{\mathbb{Z}} \mathbb{R}$. Moreover, if Q is another such polytope, then we get for the Minkowski sum

$$\|\varphi\|_{P+Q} = \|\varphi\|_P + \|\varphi\|_Q$$

Thus we get a homomorphism of groups

$$\mathfrak{N} \colon \mathcal{P}_T(H) \to \operatorname{Map}(\operatorname{Hom}(H, \mathbb{R}), \mathbb{R}), \ P - Q \mapsto (\varphi \mapsto \|\varphi\|_P - \|\varphi\|_Q)$$

where $\operatorname{Map}(\operatorname{Hom}(H, \mathbb{R}), \mathbb{R})$ denotes the group of continuous maps to $\mathbb{R}$ with the pointwise addition. In general, $\mathfrak{N}(P - Q)$ does not need to be a seminorm.

The following definition is due to McMullen [**McM**].

Definition 2.19. If G is a finitely generated group, then the *Alexander norm*

$$\| \cdot \|_A \colon H^1(G; \mathbb{R}) \to \mathbb{R}$$

is defined as the image of the Alexander polytope $P_A(G)$ under $\mathfrak{N}$.

If G is the fundamental group of a compact connected orientable 3-manifold M, the first cohomology $H^1(M; \mathbb{R}) = H^1(G; \mathbb{R})$ carries another well-known seminorm $\| \cdot \|_T$, called the *Thurston seminorm*. It was first defined and examined by Thurston [**Thu**] and is closely related to the question of whether (and how) M fibres over the circle. One of the main results of [**FL3**, Theorem 3.27] is the following.

Theorem 2.20. *Let $M \neq S^1 \times D^2$ be a compact connected aspherical 3-manifold such that $\pi_1(M)$ satisfies the Atiyah Conjecture. Then the image of the L^2-torsion polytope $P_{L^2}(\widetilde{M}; \pi_1(M))$ under $\mathfrak{N}$ is the Thurston seminorm $\| \cdot \|_T$.*

Motivated by this result, we make the following definition.

Definition 2.21. Let $G = F_n *_g$ for a monomorphism $g \colon F_n \to F_n$. We call the image of the L^2-torsion polytope $P_{L^2}(G) \in \mathcal{P}_T(H_1(G)_f)$ as defined in Definition 2.18 under $\mathfrak{N}$ the *Thurston seminorm on G* and denote it by

$$\| \cdot \|_T \colon H^1(G; \mathbb{R}) \to \mathbb{R}$$

In order for this definition to make sense, we need to argue that HNN extensions of free groups satisfy the Atiyah Conjecture.

To this end, observe that G fits into the extension

$$0 \to \langle\langle F_n \rangle\rangle \to G \to \mathbb{Z} \to 0$$

By the work of Linnell (see [**Lüc1**, Theorem 10.19]), we know that the Atiyah conjecture holds for F_n, is stable under taking directed unions, and so holds for $\langle\langle F_n \rangle\rangle$, and is stable under taking extensions with elementary amenable quotients, and thus holds for G.

The proof that the terminology *seminorm* in the above definition is justified needs to be postponed to Corollary 3.5.

In [**Har2**] Harvey generalised McMullen's work and defined *higher Alexander norms*

$$\delta_k \colon H^1(G; \mathbb{R}) \to \mathbb{R}$$

for any finitely presented group G, where $\delta_0 = \| \cdot \|_A$. While we do not need the precise definition of δ_k, the following ingredient will be needed throughout the paper.

Definition 2.22. The *rational derived series*

$$G = G_r^0 \supseteq G_r^1 \supseteq G_r^2 \supseteq \cdots$$

of a group G is inductively defined with G_r^{k+1} being the kernel of the projection

$$G_r^k \to H_1(G_r^k)_f$$

Note that the quotients $\Gamma_k := G/G_r^{k+1}$ are torsion free and solvable, and so

$$\mathrm{Wh}(\Gamma_k) = 0$$

since solvable groups satisfy the K-theoretic Farrell–Jones Conjecture by a result of Wegner [**Weg**]. Moreover, Γ_k satisfies the Atiyah Conjecture by the work of Linnell (see [**Lüc1**, Theorem 10.19]). Thus, given $G = F_n *_g$, Definition 2.17 and Definition 2.18 produce an L^2-torsion polytope $P_{L^2}(G, p_k)$ for the projections

$$p_k \colon G \to \Gamma_k$$

The next result is not explicitly stated in [**FL1, FL3**], but we will indicate how it directly follows from it.

Theorem 2.23. *Let $G = F_n *_g$ be a descending HNN extension and let*

$$p_k \colon G \to \Gamma_k := G/G_r^{k+1}$$

be the obvious projection. Then the image of the L^2-torsion polytope $P_{L^2}(G, p_k)$ under $\mathfrak{N}$ is the higher Alexander norm δ_k, unless $b_1(G) = 1$ and $k = 0$.

Proof. Let $\nu_k \colon \Gamma_k \to H_1(G)_f$ be the natural projection. There is an obvious analogue of [**FL1**, Theorem 8.4] for HNN extensions of free groups which says that for $\varphi \colon H_1(G)_f \to \mathbb{Z}$ we have an equality

$$\delta_k(\varphi) = -\chi^{(2)}(T; p_k, \varphi \circ \nu_k)$$

where T denotes the mapping torus of a realisation of g. The right-hand side denotes the twisted L^2-Euler characteristic defined and examined in [**FL1**].

On the other hand, a similar argument as in the proof Theorem 2.20 (see the proof of [**FL3**, Theorem 3.27]) shows that

$$\mathfrak{N}(P_{L^2}(G, p_k))(\varphi) = \mathfrak{N}(\mathbb{P}(-\rho_u^{(2)}(G, p_k)))(\varphi) = -\chi^{(2)}(T; p_k, \varphi \circ \nu_k) \qquad \square$$

Motivated by this result, we introduce new terminology.

Definition 2.24. Let $G = F_n *_g$ be a descending HNN extension and let

$$p_k \colon G \to \Gamma_k := G/G_r^{k+1}$$

be the obvious projection. Then we call $P_{L^2}(G, p_k)$ the *higher Alexander polytopes.*

The Thurston and higher Alexander seminorms satisfy well-known inequalities for compact orientable 3-manifolds by the work of McMullen and Harvey [**McM, Har2, Har1**]. We use their characterisation in terms of polytopes to prove an analogue in the case of HNN extensions of F_2. This will be the main result of Section 4.

2.10. The Bieri–Neumann–Strebel invariant $\Sigma(G)$.

We first recall one of the definitions of the BNS-invariant $\Sigma(G)$, see [**Str**, Chapter A2.1].

Definition 2.25 (The BNS invariant)**.** Let G be a group with finite generating set S. The positive reals $\mathbb{R}_{>0}$ act on $\mathrm{Hom}(G, \mathbb{R}) \setminus \{0\}$ by multiplication. The quotient will be denoted by

$$S(G) = (\mathrm{Hom}(G, \mathbb{R}) \setminus \{0\}) / \mathbb{R}_{>0}$$

Given a class $[\varphi] \in S(G)$, let $\mathrm{Cay}(G,S)_\varphi$ denote the subgraph of the Cayley graph of G with respect to S that is induced by the vertex subset $\{g \in G \mid \varphi(g) \geqslant 0\}$. The *BNS invariant* or Σ*-invariant* is the subset

$$\Sigma(G) = \{[\varphi] \in S(G) \mid \mathrm{Cay}(G,S)_\varphi \text{ is connected}\}$$

Note that $S(G)$, with the quotient topology, is naturally homeomorphic to the unit sphere in $H^1(G;\mathbb{R})$. The invariant $\Sigma(G)$ is an open subset thereof (see [**BNS**, Theorem A]).

For rational points in $S(G)$ we have a more tangible characterisation.

Theorem 2.26 ([**BNS**, Proposition 4.3])**.** *Let $\varphi\colon G \to \mathbb{Z}$ be an epimorphism. Then $[-\varphi] \in \Sigma(G)$ if and only if G can be identified with a descending HNN-extension over a finitely generated subgroup, so that φ is the epimorphism induced by the HNN-extension.*

Definition 2.27 (Sikorav–Novikov completion)**.** Let G be a group and $\varphi \in H^1(G;\mathbb{R})$. Then the *Sikorav–Novikov completion* $\widehat{\mathbb{Z}G}_\varphi$ is defined as the set

$$\widehat{\mathbb{Z}G}_\varphi := \left\{ \sum_{g\in G} x_g \cdot g \mid \forall C \in \mathbb{R}:\ \{g \in G \mid \varphi(g) < C \text{ and } x_g \neq 0\} \text{ is finite} \right\}$$

It is easy to verify that the usual convolution turns $\widehat{\mathbb{Z}G}_\varphi$ into a ring which contains $\mathbb{Z}G$. The reason why we are interested in the Sikorav–Novikov completion is the following criterion to detect elements in the BNS-invariant.

Theorem 2.28. *Given a finitely generated group G, for a non-zero homomorphism $\varphi\colon G \to \mathbb{R}$ we have $[-\varphi] \in \Sigma(G)$ if and only if*

$$H_0(G;\widehat{\mathbb{Z}G}_\varphi) = 0 \text{ and } H_1(G;\widehat{\mathbb{Z}G}_\varphi) = 0$$

Proof. This is originally due to Sikorav [**Sik**], see also [**FT**, Theorem 4.3] for a sketch of the proof. ☐

Remark 2.29. In fact we are only discussing the *first BNS invariant*

$$\Sigma^1(G;\mathbb{Z}) = -\Sigma(G)$$

It is easily deducible from the full result of Sikorav that for descending HNN extensions of free groups the higher BNS invariants $\Sigma^n(G;\mathbb{Z})$ all coincide with $\Sigma^1(G;\mathbb{Z})$.

Definition 2.30. We define $\mu_\varphi\colon \widehat{\mathbb{Z}G}_\varphi \to \mathbb{Z}G$ in the following way: Let

$$x = \sum_{g\in G} x_g \cdot g \in \widehat{\mathbb{Z}G}_\varphi$$

and let

$$S = \big\{g \in \mathrm{supp}(x) \mid \varphi(g) = \min\{\varphi(\mathrm{supp}(x))\}\big\}$$

Then we let

$$\mu_\varphi(x) = \sum_{g\in S} x_g \cdot g$$

It is easy to see that μ_φ respects the multiplication in $\widehat{\mathbb{Z}G}_\varphi$.

The following criterion to detect units in $\widehat{\mathbb{Z}G}_\varphi$ is well-known; we include a proof here for the sake of completeness. Note that the Sikorav-Novikov completion is a domain, so being left-invertible is equivalent to being right-invertible, and so is equivalent to being a unit.

Definition 2.31. A group G is called *indicable* if it admits an epimorphism onto $\mathbb{Z}$. The group is *locally indicable* if all of its finitely generated subgroups are indicable.

Lemma 2.32. *Let G be a locally indicable group and $x \in \widehat{\mathbb{Z}G}_\varphi$. Then x is a unit in $\widehat{\mathbb{Z}G}_\varphi$ if and only if $\mu_\varphi(x)$ is of the form $\pm h$ for some $h \in G$.*

Proof. If x has an inverse $y \in \widehat{\mathbb{Z}G}_\varphi$, then

$$1 = \mu_\varphi(1) = \mu_\varphi(x)\mu_\varphi(y)$$

The latter is an equation in $\mathbb{Z}G$, where the only units are of the form $\pm h$ since G is locally indicable [**Hig**, Theorem 13].

Conversely, write $x = \sum_{g \in G} x_g \cdot g$ and write G_k for the (finite) set of elements $g \in G$ with $g \in \operatorname{supp}(x)$ and $\varphi(g) = k$. After multiplying with the unit $\mu_\varphi(x)^{-1}$, we may assume without loss of generality that $G_k = \emptyset$ for $k < 0$, $G_0 \neq \emptyset$, and $\mu_\varphi(x) = 1$, so

$$x = 1 + \sum_{g \in G_1} x_g \cdot g + \sum_{g \in G_2} x_g \cdot g + \dots$$

It is now easy to successively build a left-inverse beginning with

$$1 - \sum_{g \in G_1} x_g \cdot g + \left(\sum_{g \in G_1} x_g \cdot g \right)^2 - \sum_{g \in G_2} x_g \cdot g + \dots \qquad \square$$

Finally we verify that the above characterisation of units in $\widehat{\mathbb{Z}G}_\varphi$ is applicable for the groups of our interest.

Lemma 2.33. *Let $g\colon F_n \to F_n$ be a monomorphism. Then the associated descending HNN extension is locally indicable.*

Proof. Let $G = F_n *_g$ denote the descending HNN extension, and let ψ be the induced epimorphism to $\mathbb{Z}$.

We start by noting that G is locally indicable if and only if the normal closure of F_n inside G is, since this normal closure is the kernel of ψ, and the image of ψ is a free-abelian group, and thus locally indicable. Now, since G is a descending HNN extension, every finitely generated subgroup of $\ker \varphi$ lies in a copy of F_n, which is locally indicable. Hence G is locally indicable. $\square$

2.11. Fox calculus. In order to start computing, we introduce as a last tool Fox derivatives (defined by Fox in [**Fox**]).

Definition 2.34. Let F_n be a free group generated by $s_1, \dots, s_n$, and let w be a word in the alphabet $\{s_1, \dots, s_n\}$. We define the *Fox derivative* $\frac{\partial w}{\partial s_i} \in \mathbb{Z}G$ of w with respect to s_i inductively: we write $w = vt$ where t is one of the generators or their inverses, and v is strictly shorter than w, and set

$$\frac{\partial w}{\partial s_i} = \begin{cases} \frac{\partial v}{\partial s_i} & & t \notin \{s_i, s_i^{-1}\} \\ \frac{\partial v}{\partial s_i} + v & \text{if} & t = s_i \\ \frac{\partial v}{\partial s_i} - w & & t = s_i^{-1} \end{cases}$$

This definition readily extends first to elements $w \in F_n$, and then linearly to elements of $\mathbb{Z}F_n$, forming a map $\frac{\partial w}{\partial s_i}\colon \mathbb{Z}F_n \to \mathbb{Z}F_n$.

The following equation is known as the fundamental formula of Fox calculus [**Fox**, Formula (2.3)].

Proposition 2.35. *Let $w \in F_n$ be any word, and let $s_1, \dots, s_n$ be a generating set of F_n. Then we have*

$$\sum_{i=1}^{n} \frac{\partial w}{\partial s_i} \cdot (1 - s_i) = 1 - w$$

3. The invariants for descending HNN extensions of free groups

In this section we describe the Alexander polynomial and the universal L^2-torsion in more explicit terms for descending HNN extensions of finitely generated free groups. The computations in this chapter follow from the general properties of the invariants, but we thought it worthwhile to collect them here in order to emphasise that a close connection between the invariants should not come as a complete surprise.

Let us first observe the following.

Lemma 3.1. *Let G be a descending HNN extension $G = F_n *_g$. Pick a finite classifying space BF_n for F_n, and a realisation $Bg\colon BF_n \to BF_n$. Then the mapping torus T_{Bg} of Bg is a classifying space for G.*

Proof. It is well-known that $\pi_1(T_{Bg}) = G$. For the higher homotopy groups we observe that any map $C \to \widetilde{T_{Bg}}$ with compact domain C can be homotoped to a map whose image lies in a copy of $\widetilde{BF_n}$, which is contractible. □

3.1. First consequences. We will always view an $m \times n$-matrix A over a ring R as an R-homomorphism $R^m \to R^n$ by *right*-multiplication since we prefer working with *left*-modules.

For a monomorphism $g\colon F_n \to F_n$, let $G = F_n *_g$, and let $s_1, \ldots, s_n$ denote generators of F_n, and t the stable letter of the HNN extension. The *Fox matrix of* g is

$$F(g) = \left(\frac{\partial g(s_i)}{\partial s_j}\right)_{i,j=1}^{n} \in \mathbb{Z}F_n{}^{n\times n}$$

Put $\mathcal{S} = \{s_1, \ldots, s_n, t\}$. We will often consider the matrix

$$A(g;\mathcal{S}) = \begin{pmatrix} & s_1 - 1 \\ \operatorname{Id} - t \cdot F(g) & \vdots \\ & s_n - 1 \end{pmatrix} \in \mathbb{Z}G^{n\times(n+1)}$$

Given $s \in \mathcal{S}$, we let $A(g;\mathcal{S},s)$ be the square matrix obtained from $A(g;\mathcal{S})$ by removing the column which contains the Fox derivates with respect to s. Let $\Gamma_k = G/G_r^{k+1}$, where G_r^k are the subgroups of the rational derived series as introduced in Definition 2.22. Denote by $p_k\colon G \to \Gamma_k$ the projection and denote the ring homomorphisms $p_k\colon \mathbb{Z}G \to \mathbb{Z}\Gamma_k$ by the same letter. Notice that

$$\Gamma_0 = H_1(G)_f =: H$$

The following theorem summarises the various invariants introduced in Section 2 for descending HNN extensions of finitely generated free groups.

Theorem 3.2. *With the notation above, let $G = F_n *_g$ and $s \in \mathcal{S}$. Then:*

(1) *For the universal L^2-torsion we have*

$$\rho_u^{(2)}(G) = -[\mathbb{Z}G^n \xrightarrow{A(g;\mathcal{S},s)} \mathbb{Z}G^n] + [\mathbb{Z}G \xrightarrow{s-1} \mathbb{Z}G]$$

and so

$$P_{L^2}(G) = P(\det_{\mathcal{D}(G)}(A(g;\mathcal{S},s))) - P(s-1) \in \mathcal{P}_T(H)$$

(2) *If $p_k(s) \neq 0$, then for the universal L^2-torsion relative to p_k we have*

$$\rho_u^{(2)}(G;p_k) = -[\mathbb{Z}\Gamma_k^n \xrightarrow{p_k(A(g;\mathcal{S},s))} \mathbb{Z}\Gamma_k^n] + [\mathbb{Z}\Gamma_k \xrightarrow{p_k(s)-1} \mathbb{Z}\Gamma_k]$$

and so

$$P_{L^2}(G,p_k) = P(\det_{\mathcal{D}(\Gamma_k)}(p_k(A(g;\mathcal{S},s)))) - P(p_k(s)-1) \in \mathcal{P}_T(H)$$

(3) In $\mathrm{Wh}^w(H) \cong (T^{-1}\mathbb{Q}H)^\times/\{\pm h \mid h \in H\}$ *we have*

$$\Delta_A(G) = \begin{cases} -\rho_u^{(2)}(G;p_0) & \text{if } b_1(G) \geqslant 2 \\ -\rho_u^{(2)}(G;p_0)\cdot(p_0(t)-1) & \text{if } b_1(G) = 1 \end{cases}$$

(4) Let $\varphi \in \mathrm{Hom}(G,\mathbb{R})$*. If* $\varphi(s) \neq 0$*, then* $[-\varphi] \in \Sigma(G)$ *if and only if the map*

$$A(g;\mathcal{S},s)\colon \widehat{\mathbb{Z}G}_\varphi^n \to \widehat{\mathbb{Z}G}_\varphi^n$$

is surjective, or equivalently, bijective.

Proof. (1) We write the relations defining the descending HNN extension $G = F_n *_g$ as

$$R_i = s_i t g(s_i)^{-1} t^{-1}$$

If we let BF_n be the wedge of n circles, then the $\mathbb{Z}G$-chain complex of the mapping torus T_{Bg} has the form

$$C_* = 0 \to \mathbb{Z}G^n \xrightarrow{c_2} \mathbb{Z}G^{n+1} \xrightarrow{c_1} \mathbb{Z}G \to 0$$

where c_1 is given by the transpose of

$$\begin{pmatrix} s_1 - 1 & s_2 - 1 & \ldots & s_n - 1 & t-1 \end{pmatrix}$$

and c_2 is given by the $n \times (n+1)$ matrix containing the Fox derivatives $\frac{\partial R_i}{\partial s_j}$ and $\frac{\partial R_i}{\partial t}$. This is precisely the matrix $A(g;\mathcal{S})$ since

$$\begin{aligned} \frac{\partial R_i}{\partial s_j} &= \delta_{ij} + s_i t\left(\frac{\partial g(s_i)^{-1}}{\partial s_j} + g(s_i)^{-1}\cdot\frac{\partial t^{-1}}{\partial s_j}\right) \\ &= \delta_{ij} - s_i t g(s_i)^{-1}\cdot\frac{\partial g(s_i)}{\partial s_j} \\ &= \delta_{ij} - t\cdot\frac{\partial g(s_i)}{\partial s_j} \\ \frac{\partial R_i}{\partial t} &= s_i - s_i t g(s_j)^{-1} t^{-1} = s_i - 1 \end{aligned}$$

where δ_{ij} denotes the Kronecker delta.

Consider the $\mathbb{Z}G$-chain complexes

$$B_* = \quad 0 \longrightarrow 0 \longrightarrow \mathbb{Z}G \xrightarrow{s-1} \mathbb{Z}G \longrightarrow 0$$

$$D_* = \quad 0 \longrightarrow \mathbb{Z}G^n \xrightarrow{A(g;\mathcal{S},s)} \mathbb{Z}G^n \longrightarrow 0 \longrightarrow 0$$

We obtain a short exact sequence of $\mathbb{Z}G$-chain complexes

$$0 \to B_* \to C_* \to D_* \to 0$$

Since B_* is L^2-acyclic by [**Lüc1**, Theorem 3.14 (6) on page 129 and (3.23) on page 136], D_* is also L^2-acyclic and we have the sum formula [**FL3**, Lemma 1.9]

$$\begin{aligned} \rho_u^{(2)}(G) &= \rho_u^{(2)}(C_*) \\ &= \rho_u^{(2)}(B_*) + \rho_u^{(2)}(D_*) \\ &= [\mathbb{Z}G \xrightarrow{s-1} \mathbb{Z}G] - [\mathbb{Z}G^n \xrightarrow{A(g;\mathcal{S},s)} \mathbb{Z}G^n] \end{aligned}$$

The statement

$$P_{L^2}(G) = P(\det_{\mathcal{D}(G)}(A(g;\mathcal{S},s))) - P(s-1) \in \mathcal{P}_T(H)$$

is obtained by applying the polytope homomorphism $\mathbb{P}\colon \mathrm{Wh}^w(G) \to \mathcal{P}_T(G)$.

(2) This follows exactly as (1) since the chain complex used to define $\rho_u^{(2)}(G; p_k)$ is

$$0 \to \mathbb{Z}\Gamma_k^n \xrightarrow{p_k(c_2)} \mathbb{Z}\Gamma_k^{n+1} \xrightarrow{p_k(c_1)} \mathbb{Z}\Gamma_k \to 0$$

(3) A $\mathbb{Z}H$-presentation of the Alexander module $A(G)$ is given by

$$\mathbb{Z}H^n \xrightarrow{p_0(A(g;\mathcal{S}))} \mathbb{Z}H^{n+1} \to A(G) \to 0$$

We now apply the same argument as in the proof of [**McM**, Theorem 5.1]: If $b_1(G) \geqslant 2$, then this yields

$$\det(p_0(A(g; \mathcal{S}, s))) = (p_0(s) - 1) \cdot \Delta_A(G)$$

for all $s \in \mathcal{S}$ such that $p_0(s) \neq 0$.

If $b_1(G) = 1$, then

$$\det(p_0(A(g; \mathcal{S}, t))) = \Delta_A(G)$$

Since the isomorphism $\mathrm{Wh}^w(G) \cong T^{-1}(\mathbb{Z}H)$ is given by the determinant over $T^{-1}(\mathbb{Z}H)$, the claim follows from part (2) for $k = 0$ (since $\Gamma_0 = H$).

(4) By Theorem 2.28, $[-\varphi] \in \Sigma(G)$ if and only if

$$H_0(G; \widehat{\mathbb{Z}G}_\varphi) = 0 \text{ and } H_1(G; \widehat{\mathbb{Z}G}_\varphi) = 0$$

The chain complex computing these homology groups is

$$0 \to \widehat{\mathbb{Z}G}_\varphi^n \xrightarrow{c_2} \widehat{\mathbb{Z}G}_\varphi^{n+1} \xrightarrow{c_1} \widehat{\mathbb{Z}G}_\varphi \to 0$$

We assume $\varphi(s) \neq 0$ for a fixed $s \in \mathcal{S}$. Since G is locally indicable (by Lemma 2.33), Lemma 2.32 shows that $s-1$ is invertible in $\widehat{\mathbb{Z}G}_\varphi$, which implies that c_1 is surjective, and therefore $H_0(G; \widehat{\mathbb{Z}G}_\varphi) = 0$ for any non-zero φ.

Assume without loss of generality that $s = s_1$. Then the kernel of d_1 is the set

$$K = \left\{ (x_1, \ldots, x_{n+1}) \in \widehat{\mathbb{Z}G}_\varphi^{n+1} \mid \sum_{k=2}^{n+1} x_k (s_k - 1)(s_1 - 1)^{-1} = -x_1 \right\}$$

By forgetting the first coordinate we see that K is $\widehat{\mathbb{Z}G}_\varphi$-isomorphic to $\widehat{\mathbb{Z}G}_\varphi^n$, and

$$H_1(G; \widehat{\mathbb{Z}G}_\varphi) = 0$$

is equivalent to

$$A(g; \mathcal{S}, s) \colon \widehat{\mathbb{Z}G}_\varphi^n \to \widehat{\mathbb{Z}G}_\varphi^n$$

being surjective.

Since $\widehat{\mathbb{Z}G}_\varphi$ is stably finite (this was shown by Kochloukova [**Koc**]), an epimorphism $\widehat{\mathbb{Z}G}_\varphi^n \to \widehat{\mathbb{Z}G}_\varphi^n$ is necessarily an isomorphism. ☐

Remark 3.3. Note that the above proof shows (and uses) that $A(g; \mathcal{S}, s)$ (resp. $p_k(A(g; \mathcal{S}, s))$ is invertible over $\mathcal{D}(G)$ (resp. $\mathcal{D}(\Gamma_k)$). We will henceforth call a $\mathbb{Z}G$-square matrix with this property *non-degenerate*.

Example 3.4. Using part (1) of the above theorem we compute the L^2-torsion polytope in a few examples. We use $a, b, c, \ldots$ to denote some fixed generators of F_n.

(1) For arbitrary n and $g = \mathrm{id}$ the polytope is just a line of length $n - 1$ between 0 and t^{n-1}.
(2) For $g \colon F_2 \to F_2$, $x \mapsto a^k x a^{-k}$ for some $k \in \mathbb{Z}$, we get a tilted line between 0 and $a^k t$.
(3) For $g \colon F_3 \to F_3$, $a \mapsto b$, $b \mapsto c$, $c \mapsto a[b, c]$ we get a triangle as shown below.

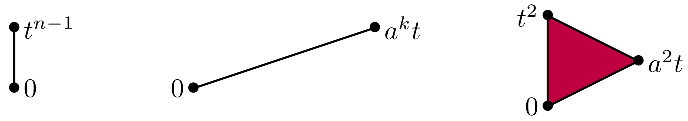

FIGURE 3.1. The L^2-torsion polytopes in Example 3.4

More importantly, we can now show that the L^2-torsion polytope of free group HNN extensions induces indeed a seminorm on the first cohomology.

Corollary 3.5. *Let $G = F_n *_g$. Then the Thurston seminorm*

$$\| \cdot \|_T \colon H^1(G;\mathbb{R}) \to \mathbb{R}$$

as defined in Definition 2.21 is indeed a seminorm.

Proof. As a difference of seminorms it is clear that $\| \cdot \|_T$ is $\mathbb{R}$-linear and continuous.

First let $\varphi \in H^1(G;\mathbb{Q})$ be a rational class. We easily find a generating set $s_1, \ldots, s_n$ of F_n such that $\varphi(s_1) = 0$. We add a stable letter to this set, and form a generating set $\mathcal{S}$ for G.

We get from the previous theorem

$$\rho_u^{(2)}(G) = -[\mathbb{Z}G^n \xrightarrow{A(g;\mathcal{S},s_1)} \mathbb{Z}G^n] + [\mathbb{Z}G \xrightarrow{s_1 - 1} \mathbb{Z}G]$$

By [**FH**, Theorem 2.2] of Friedl–Harvey applied to $\mathbb{K} = \mathcal{D}(K)$, the polytope $P(\det_{\mathcal{D}(G)}(A(g;\mathcal{S},s_1)))$ defines a seminorm on $H^1(G;\mathbb{R})$ which we denote by $\| \cdot \|_{T'}$. Then, since $\varphi(s_1) = 0$, we have

$$\|\varphi\|_T = \|\varphi\|_{T'} \geqslant 0$$

and for any $\psi \in H^1(G;\mathbb{R})$

$$\begin{aligned} \|\varphi + \psi\|_T &= \|\varphi + \psi\|_{T'} - |(\varphi + \psi)(s_1)| \\ &\leqslant \|\varphi\|_{T'} + \|\psi\|_{T'} - |\psi(s_1)| \\ &= \|\varphi\|_T + \|\psi\|_T \end{aligned}$$

This finishes the proof for rational classes.

The general case directly follows by the continuity of $\| \cdot \|_T$. □

3.2. The Determinant Comparison Problem. We borrow the following partial order on $\mathcal{P}_T(H)$ from Friedl–Tillmann [**FT**]: If $P - Q, P' - Q' \in \mathcal{P}_T(H)$, then we say that

$$P - Q \leqslant P' - Q'$$

if $P + Q' \subseteq P' + Q$ holds up to translation. If this is the case, then the norm map

$$\mathfrak{N} \colon \mathcal{P}_T(H) \to \operatorname{Map}(\operatorname{Hom}(H,\mathbb{R}),\mathbb{R})$$

clearly satisfies

$$\mathfrak{N}(P - Q)(\varphi) \leqslant \mathfrak{N}(P' - Q')(\varphi)$$

for all $\varphi \in \operatorname{Hom}(H,\mathbb{R})$.

Upon comparing parts (1), (2), and (3) of Theorem 3.2 and motivated by McMullen and Harvey's inequalities (see [**McM**, Theorem 1.1], [**Har2**, Theorem 10.1] and [**Har1**, Corollary 2.10]), we are led to the following question.

Question 3.6 (Determinant Comparison Problem)**.** *Let $G \xrightarrow{\mu} \overline{G} \xrightarrow{\nu} H$ be epimorphisms of finitely generated torsion-free groups G and $\overline{G}$ and a free-abelian group H. Assume that G and $\overline{G}$ satisfy the Atiyah Conjecture. Let A be an $n \times n$-matrix over $\mathbb{Z}G$ that becomes invertible over $\mathcal{D}(G)$ such that its image $\mu(A)$ becomes invertible over $\mathcal{D}(\overline{G})$. Consider the polytope homomorphism*

$$\mathbb{P}\colon K_1^w(\mathbb{Z}G) \to \mathcal{P}_T(H)$$

and likewise for $\overline{G}$. Is the inequality

$$\mathbb{P}([\mu(A)\colon \mathbb{Z}\overline{G}^n \to \mathbb{Z}\overline{G}^n]) \leqslant \mathbb{P}([A\colon \mathbb{Z}G^n \to \mathbb{Z}G^n])$$

satisfied in $\mathcal{P}_T(H)$?

We record the following consequence.

Lemma 3.7. *If Question 3.6 is true for a descending HNN extension $G = F_n *_g$ with stable letter t, then*

$$\delta_k(\varphi) \leq \|\varphi\|_T$$

for all $\varphi \in H^1(G;\mathbb{R})$, unless $k = 0$ and $b_1(G) = 1$. In this latter case,

$$\delta_0(\varphi) = \|\varphi\|_A \leq \|\varphi\|_T + |\varphi(t)|$$

Proof. This follows directly from Theorem 3.2. □

The following is an elementary observation.

Lemma 3.8. *If $n = 1$, then Question 3.6 is true.*

The difficulty in answering Question 3.6 comes from the fact that polytopes are hard to control when adding elements in $\mathcal{D}(G)$, but this invariably happens when calculating the Dieudonné determinant. In this case it is sometimes easier to work with one cohomology class $\varphi\colon G \to \mathbb{Z}$ at a time, rather than taking the maximal free-abelian quotient $p\colon G \to H_1(G)_f$. Because of this, it is useful to note that we can weaken the assumption of Lemma 3.7.

Proposition 3.9. *If $G = F_n *_g$ is a descending HNN extension and Question 3.6 is true whenever $H = \mathbb{Z}$, then the conclusion of Lemma 3.7 holds.*

Proof. Given epimorphisms $G \xrightarrow{p} H \xrightarrow{\nu} H' \xrightarrow{\psi} \mathbb{Z}$, where H and H' are finitely generated free-abelian, denote by

$$\mathbb{P}_p\colon \operatorname{Wh}^w(G) \to \mathcal{P}_T(H)$$
$$\mathbb{P}_{\nu\circ p}\colon \operatorname{Wh}^w(G) \to \mathcal{P}_T(H')$$

the polytope homomorphisms associated to p and $\nu \circ p$. Then [**FL1**, Lemma 6.12] states in our notation that for any element $x \in \operatorname{Wh}^w(G)$ we have

$$\mathfrak{N}(\mathbb{P}_p(x))(\psi \circ \nu) = \mathfrak{N}(\mathbb{P}_{\nu\circ p}(x))(\psi) \tag{10}$$

We apply this to the epimorphisms

$$G \xrightarrow{p} H_1(G)_f \xrightarrow{\varphi} \mathbb{Z} \xrightarrow{\mathrm{id}} \mathbb{Z}$$

Let $A = A(g;\mathcal{S},t)$ and let $x \in \operatorname{Wh}^w(G)$ be given by

$$x = [\mathbb{Z}G^n \xrightarrow{A} \mathbb{Z}G^n] - [\mathbb{Z}G \xrightarrow{t-1} \mathbb{Z}G]$$

By Theorem 3.2, (1) and the definition of the Thurston norm (see Definition 2.21), the equality (10) becomes

$$\begin{aligned} \|\varphi\|_T &= \mathfrak{N}(\mathbb{P}_p(x))(\varphi) \\ &= \mathfrak{N}(\mathbb{P}_{\varphi\circ p}(x))(\mathrm{id}) \\ &= \mathfrak{N}(\mathbb{P}_{\varphi\circ p}(A))(\mathrm{id}) - \mathfrak{N}(\mathbb{P}_{\varphi\circ p}(t-1))(\mathrm{id}) \\ &= f(\mathbb{P}_{\varphi\circ p}(A)) - |\varphi(t)| \end{aligned} \tag{11}$$

where $f\colon \mathcal{P}_T(\mathbb{Z}) \xrightarrow{\cong} \mathbb{Z}$ denotes the isomorphism given by mapping an interval $[m,n] \subseteq \mathbb{R}$ with $m,n \in \mathbb{Z}$ to $n-m$. By the same arguments we get the equality

$$\delta_k(\varphi) = f(\mathbb{P}_{\varphi\circ p}(p_k(A))) - |\varphi(t)| \tag{12}$$

from Theorem 2.23, unless $b_1(G) = 1$ and $k = 0$. In this latter case, we get from Theorem 3.2, (3)

$$\delta_0(\varphi) = \|\varphi\|_A = f(\mathbb{P}_{\varphi\circ p}(p_0(A))) \tag{13}$$

The assumption that Question 3.6 is true whenever $H = \mathbb{Z}$ implies

$$\mathbb{P}_{\varphi\circ p}(p_k(A)) \leqslant \mathbb{P}_{\varphi\circ p}(A)$$

for all $k \geqslant 0$. Since $f\colon \mathcal{P}_T(\mathbb{Z}) \xrightarrow{\cong} \mathbb{Z}$ is order-preserving when $\mathbb{Z}$ is equipped with the usual order, we obtain the desired inequalities from (11), (12), and (13). □

Remark 3.10. Under the isomorphism $\mathcal{D}(G) \cong T^{-1}(\mathcal{D}(\ker\varphi) * \mathbb{Z})$ induced by

$$\varphi\colon G \to \mathbb{Z}$$

it is proved in [**FL1**, Lemma 6.16] that $\det_{\mathcal{D}(G)}(A)$ already lives over $\mathcal{D}(\ker\varphi) * \mathbb{Z}$ (this is in essence an application of Euclid's algorithm). Then $\mathfrak{N}(\mathbb{P}_{\varphi\circ p}(A))(\mathrm{id})$ is precisely the degree of $\det_{\mathcal{D}(G)}(A)$ as a (twisted) Laurent polynomial. The same comment holds for $\det_{\mathcal{D}(\Gamma_k)}(p_k(A))$ and $\mathcal{D}(\Gamma_k)$.

4. Thurston, Alexander and higher Alexander norms

In this section we will circumvent the Determinant Comparison Problem to prove the inequalities announced in Lemma 3.7 for descending HNN extensions of F_2. As before, we denote by $\Gamma_k = G/G_r^{k+1}$ the quotient of the rational derived series and the natural projections by $p_k\colon G \to \Gamma_k$. We also write $H = \Gamma_0 = H_1(G)_f$.

Definition 4.1. Let $\leqslant$ be a biorder on H. For every $k \geqslant 0$ we let K_k be the kernel of the projection $\Gamma_k \to H$. We define maps

$$\mu_{\leqslant}\colon (\mathcal{D}(K_k) * H) \smallsetminus \{0\} \to (\mathcal{D}(K_k) * H)^\times$$

by

$$\sum_{h\in H} \lambda_h \cdot h \mapsto \lambda_{h_0} \cdot h_0$$

where h_0 is the $\leqslant$-minimal element in the support of $\sum \lambda_h \cdot h$.

It is easy to see that $\mu_{\leqslant}$ is multiplicative and so extends to a group homomorphism on the Ore localisation

$$\mu_{\leqslant}\colon \mathcal{D}(\Gamma_k)^\times = T^{-1}(\mathcal{D}(K_k) * H)^\times \to (\mathcal{D}(K_k) * H)^\times$$

Since the units in $\mathcal{D}(K_k) * H$ are precisely those elements whose support is a singleton, there is a canonical group homomorphism

$$\mathrm{supp}\colon (\mathcal{D}(K_k) * H)^\times \to H$$

As H is abelian, the composition

$$\mathrm{supp}\,\mu_{\leq}\colon \mathcal{D}(\Gamma_k)^\times \to H$$

factors through the abelianisation of the source to give a map denoted by the same name

$$\operatorname{supp}\mu_{\leq}\colon \mathcal{D}(\Gamma_k)^{\times}_{\text{ab}} \to H$$

Similarly, we can define a map

$$\operatorname{supp}\mu_{\leq}\colon \mathcal{D}(G)^{\times}_{\text{ab}} \to H$$

For the behaviour of $\mu_{\leqslant}$ under addition we have the following.

Lemma 4.2. *Let $x, y \in \mathcal{D}(G)^{\times}$.*

(1) If $\operatorname{supp}\mu_{\leqslant}(x) < \operatorname{supp}\mu_{\leqslant}(y)$, then

$$\mu_{\leqslant}(x+y) = \mu_{\leqslant}(x)$$

(2) If $\operatorname{supp}\mu_{\leqslant}(x) = \operatorname{supp}\mu_{\leqslant}(y)$ and $\mu_{\leqslant}(x) \neq \mu_{\leqslant}(y)$, then

$$\mu_{\leqslant}(x-y) = \mu_{\leqslant}(x) - \mu_{\leqslant}(y)$$

(3) If $\mu_{\leqslant}(x) = \mu_{\leqslant}(y)$, then

$$\operatorname{supp}\mu_{\leqslant}(x-y) > \operatorname{supp}\mu_{\leqslant}(x) = \operatorname{supp}\mu_{\leqslant}(y)$$

The same statements hold for $x, y \in \mathcal{D}(\Gamma_k)^{\times}$, $k \geqslant 0$.

Proof. Each of the claims is obvious if both x and y lie in the subring $\mathcal{D}(K) * H$.

For the general case, write $x = t^{-1}s$, $y = v^{-1}u$ with $s, t, u, v \in \mathcal{D}(K) * H$. Write $d^{-1}c = tv^{-1}$ for some $c, d \in \mathcal{D}(K) * H$. Then

$$x + y = t^{-1}(s + tv^{-1}u) = t^{-1}d^{-1}(ds + cu)$$

and for the first claim it thus suffices to prove

$$\mu_{\leqslant}(ds + cu) = \mu_{\leqslant}(ds) \tag{14}$$

But by assumption we have

$$\mu_{\leqslant}(ds) = \mu_{\leqslant}(cvt^{-1}s) = \mu_{\leqslant}(cv)\mu_{\leqslant}(x)$$

and

$$\mu_{\leqslant}(cv)\mu_{\leqslant}(y) = \mu_{\leqslant}(cv)\mu_{\leqslant}(v^{-1}u) = \mu_{\leqslant}(cu)$$

and so the first observation in this proof is applicable and yields (14).

The other claims follow in precisely the same way. □

Recall that we have introduced a non-degeneration condition in Remark 3.3. Under this assumption the following definition is meaningful.

Definition 4.3 (Well-behaved matrices)**.** Let $\leqslant$ be a biorder on H. A non-degenerate square matrix A over $\mathbb{Z}G$ is *well behaved* with respect to $\leqslant$ if for every $k \geqslant 0$

$$\operatorname{supp}\mu_{\leqslant}(\det_{\mathcal{D}(\mathrm{G})} A) \leqslant \operatorname{supp}\mu_{\leqslant}(\det_{\mathcal{D}(\Gamma_k)} p_k(A))$$

If $\operatorname{supp}\mu_{\leqslant}(\det_{\mathcal{D}(\mathrm{G})} A) = \operatorname{supp}\mu_{\leqslant}(\det p_k(A))$, we say that A is *very well behaved.*

Lemma 4.4. *The product of two well-behaved matrices is itself well-behaved. Also, a matrix is well-behaved if and only if it is so after being multiplied on either side by a very well-behaved matrix.*

Proof. This follows immediately from the observations that the Dieudonné determinant and $\mu_{\leqslant}$ are multiplicative, and $\leqslant$ is multiplication invariant. □

Lemma 4.5. *Let A be a non-degenerate 2×2 matrix over $\mathbb{Z}G$. Then A is well-behaved provided that $\det_{\mathcal{D}(\mathrm{G})}\mu_{\leqslant}(A) \neq 0$.*

Proof. Let

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

and let us fix a k.

Since A is non-degenerate, it contains at least one entry which does not become zero after applying p_k; without loss of generality let us suppose that d is such an entry.

We have

$$\det{}^{\mathrm{c}}_{\mathcal{D}(\mathrm{G})} A = ad - bd^{-1}cd$$

Note that $\operatorname{supp} \mu_{\leqslant}(bd^{-1}cd) = \operatorname{supp} \mu_{\leqslant}(bc)$ and likewise after applying p_k (we are also allowing empty supports here).

We need to consider three cases. If $\operatorname{supp} \mu_{\leqslant}(ad) < \operatorname{supp} \mu_{\leqslant}(bc)$, then by Lemma 4.2

$$\mu_{\leqslant}(\det{}^{\mathrm{c}}_{\mathcal{D}(\mathrm{G})} A) = \mu_{\leqslant}(ad - bd^{-1}cd) = \mu_{\leqslant}(ad)$$

But then

$$\operatorname{supp} \mu_{\leqslant}(ad) \leqslant \operatorname{supp} \mu_{\leqslant}(p_k(ad))$$

and

$$\operatorname{supp} \mu_{\leqslant}(ad) \leq \operatorname{supp} \mu_{\leqslant}(p_k(bc))$$

and thus

$$\operatorname{supp} \mu_{\leqslant}(ad) \leqslant \operatorname{supp} \mu_{\leqslant}(\det_{\mathcal{D}(\Gamma_k)} p_k(A))$$

The case $\operatorname{supp} \mu_{\leqslant}(bc) < \operatorname{supp} \mu_{\leqslant}(ad)$ is analogous.

Now let us suppose that $\operatorname{supp} \mu_{\leqslant}(bc) = \operatorname{supp} \mu_{\leqslant}(ad)$. By assumption we have

$$\det{}^{\mathrm{c}}_{\mathcal{D}(\mathrm{G})} \mu_{\leqslant}(A) = \mu_{\leqslant}(ad) - \mu_{\leqslant}(bd^{-1}cd) \neq 0$$

and so by the second part of Lemma 4.2

$$\mu_{\leqslant}(\det{}^{\mathrm{c}}_{\mathcal{D}(\mathrm{G})} A) = \mu_{\leqslant}(ad - bd^{-1}cd) = \mu_{\leqslant}(ad) - \mu_{\leqslant}(bd^{-1}cd)$$

Hence

$$\operatorname{supp} \mu_{\leqslant}(\det_{\mathcal{D}(\mathrm{G})} A) = \operatorname{supp} \mu_{\leqslant}(ad)$$

But as before

$$\operatorname{supp} \mu_{\leqslant}(ad) \leqslant \operatorname{supp} \det p_k(A)$$

which completes the proof. □

Lemma 4.6. *Let $G = F_2 *_g$ be a descending HNN extension with generating set $\mathcal{S} = \{x, y, t\}$. Let $\leqslant$ be a biorder on $H_1(G)_f$ and suppose that $p(y), p(x) > 0$ or $p(y), p(x) < 0$. Then for every $s \in \mathcal{S}$ the matrix $A(g; \mathcal{S}, s)$ as defined in Section 3.1 is well-behaved with respect to $\leqslant$.*

Proof. Let $v = g(x)$ and $w = g(y)$. Consider

$$B_y = A(g; \mathcal{S}, t) \cdot \begin{pmatrix} 1-x & 0 \\ 0 & 1-y \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

Note that B_y is well-behaved if and only if so is $A(g; \mathcal{S}, t)$ since both of the matrices on the right are very well-behaved; for the middle matrix we are using the fact that

$$p(x) \neq 0 \neq p(y)$$

Using Proposition 2.35 we compute

$$\begin{aligned} B_y &= \begin{pmatrix} 1-x-t(1-v) & -t\frac{\partial v}{\partial y}(1-y) \\ 1-y-t(1-w) & 1-y-t\frac{\partial w}{\partial y}(1-y) \end{pmatrix} \\ &= \begin{pmatrix} 1-x-(1-x)t & -t\frac{\partial v}{\partial y}(1-y) \\ 1-y-(1-y)t & 1-y-t\frac{\partial w}{\partial y}(1-y) \end{pmatrix} \\ &= \begin{pmatrix} (1-x)(1-t) & -t\frac{\partial v}{\partial y}(1-y) \\ (1-y)(1-t) & 1-y-t\frac{\partial w}{\partial y}(1-y) \end{pmatrix} \end{aligned}$$

Now B_y is clearly a product of a very well-behaved matrix and the matrix

$$A(g;\mathcal{S},x) = \begin{pmatrix} -t\frac{\partial v}{\partial y} & 1-x \\ 1-t\frac{\partial w}{\partial y} & 1-y \end{pmatrix}$$

We form B_x in the analogous manner:

$$B_x = A(g;\mathcal{S},t)\cdot\begin{pmatrix} 1-x & 0 \\ 0 & 1-y \end{pmatrix}\cdot\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

and observe as before that it is a product of a very well-behaved matrix and the matrix

$$A(g;\mathcal{S},y) = \begin{pmatrix} 1-t\frac{\partial v}{\partial x} & 1-x \\ -t\frac{\partial w}{\partial x} & 1-y \end{pmatrix}$$

From all this we see that one of $A(g;\mathcal{S},x), A(g;\mathcal{S},y), A(g;\mathcal{S},t)$ is well-behaved if and only if the others are.

We now show that the matrices $A(g;\mathcal{S},x), A(g;\mathcal{S},y), A(g;\mathcal{S},t)$ are indeed well-behaved for any monomorphism $h\colon F_2 \to F_2$ by induction. Depending on whether $p(x), p(y)$ are both positive or both negative, we need to consider two cases.

Case 1: $p(y), p(x) > 0$.

The induction in this case is over n, the length of the maximal common prefix of v and w.

Suppose first that $n = 0$. If

$$\mu_{\leqslant}\left(1 - t\frac{\partial w}{\partial y}\right) = 1$$

then $\det_{\mathcal{D}(\mathrm{G})} \mu_{\leqslant}(A(g;\mathcal{S},x)) \neq 0$, and we are already done by Lemma 4.5.

Otherwise, we have

$$\mu_{\leqslant}(A(g;\mathcal{S},x)) = \begin{pmatrix} -t\mu_{\leqslant}(\frac{\partial v}{\partial y}) & 1 \\ -t\mu_{\leqslant}(\frac{\partial w}{\partial y}) & 1 \end{pmatrix}$$

The determinant of this matrix is trivial if and only if

$$\mu_{\leqslant}\left(\frac{\partial v}{\partial y}\right) = \mu_{\leqslant}\left(\frac{\partial w}{\partial y}\right) \tag{15}$$

Let us assume that this holds. Since v and w have no common prefix, and all elements in the support of $\frac{\partial v}{\partial y}$ are either trivial or a prefix of v (and similarly for w), we must have

$$1 = \mu_{\leqslant}\left(\frac{\partial v}{\partial y}\right) = \mu_{\leqslant}\left(\frac{\partial w}{\partial y}\right)$$

But $1 \in \operatorname{supp}\frac{\partial v}{\partial y}$ implies that v begins with y. The analogous statement holds for w, and so v and w have a common prefix. This is the desired contradiction, so (15) cannot happen. Hence $\det^{\mathrm{c}}_{\mathcal{D}(\mathrm{G})} \mu_{\leqslant}(A(g;\mathcal{S},x)) \neq 0$, and we are again done by Lemma 4.5.

For the induction step, let us suppose that v and w have a common prefix. We let l be the first letter of v and w; without loss of generality let us assume that $l \in \{y^{\pm 1}\}$.

Let $s = tl \in G$, and also write $v' = s^{-1}xs = l^{-1}vl$ and $w' = s^{-1}ys = l^{-1}wl$. Then

$$A(g;\mathcal{S},x) = \begin{pmatrix} -sl^{-1}\frac{\partial l v' l^{-1}}{\partial y} & 1-x \\ 1 - sl^{-1}\frac{\partial l w' l^{-1}}{\partial y} & 1-y \end{pmatrix}$$
$$= \begin{pmatrix} -sl^{-1}l\frac{\partial v' l^{-1}}{\partial y} & 1-x \\ 1 - sl^{-1}l\frac{\partial w' l^{-1}}{\partial y} & 1-y \end{pmatrix}$$
$$= \begin{pmatrix} -s\frac{\partial v'}{\partial y} & 1-x \\ 1 - s\frac{\partial w'}{\partial y} & 1-y \end{pmatrix}$$

If we let $g' \colon F_2 \to F_2$ be g followed by conjugation with l, then G is isomorphic to the HNN extensions $F_2 *_{g'}$ with stable letter s. From the above calculation we see for $\mathcal{S}' = \mathcal{S} \cup \{s\} \smallsetminus \{t\}$ that

$$A(g';\mathcal{S}',x) = A(g;\mathcal{S},x)$$

Note that when $l \in \{x^{\pm 1}\}$, we show that $A(g';\mathcal{S}',y) = A(g;\mathcal{S},y)$.

The claim now follows from the induction hypothesis since $g'(x)$ and $g'(y)$ have a shorter common prefix, and from the fact that $A(g;\mathcal{S},x)$ being well behaved implies the same for $A(g;\mathcal{S},y)$.

Case 2: $p(y), p(x) < 0$.

This case is completely analogous, except now we induct on the length of the maximal common suffix. Let us look at the base case of the induction. Recall that

$$A(g;\mathcal{S},x) = \begin{pmatrix} -t\frac{\partial v}{\partial y} & 1-x \\ 1 - t\frac{\partial w}{\partial y} & 1-y \end{pmatrix}$$

Assuming $\det^{c}_{\mathcal{D}(G)} \mu_{\leqslant}(A(g;\mathcal{S},x)) = 0$ immediately yields

$$x^{-1}\mu_{\leqslant}\left(t\frac{\partial v}{\partial y}\right) = y^{-1}\mu_{\leqslant}\left(t\frac{\partial w}{\partial y}\right)$$

since now we have $\mu_{\leqslant}(1-y) = -y$ and $\mu_{\leqslant}(1-x) = -x$. The equation is equivalent to

$$\mu_{\leqslant}\left(v^{-1}\frac{\partial v}{\partial y}\right) = \mu_{\leqslant}\left(w^{-1}\frac{\partial w}{\partial y}\right)$$

which implies, as w and v have no common suffix, that

$$1 = \mu_{\leqslant}\left(v^{-1}\frac{\partial v}{\partial y}\right) = \mu_{\leqslant}\left(w^{-1}\frac{\partial w}{\partial y}\right)$$

This in turn implies that both v and w end with y^{-1}, which contradicts the lack of common suffix.

Now suppose that w and v do indeed have a common suffix. We let l denote the last letter of w and v, declare $s = tl^{-1}$, and proceed exactly as in the first case. □

The following is an analogue of McMullen's [**McM**, Theorem 1.1], and more generally Harvey's [**Har2**, Theorem 10.1], for the newly defined Thurston norm of descending HNN extensions of free groups.

Theorem 4.7. *Let G be a descending HNN extension of F_2 with $b_1(G) \geqslant 2$. Then we have for the Thurston and higher-order Alexander semi-norms on $H^1(G;\mathbb{R})$ the inequality*

$$\delta_k(\varphi) \leqslant \|\varphi\|_T$$

for every $\varphi \in H^1(G;\mathbb{R})$ *and* $k \geqslant 0$.

Proof. Let $\varphi \in H^1(G;\mathbb{R})$ be a non-trivial class. There exist two biorders on $H_1(G)_f$, say $\leqslant_+$ and $\leqslant_-$, such that the former makes $\varphi\colon H_1(G)_f \to \mathbb{R}$ into an order-preserving and the latter into an order-reversing map. We will write $\mu_\pm$ for $\mu_{\leqslant_\pm}$.

Since $b_1(G) \geqslant 2$, the projection $p\colon G \to H_1(G)_f$ is non-trivial on F_2, and hence one easily finds generators x, y for F_2 such that $p(x), p(y) <_+ 0$. Put $A = A(g;\mathcal{S},t)$.

Theorem 3.2 tells us that

$$P_{L^2}(G) = P(\det_{\mathcal{D}(G)}(A)) - P(t-1)$$

and

$$P_{L^2}(G;p_k) = P(\det_{\mathcal{D}(\Gamma_k)}(p_k(A)) - P(p_k(t)-1)$$

Note that $P(t-1) = P(p_k(t)-1)$. Thus by Theorems 2.20 and 2.23, it suffices to show

$$\mathfrak{N}\big(P(\det_{\mathcal{D}(\Gamma_k)}(p_k(A)))\big)(\varphi) \leqslant \mathfrak{N}\big(P(\det_{\mathcal{D}(G)}(A))\big)(\varphi)$$

Let $\det_{\mathcal{D}(\mathrm{G})}(A) = s^{-1}r$ with $r, s \in \mathcal{D}(K) * H_1(G)_f$. By the choices of $\leq_\pm$ we have

$$\begin{aligned}\mathfrak{N}\big(P(\det_{\mathcal{D}(G)}(A))\big)(\varphi) &= \|\varphi\|_{P(r)} - \|\varphi\|_{P(s)} \\ &= \varphi(\mu_-(r) - \mu_+(r) + \mu_+(s) - \mu_-(s)) \\ &= \varphi(\mu_-(s^{-1}r) - \mu_+(s^{-1}r)) \\ &= \varphi(\mu_-(\det_{\mathcal{D}(\mathrm{G})}(A)) - \mu_+(\det_{\mathcal{D}(\mathrm{G})}(A)))\end{aligned}$$

and similarly for $\det_{\mathcal{D}(\Gamma_k)}(p_k(A))$ (note that, formally speaking, each $\mu_\pm$ should be replaced by $\operatorname{supp}\mu_\pm$ in the above expression; we omitted the supp for the sake of clarity, and we will continue to do so).

By Lemma 4.6, the matrix A is well-behaved with respect to both $\leqslant_+$ and $\leqslant_-$. Since φ is order-preserving when we consider μ_+ and order-reversing when we consider μ_-, this means that

$$\varphi(\mu_+(\det_{\mathcal{D}(G)}(A))) \leqslant \varphi(\mu_+(\det_{\mathcal{D}(\Gamma_k)} p_k(A))$$

and

$$\varphi(\mu_-(\det_{\mathcal{D}(G)}(A))) \geqslant \varphi(\mu_-(\det_{\mathcal{D}(\Gamma_k)} p_k(A))$$

and the result follows. □

4.1. Fibred cohomology classes. In this short section we look at a cohomology class $\varphi\colon G \to \mathbb{Z}$ that is *fibred* in the sense that its kernel is finitely generated.

Corollary 4.8. *Let* $G = F_2 *_g$ *be a descending HNN extension with* $b_1(G) \geqslant 2$. *If* $\varphi\colon G \to \mathbb{Z}$ *is surjective and fibred, then we have* $[\pm\varphi] \in \Sigma(G)$ *and*

$$\|\varphi\|_T = \|\varphi\|_A = b_1(\ker\varphi) - 1$$

where b_1 *denotes the usual first Betti number.*

Proof. The claim about the Σ-invariant is well-known [**BNS**, Theorem B1], and is in fact an equivalent characterisation of φ being fibred.

Since φ has finitely generated kernel K, it follows from the work of Geoghegan–Mihalik–Sapir–Wise [**GMSW**, Theorem 2.6 and Remark 2.7] that K is finitely generated free itself, say of rank m. Denote the inclusion by $i\colon K \to G$.

By claim (3.26) made in the proof of [**FL3**, Theorem 3.24], we have

$$\|\varphi\|_T = \mathfrak{N}(P_{L^2}(G))(\varphi) = \mathfrak{N}(\mathbb{P}(-\rho_u^{(2)}(G)))(\varphi) = -\chi^{(2)}(i^*\widetilde{T};\mathcal{N}(K))$$

where T is the mapping telescope of a realisation of g. Recall that the K-CW-complex $i^*\widetilde{T}$ is a model for EK and that K is finitely generated free, so

$$-\chi^{(2)}(i^*\widetilde{T};\mathcal{N}(K)) = -\chi^{(2)}(K) = b_1^{(2)}(K) = b_1(K) - 1$$

McMullen showed in [**McM**, Theorem 4.1] that for fibred classes we have

$$b_1(\ker\varphi) - 1 \leqslant \|\varphi\|_A$$

(In fact, McMullen showed that this is an equality when φ lies in the cone over an open face of the Alexander polytope.)

Combining the above results with Theorem 4.7 we obtain

$$b_1(\ker\varphi) - 1 \leqslant \|\varphi\|_A \leqslant \|\varphi\|_T = b_1(\ker\varphi) - 1 \qquad \square$$

5. The L^2-torsion polytope and the BNS-invariant

In this section we relate the L^2-torsion polytope of a descending HNN extension of F_2 with the BNS-invariant introduced in Section 2.10. This approach is motivated by the following results: If M is a compact orientable 3-manifold, the unit norm ball of the Thurston norm is a polytope, and there are certain maximal faces such that a cohomology class comes from a fibration over the circle if and only if it lies in the positive cone over these faces [**Thu**]. Bieri-Neumann-Strebel [**BNS**, Theorem E] showed that the BNS-invariant $\Sigma(\pi_1(M))$ is precisely the projection of these *fibered faces* to the sphere $S(G) = (\mathrm{Hom}(G,\mathbb{R}) \smallsetminus \{0\})/\mathbb{R}_{>0}$. Since the L^2-torsion polytope induces the Thurston norm for descending HNN extensions of F_n, we expect a similar picture in this setting. The work of Friedl-Tillmann [**FT**, Theorem 1.1] provides further evidence for this expectation.

Definition 5.1. Let H be an abelian group with a total ordering $\leqslant$, which is invariant under multiplication. Let R be a skew-field. We define $R(H,\leqslant)$ to be the set of functions $H \to R$ with well-ordered support, that is $f\colon H \to R$ belongs to $R(H,\leqslant)$ if every subset of H whose image under f misses zero has a $\leqslant$-minimal element.

Theorem 5.2 (Malcev, Neumann [**Mal1, Neu**])**.** *Convolution is well-defined on*

$$R(H,\leqslant)$$

and turns it into a skew-field.

Remark 5.3. In fact, given structure maps $\varphi\colon H \to \mathrm{Aut}(R)$ and $\mu\colon H\times H \to R^\times$ of a crossed product $R*H$, one can also define a crossed-product convolution on $R(H,\leqslant)$ in a way completely analogous to the usual construction of crossed product rings (see Definition 2.8). The resulting ring is still a skew-field, and we will denote it by $R*(H,\leqslant)$ for emphasis.

Remark 5.4. In fact the Malcev–Neumann construction works for all biorderable groups, and not merely abelian ones.

In order to relate the L^2-torsion polytope to the BNS-invariant, we first need to put the skew-field $\mathcal{D}(G)$ and the Novikov-Sikorav completion $\widehat{\mathbb{Z}G}_\varphi$ (introduced in Definition 2.27) under the same roof.

Lemma 5.5. *Let $K = \ker(p_0\colon G \to \Gamma_0 = H_1(G)_f)$. Given $\varphi \in \mathrm{Hom}(G,\mathbb{R}) \smallsetminus \{0\}$ with $L = \ker(\varphi)$, let $\leq_\varphi$ be a multiplication invariant total order on $H_1(G)_f$ such that φ is order-preserving (we endow $\mathbb{R}$ with the standard ordering $\leqslant$). We define*

$$\mathfrak{F}(G,\varphi) := \mathcal{D}(K) * (H_1(G)_f, \leq_\varphi)$$

in the sense of Remark 5.3. Then there is a commutative diagram of rings

$$\begin{array}{ccccc} \mathbb{Z}K * H_1(G)_f & \longrightarrow & \mathcal{D}(K) * H_1(G)_f & \longrightarrow & \mathcal{D}(G) \\ \uparrow{\scriptstyle\cong} & & & & \downarrow{\scriptstyle i_\varphi} \\ \mathbb{Z}G & & & & \mathfrak{F}(G,\varphi) \\ \downarrow{\scriptstyle\cong} & & & & \uparrow{\scriptstyle j_\varphi} \\ \mathbb{Z}L * \operatorname{im}\varphi & \longrightarrow & \widehat{\mathbb{Z}L * \operatorname{im}\varphi}_\iota & \xrightarrow{\cong} & \widehat{\mathbb{Z}G}_\varphi \end{array}$$

such that all maps are inclusions, where ι denotes the inclusion $\operatorname{im}\varphi \hookrightarrow \mathbb{R}$, and $\widehat{\mathbb{Z}L * \operatorname{im}\varphi}_\iota$ denotes the Sikorav–Novikov completion of $\mathbb{Z}L * \operatorname{im}\varphi$ with respect to $\iota\colon \operatorname{im}\varphi \to \mathbb{R}$.

Proof. All maps apart from i_φ and j_φ are either obvious or have already been explained. The commutativity of the upper and lower triangle is clear.

Since $\mathfrak{F}(G,\varphi)$ is a skew-field, the universal property of the Ore localisation allows us to define

$$i_\varphi\colon \mathcal{D}(G) \cong T^{-1}(\mathcal{D}(K) * H_1(G)_f) \to \mathfrak{F}(G,\varphi)$$

as the localisation of the obvious inclusion

$$\mathcal{D}(K) * H_1(G)_f \to \mathfrak{F}(G,\varphi)$$

The definition of

$$j_\varphi\colon \widehat{\mathbb{Z}G}_\varphi \cong \widehat{\mathbb{Z}L * \operatorname{im}}\,\varphi_\iota \to \mathfrak{F}(G,\varphi)$$

uses the same formulae as the composition

$$\mathbb{Z}L * \operatorname{im}\varphi \xrightarrow{\cong} \mathbb{Z}G \xrightarrow{\cong} \mathbb{Z}K * H_1(G)_f$$

and we need to verify that this indeed maps to formal sums with well-ordered support with respect to $\leq_\varphi$. But this follows directly from the fact that

$$\varphi\colon H_1(G)_f \to \mathbb{R}$$

is order-preserving. The commutativity of the right-hand triangle follows immediately. □

Definition 5.6. Given $\varphi \in \operatorname{Hom}(G,\mathbb{R})$ and

$$x = \sum_{h \in H_1(G)_f} x_h \cdot h \in \mathcal{D}(K) * (H_1(G)_f, \leq_\varphi)$$

we set

$$S_\varphi(x) = \operatorname{minsupp}_\varphi(x) = \big\{h \in \operatorname{supp}(x) \mid \varphi(h) = \min\{\varphi(\operatorname{supp}(x))\}\big\}$$

and define $\mu_\varphi\colon \mathfrak{F}(G,\varphi)^\times \to \mathfrak{F}(G,\varphi)^\times$ by

$$\mu_\varphi\big(\sum_{h \in H_1(G)_f} x_h \cdot h\big) = \sum_{h \in S_\varphi(x)} x_h \cdot h$$

We record the following properties.

Lemma 5.7. *Let $\varphi \in \operatorname{Hom}(G,\mathbb{R})$.*

(1) The map μ_φ is a group homomorphism.

(2) It restricts to maps (denoted by the same name)

$$\mu_\varphi\colon \mathcal{D}(G)^\times \to \mathcal{D}(G)^\times$$
$$\mu_\varphi\colon \widehat{\mathbb{Z}G}_\varphi^\times \to \mathbb{Z}G \smallsetminus \{0\}$$

and the latter map agrees with $\mu_\varphi\colon \widehat{\mathbb{Z}G}_\varphi^\times \to \mathbb{Z}G\smallsetminus\{0\}$ *from Definition 2.30.*

Proof. This is obvious. □

We now give a practical method for calculating the BNS invariant for descending HNN-extensions of F_2.

Theorem 5.8. *Let G be a descending HNN extension of F_2. Let*

$$\varphi \in \mathrm{Hom}(G, \mathbb{R}) \smallsetminus \{0\}$$

Suppose that x, y are generators of F_2 for which $\varphi(x), \varphi(y) > 0$, and let $g\colon F_2 \to F_2$ be a monomorphism such that $G = F_2{*_g}$, and such that $g(x), g(y)$ have no common prefix. Then $[-\varphi] \in \Sigma(G)$ if and only if

$$\mu_\varphi(1 + t\frac{\partial g(x)}{\partial y} - t\frac{\partial g(y)}{\partial y}) = \pm z$$

for some $z \in G$.

Proof. By Theorem 3.2, (4), we have $-\varphi \in \Sigma(G)$ if and only if the map

$$A\colon \widehat{\mathbb{Z}G_\varphi}^2 \to \widehat{\mathbb{Z}G_\varphi}^2$$

is an isomorphism, where

$$A = A(g; \mathcal{S}, x) = \begin{pmatrix} -t\frac{\partial g(x)}{\partial y} & x-1 \\ 1 - t\frac{\partial g(y)}{\partial y} & y-1 \end{pmatrix}$$

Since $\varphi(y) \neq 0$, the element $y-1$ is invertible in $\widehat{\mathbb{Z}G_\varphi}$, and thus we may perform an elementary row operation over $\widehat{\mathbb{Z}G_\varphi}$ to obtain a triangular $\widehat{\mathbb{Z}G_\varphi}$-matrix

$$B = \begin{pmatrix} -t\frac{\partial g(x)}{\partial y} - (1 - t\frac{\partial g(y)}{\partial y})(y-1)^{-1}(x-1) & 0 \\ 1 - t\frac{\partial g(y)}{\partial y} & y-1 \end{pmatrix}$$

Note that A is invertible over $\widehat{\mathbb{Z}G_\varphi}$ if and only if the diagonal entries of B are invertible in $\widehat{\mathbb{Z}G_\varphi}$. One of the diagonal entries is $y-1$, which we already know to be invertible. The other one is invertible if and only if

$$\mu_\varphi\big(-t\frac{\partial g(x)}{\partial y} - (1 - t\frac{\partial g(y)}{\partial y})(y-1)^{-1}(x-1)\big) = \pm z$$

for some $z \in G$, thanks to Lemma 2.32. But

$$\mu_\varphi\big((1 - t\frac{\partial g(y)}{\partial y})(y-1)^{-1}(x-1)\big) = \mu_\varphi(1 - t\frac{\partial g(y)}{\partial y})$$

and the supports of $1 - t\frac{\partial g(y)}{\partial y}$ and $t\frac{\partial g(x)}{\partial y}$ have a trivial intersection: the lack of common prefixes of $g(x)$ and $g(y)$ implies that the only element in G which could lie in both supports is t, but then we would need to have both $g(x)$ and $g(y)$ starting with y, which would yield a non-trivial common prefix.

This implies

$$\begin{aligned}&\mu_\varphi\big(-t\frac{\partial g(x)}{\partial y} - (1 - t\frac{\partial g(y)}{\partial y})(y-1)^{-1}(x-1)\big)\\ &= \mu_\varphi\big(-t\frac{\partial g(x)}{\partial y} - 1 + t\frac{\partial g(y)}{\partial y}\big)\end{aligned}$$ □

Remark 5.9. The above theorem does not apply to $\varphi \in H^1(G;\mathbb{R}) \smallsetminus \{0\}$ which have $F_2 \leqslant \ker\varphi$. There are however only two such cohomology classes (up to scaling): ψ, the class induced by the HNN-extension $G = F_2 *_g$, which lies in $\Sigma(G)$ if and only if $g\colon F_2 \to F_2$ is an isomorphism, and $-\psi$, which always lies in $\Sigma(G)$.

For every other $\varphi \in H^1(G;\mathbb{R}) \smallsetminus \{0\}$ one easily finds appropriate generators x and y, and then any monomorphism $F_2 \to F_2$ inducing G can be made into the desired form by postcomposing it with a conjugation of F_2. Such a postcomposition does not alter the isomorphism type of G.

Next we are going to relate the L^2-torsion polytope $P_{L^2}(G)$ to the BNS invariant for $G = F_2 *_g$. For this we need some more preparations.

Definition 5.10. Let H be a finitely generated free-abelian group. Let $P \subseteq H \otimes_{\mathbb{Z}} \mathbb{R}$ be a polytope and take $\varphi \in \operatorname{Hom}(H, \mathbb{R})$. We define the *minimal face* of P for φ to be

$$F_\varphi(P) = \{p \in P \mid \varphi(p) = \min\{\varphi(q) \mid q \in P\}\}$$

It is easy to see that F_φ respects Minkowski sums and hence induces group homomorphisms

$$F_\varphi\colon \mathcal{P}(H) \to \mathcal{P}(H)$$
$$F_\varphi\colon \mathcal{P}_T(H) \to \mathcal{P}_T(H)$$

Definition 5.11. Let $K = \ker(p_0\colon G \to H_1(G)_f =: H)$, and let $x \in \mathcal{D}(G) = T^{-1}(\mathcal{D}(K) * H)$ and $\varphi, \psi \in \operatorname{Hom}(G, \mathbb{R}) = \operatorname{Hom}(H, \mathbb{R})$. We call φ and ψ *x-equivalent* if we can write $x = u^{-1}v$ with $u, v \in \mathcal{D}(K) * H$ in such a way that

$$F_\varphi(P(u)) = F_\psi(P(u)) \text{ and } F_\varphi(P(v)) = F_\psi(P(v))$$

We are aiming at proving that the universal L^2-torsion determines the BNS-invariant for descending HNN extensions of free groups. In this process the following lemma is crucial in order to extract algebraic information about Dieudonné determinants from geometric properties of their polytopes.

Lemma 5.12. *Let $x \in \mathcal{D}(G)^\times$ and $\varphi, \psi \in \operatorname{Hom}(G, \mathbb{R})$. If φ and ψ are x-equivalent, then*

$$\mu_\varphi(x) = \mu_\psi(x)$$

Proof. Write $x = u^{-1}v$ with $u, v \in \mathcal{D}(K) * H_1(G)_f$, so that by assumption we have

$$F_\varphi(P(u)) = F_\psi(P(u)) \text{ and } F_\varphi(P(v)) = F_\psi(P(v))$$

But $F_\varphi(P(u)) = F_\psi(P(u))$ implies

$$\operatorname{minsupp}_\varphi(u) = \operatorname{minsupp}_\psi(u)$$

and so

$$\mu_\varphi(u) = \mu_\psi(u)$$

The same argument applies to v and so the claim follows from

$$\mu_\varphi(x) = \mu_\varphi(u)^{-1} \cdot \mu_\varphi(v)$$

□

The following is similar to [**FT**, Theorem 1.1]; although we do not provide markings on the polytopes which fully detect the BNS-invariant, Theorem 5.8 makes up for this lack. The crucial point now is that the BNS invariant is locally determined by a polytope.

Theorem 5.13. *Let $g\colon F_2 \to F_2$ be a monomorphism and let $G = F_2 *_g$ be the associated descending HNN extension. Given $\varphi \in \operatorname{Hom}(G, \mathbb{R}) \smallsetminus \{0\}$ such that $-\varphi$ is not the epimorphism induced by $F_2 *_g$, there exists an open neighbourhood U of $[\varphi]$ in $S(G)$ and an element $d \in \mathcal{D}(G)^\times$ such that:*

(1) The image of d under the quotient maps

$$\mathcal{D}(G)^\times \to \mathcal{D}(G)^\times/[\mathcal{D}(G)^\times, \mathcal{D}(G)^\times] \cong K_1^w(\mathbb{Z}G) \to \mathrm{Wh}^w(G)$$

is $-\rho_u^{(2)}(G)$. In particular $P_{L^2}(G) = P(d)$ in $\mathcal{P}_T(H_1(G)_f)$.

(2) For every $\psi \in \mathrm{Hom}(G,\mathbb{R}) \smallsetminus \{0\}$ which satisfies $[\psi] \in U$ and is d-equivalent to φ, we have $[-\varphi] \in \Sigma(G)$ if and only if $[-\psi] \in \Sigma(G)$.

Proof. Suppose that $\ker\varphi \neq F_2$. We easily find generators x, y of F_2 for which $\varphi(x), \varphi(y) > 0$. Set

$$U = \{[\psi] \mid \psi(x) > 0 \text{ and } \psi(y) > 0\} \subseteq S(G)$$

This is clearly an open neighbourhood of $[\varphi]$. Suppose that $[\psi] \in U$.

Let $A = A(g; \mathcal{S}, x)$, as in the proof of Theorem 5.8. Since $\varphi(y) \neq 0$, we can still form the matrix B from Theorem 5.8, and $[-\varphi] \in \Sigma(G)$ if and only if A is invertible over $\widehat{\mathbb{Z}G_\varphi}$ if and only if B is invertible over $\widehat{\mathbb{Z}G_\varphi}$.

Since B is obtained from A by an elementary row operation over $\mathfrak{F}(G,\varphi)$ in which we add a multiple of the last row to another row, and such operations do not affect the canonical representative of the Dieudonné determinant, we have

$$i_\varphi(\det{}^c_{\mathcal{D}(G)}(A)) = \det{}^c_{\mathfrak{F}(G,\varphi)}(A) = \det{}^c_{\mathfrak{F}(G,\varphi)}(B)$$

which is the product of the diagonal entries of B. Note that B is invertible over $\widehat{\mathbb{Z}G_\varphi}$ if and only if the diagonal entries are invertible in $\widehat{\mathbb{Z}G_\varphi}$, which is the case if and only if their product is invertible in $\widehat{\mathbb{Z}G_\varphi}$ since $\widehat{\mathbb{Z}G_\varphi}$ is a domain. Thus, by Lemma 5.7, $[-\varphi] \in \Sigma(G)$ if and only if $\mu_\varphi(\det{}^c_{\mathfrak{F}(G,\varphi)}(B)) = \mu_\varphi(i_\varphi(\det{}^c_{\mathcal{D}(G)}(A))$ is of the form $\pm z$ for some $z \in G$.

The same arguments apply to ψ since $\psi(y) \neq 0$. By Lemma 5.7, it therefore suffices to prove

$$\mu_\varphi(i_\varphi(\det{}^c_{\mathcal{D}(G)}(A))) = \mu_\psi(i_\psi(\det{}^c_{\mathcal{D}(G)}(A)))$$

If we put $d := \det{}^c_{\mathcal{D}(G)}(A) \cdot (x-1)^{-1}$, then this is equivalent to

$$\mu_\varphi(i_\varphi(d)) = \mu_\psi(i_\psi(d))$$

since $\varphi(x), \psi(x) > 0$. But this is true by Lemma 5.12 if we assume that φ and ψ are d-equivalent.

On the other hand, Theorem 3.2 (1) says that d_U maps under the quotient maps

$$\mathcal{D}(G)^\times \to \mathcal{D}(G)^\times/[\mathcal{D}(G)^\times, \mathcal{D}(G)^\times] \cong K_1^w(\mathbb{Z}G) \to \mathrm{Wh}^w(G)$$

to $-\rho_u^{(2)}(G)$, as desired. This finishes the proof in the case that $\ker\varphi \neq F_2$.

Now suppose that $F_2 \leqslant \ker\varphi$. Since $-\varphi$ is not induced by the HNN extension, we must have $\varphi(t) > 0$.

Let us choose a generating set x, y for F_2, and set

$$U = \left\{[\psi] \mid \psi(t) > |\psi(z)|, z \in \operatorname{supp} \frac{\partial g(y)}{\partial y}\right\}$$

Again, this is an open neighbourhood of $[\varphi]$.

We proceed similarly to the previous case. Observing that $1-t$ is invertible over $\widehat{\mathbb{Z}G_\varphi}$ and $\widehat{\mathbb{Z}G_\psi}$ reduces the problem to verifying whether the matrix $A(g,\mathcal{S},t)$ is invertible over $\widehat{\mathbb{Z}G_\varphi}$ and $\widehat{\mathbb{Z}G_\psi}$. The bottom-right entry of $A(g,\mathcal{S},t)$ is $1 - t\frac{\partial g(y)}{\partial y}$, which is invertible for φ and every ψ with $[\psi] \in U$ by construction. If φ and ψ are additionally d-equivalent for $d := \det{}^c_{\mathcal{D}(G)}(A(g,\mathcal{S},t)) \cdot (t-1)^{-1}$, we now continue in precisely the same way as before. ∎

Remark 5.14. Note that the result in the latter case also follows from the observation that $\Sigma(G)$ is open, since $[-\varphi] \in \Sigma(G)$.

Note also that our neighbourhood U is very explicit, and rather large, especially when $\ker \varphi \neq F_2$.

6. UPG automorphisms

In this section we will strengthen Theorem 4.7 and Theorem 5.13 for a class of free group automorphisms.

Definition 6.1 (Polynomially growing and UPG automorphism)**.** An automorphism $f : F_n \to F_n$ is *polynomially growing* if the quantity $d(1, f^n(g))$ grows at most polynomially in n for every $g \in F_n$, where 1 denotes the identity in G and d is some word metric on F_n. If, additionally, the image $\overline{f}$ of f under the obvious map $\mathrm{Aut}(F_n) \to \mathrm{GL}(n, \mathbb{Z})$ is unipotent, i.e. $\mathrm{id} - \overline{f}$ is nilpotent, then f will be called *UPG*.

The main result of Cashen-Levitt [**CL**, Theorem 1.1] reads as follows.

Theorem 6.2. *Let $G = F_n \rtimes_g \mathbb{Z}$ with $n \geqslant 2$ and g polynomially growing. Then there are elements $t_1, ..., t_{n-1} \in G \smallsetminus F_n$ such that*

$$\Sigma(G) = -\Sigma(G) = \{[\varphi] \in S(G) \mid \varphi(t_i) \neq 0 \text{ for all } 1 \leqslant i \leqslant n-1\}$$

Motivated by this, we prove

Theorem 6.3. *Let $G = F_n \rtimes_g \mathbb{Z}$ with $n \geqslant 1$ and g a UPG automorphism. Denote by $p_k \colon G \to \Gamma_k = G/G_r^{k+1}$ the projection, where G_r^k denotes the k-th subgroup of the rational derived series. For simplicity write Γ_∞ for G and p_∞ for id_G.*

Then there are elements $t_1, ..., t_{n-1} \in G \smallsetminus F_n$ which can be chosen to coincide with those of Theorem 6.2 such that for $k \in \mathbb{N} \cup \{\infty\}$

$$\rho_u^{(2)}(G; p_k) = -\sum_{i=1}^{n-1} [\mathbb{Z}\Gamma_k \xrightarrow{p_k(1-t_i)} \mathbb{Z}\Gamma_k] \tag{16}$$

In particular,

$$P_{L^2}(G; p_k) = \sum_{i=1}^{n-1} P(1 - t_i) \in \mathcal{P}(H_1(G)_f)$$

is a polytope (and not merely a difference of polytopes) which is independent of $k \in \mathbb{N} \cup \{\infty\}$.

Combining the previous two results, we see that the BNS-invariant of UPG automorphisms is easily determined by their L^2-torsion polytope. More precisely, we have the following analogue of [**FT**, Theorem 1.1].

Corollary 6.4. *Let $G = F_n \rtimes_g \mathbb{Z}$ with $n \geqslant 2$ and g a UPG automorphism. Let $\varphi \in H^1(G; \mathbb{R})$. Then $[\varphi] \in \Sigma(G)$ if and only if $F_\varphi(P_{L^2}(G)) = 0$ in $\mathcal{P}_T(H_1(G)_f)$.*

Proof. Any one-dimensional face of

$$P_{L^2}(G) = \sum_{i=1}^{n-1} P(1 - t_i)$$

contains a translate of $P(1 - t_i)$ for some $1 \leqslant i \leqslant n-1$.

Now $F_\varphi(P_{L^2}(G)) \neq 0$ if and only if $F_\varphi(P_{L^2}(G))$ contains a one-dimensional face, i.e. a translate of $P(1-t_i)$ for some i. This is equivalent to $\varphi(t_i) = 0$ for some i, which by Theorem 6.2 is equivalent to $[\varphi] \notin \Sigma(G)$. ∎

Remark 6.5. We suspect Theorem 6.3 to hold as well for polynomially growing automorphisms. It is well-known that any polynomially growing automorphism has a power that is UPG, see Bestvina–Feighn–Handel's [**BFH**, Corollary 5.7.6]. Thus, in order to reduce Theorem 6.3 for polynomially growing automorphisms to the case of UPG automorphisms, one needs a better understanding of the restriction homomorphism

$$i^*\colon \operatorname{Wh}^w(F_n \rtimes_g \mathbb{Z}) \to \operatorname{Wh}^w(F_n \rtimes_{g^k} \mathbb{Z})$$

(induced by the obvious inclusion $i\colon F_n \rtimes_{g^k} \mathbb{Z} \to F_n \rtimes_g \mathbb{Z}$) since it maps $\rho_u^{(2)}(F_n \rtimes_g \mathbb{Z})$ to $\rho_u^{(2)}(F_n \rtimes_{g^k} \mathbb{Z})$ (see Lemma 2.16 (3)).

We also obtain

Corollary 6.6. *Let $G = F_n \rtimes_g \mathbb{Z}$ with $n \geqslant 2$ and g a UPG automorphism. Let $\varphi \in H^1(G;\mathbb{R})$. Then for all $k \in \mathbb{N} \cup \{\infty\}$ we have*

$$\|\varphi\|_A = \delta_k(\varphi) = \|\varphi\|_T.$$

Proof. This follows directly from the fact that $P_{L^2}(G;p_k)$ is independent of $k \in \mathbb{N}\cup\{\infty\}$ as stated in Theorem 6.3. Note that $b_1(G) \geqslant 2$ by [**CL**, Remark 5.6]. Hence we get as special cases $P_{L^2}(G;p_0) = P_A(G)$ by Theorem 3.2 (3) and this polytope determines the Alexander norm, and on the other hand $P_{L^2}(G;p_\infty) = P_{L^2}(G)$ which determines the Thurston norm. □

Theorems 6.2 and 6.3 both rely on the following lemma which follows from the train track theory of Bestvina–Feighn–Handel [**BFH**]; see [**CL**, Proposition 5.9] for the argument.

Lemma 6.7. *For $n \geqslant 2$ and a UPG automorphism $g \in \operatorname{Aut}(F_n)$, there exists $h \in \operatorname{Aut}(F_n)$ representing the same outer automorphism class as g, such that either*

(1) there is an h-invariant splitting $F_n = B_1 * B_2, h = h_1 * h_2$; or
(2) there is a splitting $F_n = B_1 * \langle x\rangle$ such that B_1 is h-invariant and $h(x) = xu$ for some $u \in B_1$.

This lemma allows us two write the semi-direct product associated to a UPG automorphism as an iterated splitting over infinite cyclic subgroups with prescribed vertex groups. This is explained in [**CL**, Lemma 5.10] and will be repeated in the following proof.

Proof of Theorem 6.3 . We prove the statement by induction on n. For the base case $n = 1$ we have $F_1 \rtimes_g \mathbb{Z} \cong \mathbb{Z}^2$ and $\rho_u^{(2)}(\mathbb{Z}^2;p_k) = 0$ for all $k \in \mathbb{N}\cup\{\infty\}$ by [**FL3**, Example 2.7] which already verifies (16).

For the inductive step, we may assume that $g = h$ in the notation of Lemma 6.7 since the isomorphism class of $F_n \rtimes_g \mathbb{Z}$ only depends on the outer automorphism class of g. We analyse the two cases appearing in Lemma 6.7 separately.

Case 1: There is a g-invariant splitting $F_n = B_1 * B_2, g = g_1 * g_2$. Write

$$G_i = B_i \rtimes_{g_i} \mathbb{Z}$$

and let $G_0 = \mathbb{Z} \hookrightarrow G_i$ be the inclusion of the second factor. Then we have

$$G = F_n \rtimes_g \mathbb{Z} \cong G_1 *_{G_0} G_2$$

and the Fox matrix of g is of the form

$$F(g) = \begin{pmatrix} F(g_1) & 0 \\ 0 & F(g_2) \end{pmatrix}$$

Let $j_i\colon G_i \to G$ be the inclusions, and denote a generator of G_0 and its image in the various groups G_i by t.

By [**CL**, Remark 5.6], we have $b_1(G) \geqslant 2$ and similarly for G_1 and G_2. Hence by Theorem 3.2 (2) and (3) as well as the above matrix decomposition, we compute in $\mathrm{Wh}^w(\Gamma_k)$

$$\begin{aligned} \rho_u^{(2)}(G;p_k) &= -[p_k(I - t\cdot F(g))] + [p_k(t-1)] \\ &= -[p_k(I - t\cdot F(g_1))] - [p_k(I - t\cdot F(g_2))] + [p_k(t-1)] \\ &= (j_1)_*(\rho_u^{(2)}(G_1;p_k^1)) + (j_2)_*(\rho_u^{(2)}(G_2;p_k^2)) - [p_k(t-1)] \end{aligned} \tag{17}$$

where p_k^i denote the projections on the quotients of the rational derived series of G_i. Here we have used that in our setting p_k^i can be seen as a restriction of p_k.

Denote the rank of B_i by r_i. By the inductive hypothesis applied to G_i, there are elements

$$t'_1, \dots, t'_{r_1-1} \in G_1 \setminus B_1$$

and

$$t''_1, \dots, t''_{r_2-1} \in G_2 \setminus B_2$$

such that

$$\rho_u^{(2)}(G_1;p_k^1) = -\sum_{i=1}^{r_1-1} [p_k^1(1-t'_i)] \tag{18}$$

and

$$\rho_u^{(2)}(G_2;p_k^2) = -\sum_{i=1}^{r_2-1} [p_k^2(1-t''_i)] \tag{19}$$

Notice that $r_1 + r_2 = n$. Moreover, the corresponding induction step in the proof of Theorem 6.2 adds t to the union of the t'_i and the t''_i. Thus the desired statement (16) follows by combining (17), (18), and (19).

Case 2: There is a splitting $F_n = B_1 * \langle x \rangle$ such that B_1 is g-invariant and $g(x) = xu$ for some $u \in B_1$. In this case, let $g_1 = g|_{G_1}$, $G_1 = B_1 \rtimes_{g_1} \mathbb{Z} \subseteq G$, and denote the stable letter of G_1 and G by t.

In this case, the Fox matrix of g takes the form

$$F(g) = \begin{pmatrix} F(g_1) & 0 \\ * & 1 \end{pmatrix}$$

From this we compute in $\mathrm{Wh}^w(\Gamma_k)$ similarly as in the first case

$$\begin{aligned} \rho_u^{(2)}(G;p_k) &= -[p_k(I - t\cdot F(g))] + [p_k(t-1)] \\ &= -[p_k(I - t\cdot F(g_1))] - [p_k(1-t)] + [p_k(t-1)] \\ &= \rho_u^{(2)}(G_1;p_k^1) - [p_k(1-t)] \end{aligned} \tag{20}$$

The corresponding induction step in the proof of Theorem 6.2 adds t to the elements t'_i belonging to G_1 which we get from the induction hypothesis.

This finishes the proof of Theorem 6.3. □

Remark 6.8. The same strategy as above can be used to prove that the ordinary L^2-torsion $\rho^{(2)}(g) := \rho^{(2)}(F_n \rtimes_g \mathbb{Z}) \in \mathbb{R}$ vanishes for all polynomially growing automorphisms. Here the reduction to UPG automorphisms explained in Remark 6.5 is simpler since we have $\rho^{(2)}(g^k) = k \cdot \rho^{(2)}(g)$, so that the vanishing of the L^2-torsion of some power of g implies the vanishing of the L^2-torsion of g. This is a special case of a result of Clay [**Cla**, Theorem 5.1].

CHAPTER IV

Nielsen Realisation by Gluing: Limit Groups and Free Products

This is joint work with Sebastian Hensel.

We generalise the Karrass–Pietrowski–Solitar and the Nielsen realisation theorems from the setting of free groups to that of free products. As a result, we obtain a fixed point theorem for finite groups of outer automorphisms acting on the relative free splitting complex of Handel–Mosher and on the outer space of a free product of Guirardel–Levitt, as well as a relative version of the Nielsen realisation theorem, which in the case of free groups answers a question of Karen Vogtmann. We also prove Nielsen realisation for limit groups, and as a byproduct obtain a new proof that limit groups are CAT(0).
The proofs rely on a new version of Stallings' theorem on groups with at least two ends, in which some control over the behaviour of virtual free factors is gained.

1. Introduction

In its original form, the *Nielsen realisation problem* asks which finite subgroups of the mapping class group of a surface can be realised as groups of homeomorphisms of the surface. A celebrated result of Kerckhoff [**Ker1, Ker2**] answers this in the positive for all finite subgroups, and even allows for realisations by isometries of a suitable hyperbolic metric.

Subsequently, similar realisation results were found in other contexts, perhaps most notably for realising finite groups in $\mathrm{Out}(F_n)$ by isometries of a suitable graph (independently by [**Cul**], [**Khr1**], [**Zim1**]; compare [**HOP**] for a different approach).

In this article, we begin to develop a *relative* approach to Nielsen realisation problems. The philosophy here is that if a group G allows for a natural decomposition into pieces, then Nielsen realisation for $\mathrm{Out}(G)$ may be reduced to realisation in the pieces, and a *gluing problem.* In addition to just solving Nielsen realisation for finite subgroups of $\mathrm{Out}(G)$, such an approach yields more explicit realisations, which also exhibit the structure of pieces for G.

We demonstrate this strategy for two classes of groups: free products and limit groups. In another article, we use the results presented here, together with the philosophy of relative Nielsen realisation, to prove Nielsen realisation for certain right-angled Artin groups ([**HK1**]).

The early proofs of Nielsen realisation for free groups rely in a fundamental way on a result of Karrass–Pietrowski–Solitar [**KPS**], which states that every finitely generated virtually free group acts on a tree with finite edge and vertex stabilisers. In the language of Bass–Serre theory, it amounts to saying that such a virtually free group is a fundamental group of a finite graph of groups with finite edge and vertex groups.

This result of Karrass–Pietrowski–Solitar in turn relies on the celebrated theorem of Stallings on groups with at least two ends [**Sta1, Sta2**]. Stallings' theorem states that any finitely generated group with at least two ends splits over a finite group, which means that it acts on a tree with a single edge orbit and finite edge stabilisers. Equivalently: it is a fundamental group of a graph of groups with a single edge and a finite edge group.

In the first part of this article, we generalise these results to the setting of a free product

$$A = A_1 * \cdots * A_n * B$$

in which we (usually) require the factors A_i to be finitely generated, and B to be a finitely generated free group. Consider any finite group H acting on A by outer automorphisms in a way preserving the given free-product decomposition, by which we mean that each element of H sends each subgroup A_i to some A_j (up to conjugation); note that we do not require the action of H to preserve B in any way. We then obtain a corresponding group extension

$$1 \to A \to \overline{A} \to H \to 1$$

In this setting we prove (for formal statements, see the appropriate sections)

Relative Stallings' theorem (Theorem 2.9): $\overline{A}$ splits over a finite group, in such a way that each A_i fixes a vertex in the associated action on a tree.

Relative Karrass–Pietrowski–Solitar theorem (Theorem 4.1): $\overline{A}$ acts on a tree with finite edge stabilisers, and with each A_i fixing a vertex of the tree, and with, informally speaking, all other vertex groups finite.

Relative Nielsen realisation theorem (Theorem 7.5): Suppose that we are given complete non-positively curved (i.e. locally CAT(0)) spaces X_i realising the induced actions of H on the factors A_i. Then the action of H can be realised by a complete non-positively curved space X; in fact X can be chosen to contain the X_i in an equivariant manner.

We emphasise that such a relative Nielsen realisation is new even if all A_i are free groups, in which case it answers a question of Karen Vogtmann.

The classical Nielsen realisation for graphs immediately implies that a finite subgroup $H < \mathrm{Out}(F_n)$ fixes points in the Culler–Vogtmann Outer Space (defined in [**CV**]), as well as in the complex of free splittings of F_n (which is a simplicial closure of Outer Space). As an application of the work in this article, we similarly obtain fixed point statements (Corollaries 5.1 and 6.1) for the graph of relative free splittings defined by Handel and Mosher [**HM**], and the outer space of a free product defined by Guirardel and Levitt [**GL**].

In the last section of the paper we prove

Theorem 8.11. *Let A be a limit group, and let*

$$A \to \overline{H} \to H$$

be an extension of A by a finite group H. Then there exists a complete locally CAT(κ) space X realising the extension $\overline{H}$, where $\kappa = -1$ when A is hyperbolic, and $\kappa = 0$ otherwise.

This theorem is obtained by combining the classical Nielsen realisation theorems (for free, free-abelian and surface groups – see Theorems 8.1 to 8.3) with the existence of an invariant JSJ decomposition shown by Bumagin–Kharlampovich–Myasnikov [**BKM**].

Observe that we obtain optimal curvature bounds for our space X – it has been proved by Alibegović–Bestvina [**AB**] that limit groups are CAT(0), and by Sam Brown [**Bro2**] that a limit group is CAT(-1) if and only if it is hyperbolic.

Also, taking H to be the trivial group gives a new (more direct) proof of the fact that limit groups are CAT(0).

Throughout the paper, we are going to make liberal use of the standard terminology of graphs of groups. The reader may find all the necessary information in Serre's book [**Ser1**]. We are also going to make use of standard facts about CAT(0) and non-positively curved (NPC) spaces, as well as more general CAT(κ) spaces; the standard reference here is the book by Bridson–Haefliger [**BH**].

Acknowledgements. The authors would like to thank Karen Vogtmann for discussions and suggesting the statement of relative Nielsen realisation for free groups, Stefan Witzel for pointing out the work of Sam Brown, and the referee for extremely valuable comments.

2. Relative Stallings' theorem

In this section we will prove the relative version of Stallings' theorem. Before we can begin with the proof, we need a number of definitions to formalise the notion of a free splitting that is preserved by a finite group action.

Convention 2.1. When talking about free factor decompositions

$$A = A_1 * \cdots * A_n * B$$

of some group A, we will always assume that at least two of the factors $\{A_1, \dots, A_n, B\}$ are non-trivial.

Definition 2.2. Suppose that $\varphi\colon H \to \mathrm{Out}(A)$ is a homomorphism with a finite domain. Let $A = A_1 * \cdots * A_n * B$ be a free factor decomposition of A. We say that this decomposition is *preserved by* H if and only if for every i and every $h \in H$, there is some j such that $h(A_i)$ is conjugate to A_j.

We say that a factor A_i is *minimal* if and only if for any $h \in H$ the fact that $h(A_i)$ is conjugate to A_j implies that $j \geqslant i$.

Remark 2.3. Note that when the decomposition is preserved, we obtain an induced action $H \to \mathrm{Sym}(n)$ on the indices $1, \dots, n$. We may thus speak of the stabilisers $\mathrm{Stab}_H(i)$ inside H. Furthermore, we obtain an induced action

$$\mathrm{Stab}_H(i) \to \mathrm{Out}(A_i)$$

The minimality of factors is merely a way of choosing a representative of each H orbit in the action $H \to \mathrm{Sym}(n)$.

Remark 2.4. Given an action $\varphi\colon H \to \mathrm{Out}(A)$, with φ injective and A with trivial centre, we can define $\overline{A} \leqslant \mathrm{Aut}(A)$ to be the preimage of $H = \mathrm{im}\,\varphi$ under the natural map $\mathrm{Aut}(A) \to \mathrm{Out}(A)$. We then note that $\overline{A}$ is an extension of A by H:

$$1 \to A \to \overline{A} \to H \to 1$$

and the left action of H by outer automorphisms agrees with the left conjugation action inside the extension $\overline{A}$.

Observe that then for each i we also obtain an extension

$$1 \to A_i \to \overline{A_i} \to \mathrm{Stab}_H(i) \to 1$$

where $\overline{A_i}$ is the normaliser of A_i in $\overline{A}$.

We emphasise that this construction works even when A_i itself is not centre-free. In this case it carries more information than the induced action $\mathrm{Stab}_H(i) \to$

$\mathrm{Out}(A_i)$ (e.g. consider the case of $A_i = \mathbb{Z}$ – there are many different extensions corresponding to the same map to $\mathrm{Out}(\mathbb{Z})$).

We will now begin the proof of the relative version of Stallings' theorem. It will use ideas from both Dunwoody's proof [**Dun1**] and Krön's proof [**Krö**][1] of Stallings' theorem, which we now recall.

Convention 2.5. If E is a set of edges in a graph Θ, we write $\Theta - E$ to mean the graph obtained from Θ by removing the interiors of edges in E.

Definition 2.6. Let Θ be a graph. A finite subset E of the edge set of Θ is called a set of *cutting edges* if and only if $\Theta - E$ is disconnected and has at least two infinite components.

A *cut* C is the union of all vertices contained in an infinite connected complementary component of some set of cutting edges. The *boundary* of C consists of all edges with exactly one endpoint in C.

Given two cuts C and D, we call them *nested* if and only if C or its complement C^* is contained in D or its complement D^*. Note that C^* and D^* do not need to be cuts.

We first aim to show the following theorem which is implicit in [**Krö**].

Theorem 2.7 ([**Krö**])**.** *Suppose that Θ is a connected graph on which a group G acts. Let $\mathcal{P}$ be a subset of the edge set of Θ, which is stable under the G-action. If there exists a set of cutting edges lying in $\mathcal{P}$, then there exists a cut C whose boundary lies in $\mathcal{P}$, such that C^* is also a cut, and such that furthermore for any $g \in G$ the cuts C and $g.C$ are nested.*

Sketch of proof. In order to prove this, we recall the following terminology, roughly following Dunwoody. We say that C is a $\mathcal{P}$-cut if and only if its boundary lies in $\mathcal{P}$. Say that a $\mathcal{P}$-cut is *$\mathcal{P}$-narrow*, if and only if its boundary contains the minimal number of elements among all $\mathcal{P}$-cuts. Note that for each $\mathcal{P}$-narrow cut C, the complement C^* is also a cut, as otherwise we could remove some edges from the boundary of C and get another $\mathcal{P}$-cut.

Given any edge $e \in \mathcal{P}$, there are finitely many $\mathcal{P}$-narrow cuts which contain e in its boundary. This is shown by Dunwoody [**Dun1**, 2.5] for narrow cuts, and the proof carries over to the $\mathcal{P}$-narrow case. Alternatively, Krön [**Krö**, Lemma 2.1] shows this for sets of cutting edges which cut the graph into exactly two connected components, and $\mathcal{P}$-narrow cuts have this property.

Now, consider for each $\mathcal{P}$-narrow cut C the number $m(C)$ of $\mathcal{P}$-narrow cuts which are not nested with C (this is finite by [**Dun1**, 2.6]). Call a $\mathcal{P}$-narrow cut *optimally nested* if $m(C)$ is smallest amongst all $\mathcal{P}$-narrow cuts. The proof of Theorem 3.3 of [**Krö**] now shows that optimally nested $\mathcal{P}$-cuts are all nested with each other. This shows Theorem 2.7. □

To use that theorem, recall

Theorem 2.8 ([**Dun1**, Theorem 4.1])**.** *Let G be a group acting on a graph Θ. Suppose that there exists a cut C, such that*

1. *C^* is also a cut; and*
2. *there exists $g \in G$ such that $g.C$ is properly contained in C or C^*; and*
3. *C and $h.C$ are nested for any $h \in G$.*

Let E be the boundary of C. Then G splits over the stabiliser of E , and the stabiliser of any component of $\Theta - G.E$ is contained in a conjugate of a vertex group.

[1] We warn the reader that later parts of Krön's paper are not entirely correct; we only rely on the early, correct sections.

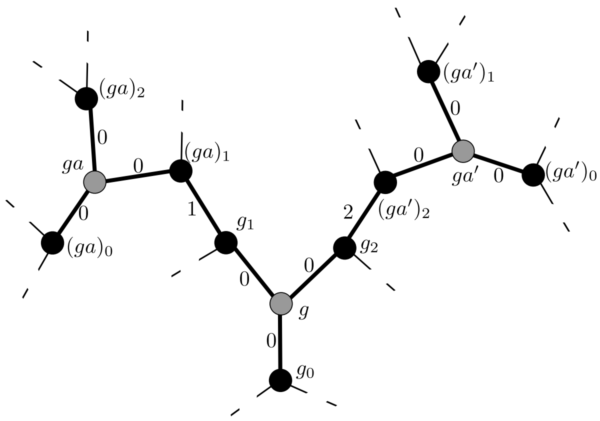

FIGURE 2.1. A local picture of the graph Θ.

Now we are ready for our main splitting result.

Theorem 2.9 (Relative Stallings' Theorem). *Let $\varphi\colon H \to \mathrm{Out}(A)$ be a monomorphism with a finite domain. Let $A = A_1 * \cdots * A_n * B$ be a free product decomposition with each A_i and B finitely generated, and suppose that it is preserved by H. Let $\overline{A}$ be the preimage of $H = \mathrm{im}\ \varphi$ in $\mathrm{Aut}(A)$. Then $\overline{A}$ acts on a tree with finite quotient so that each A_i fixes a vertex, and no non-trivial subgroup of A fixes any edge.*

Note in particular that the quotient of the associated tree by $\overline{A}$ has a single edge.

Proof. Before we begin the proof in earnest, we will give a brief outline of the strategy. First, we will define a variant of the Cayley graph for $\overline{A}$ in which the free product structure of A will be visible (in fact, a subgraph will collapse to the Bass-Serre tree of the free product decomposition of A). This graph will contain the different copies if A_i disjointly, separated by edges labelled with a certain label. We will then aim to show that there is a set of cutting edges just using edges with that label – which, using Theorem 2.8, will yield the desired action on a tree.

Let $\mathcal{A}_i$ and $\mathcal{B}$ be finite generating sets of A_i and B, respectively (for all $i \leqslant n$). We also choose a finite set $\mathcal{H} \subset \overline{A}$ which maps onto H under the natural epimorphism $\overline{A} \to H$. Note that $\bigcup_i \mathcal{A}_i \cup \mathcal{B} \cup \mathcal{H}$ is a generating set of $\overline{A}$.

We define Θ to be a variation of the (right) Cayley graph of $\overline{A}$ with respect to the generating set $\bigcup_i \mathcal{A}_i \cup \mathcal{B} \cup \mathcal{H}$. Intuitively, every vertex of the Cayley graph will be "blown up" to a finite tree (see Figure 2.1). More formally, the vertex set of Θ is

$$V(\Theta) = \overline{A} \sqcup (\overline{A} \times \{0, \ldots, n\})$$

We adopt the notation that a vertex corresponding to an element in $\overline{A}$ will simply be denoted by g, whereas a vertex (g, i) in the second factor will be denoted by g_i.

We now define the edge set, together with a labelling of the edges by integers $0, 1, \ldots, n$, as follows:

- for each $g \in \overline{A}$ and each $i \in \{0, \ldots, n\}$ we have an edge labelled by 0 connecting g to g_i;
- for each $g \in \overline{A}$, each $i \geqslant 1$ and each $a \in \mathcal{A}_i$, we have an edge labelled by i from g_i to $(ga)_i$;
- for each $g \in \overline{A}$, and each $b \in \mathcal{B} \cup \mathcal{H}$, we have an edge labelled by 0 from g_0 to $(gb)_0$.

The group $\overline{A}$ acts on Θ on the left, preserving the labels. The action is free and co-compact. The graph Θ retracts via a quasi-isometry onto a usual Cayley graph of $\overline{A}$ by collapsing edges connecting g to g_i. Also note that there are copies of the Cayley graphs of the A_i with respect to the generating set $\mathcal{A}_i$ in Θ, where each edge has the label i.

Let Ω denote a graph constructed in the same way for the group A with respect to the generating set $\bigcup \mathcal{A}_i \cup \mathcal{B}$. There is a natural embedding of Ω into Θ, and hence we will consider Ω as a subgraph of Θ. Note that this embedding is also a quasi-isometry.

We will now construct n specific quasi-isometric retractions of Θ onto Ω. These will be used later to modify paths in order to avoid edges with certain labels.

Let us fix $i \in \{1, \ldots, n\}$. For each $h \in H$ we pick a representative $h_i \in \overline{A}$ thereof, such that $h_i A_i {h_i}^{-1} = A_j$ for a suitable (and unique) j; for $1 \in H$ we pick $1 \in \overline{A}$ as a representative. These elements h_i are coset representatives of the normal subgroup A of $\overline{A}$.

Such a choice defines a retraction $\rho_i \colon \Theta \to \Omega$ in the following way: each vertex g is mapped to the unique vertex g' where $g' \in A$ and $g' h_i = g$ for some h_i; the vertex g_k is then mapped to $(g')_k$. An edge labelled by 0 connecting g to g_k is sent to the edge connecting g' to $(g')_k$. The remaining edges with label 0 are sent in an A-equivariant fashion to paths connecting the image of their endpoints; the lengths of such paths are uniformly bounded, since (up to the A-action) there are only finitely many edges with label 0.

Similarly, the edges of label $k \notin \{0, i\}$ are mapped in an A-equivariant manner to paths connecting the images of their endpoints; again, their length is uniformly bounded.

Each edge labelled by i is sent A-equivariantly to a path connecting the images of its endpoints, such that the path contains edges labelled only by some j (where j is determined by the coset of A the endpoints lie in); such a path exist by the choice of the representatives h_i.

Note that each such retraction ρ_i is a (κ_i, κ_i)-quasi-isometry for some $\kappa_i \geqslant 1$; we set $\kappa = \max_i \kappa_i$.

Now we are ready to construct a set of cutting edges in Θ.

Consider the ball $B_\Omega(1,1)$ of radius 1 around the vertex 1 in Ω (all of whose edges are labelled by 0). Since A is a nontrivial free product, the identity element disconnects the Cayley graph into at least two infinite components. Hence, $B_\Omega(1,1)$ disconnects Ω also into at least two infinite components; let us take two vertices of Ω, x and y, lying in distinct infinite components of $\Omega - B_\Omega(1,1)$, and such that

$$d_\Omega(1,x) = d_\Omega(1,y) \geqslant \kappa^2 + 4$$

Now let E denote the set of all edges lying in the ball $B_\Theta(1, \kappa^2+4)$ labelled by 0. We claim that E disconnects Θ into at least two infinite components. Note that $\Theta - E$ has finitely many components, since E is finite. By possibly choosing x, y even further from each other, it therefore suffices to show that E disconnects x from y (viewed as vertices of Θ).

Suppose for a contradiction that there exists a path γ in $\Theta - E$ connecting x to y. Using any of the quasi-isometries ρ_i we immediately see that γ has to go through $B_\Theta(1, \kappa^2+4)$, since $\rho_i(\gamma)$ must intersect $B_\Omega(1,1)$. Note that if $\gamma' \subset \gamma$ is a subpath lying completely in $B_\Theta(1, \kappa^2+4)$, then γ' only traverses edges with the same label (as γ does not intersect E). Thus, we can write γ as a concatenation

$$\gamma = \gamma_1 * \cdots * \gamma_m$$

where each γ_i intersects $B_\Theta(1, \kappa^2 + 4)$ only at edges of one label, and its endpoints lie outside of $B_\Theta(1, \kappa^2 + 4)$. We modify each γ_i by pre- and post-concatenating it with a path of length at most 4 (note that all the elements of $\mathcal{H}$ correspond to edges), so that it now starts and ends at Ω. Still, the new path (which we will continue to call γ_i) intersects $B_\Theta(1, \kappa^2 + 1)$ only at edges labelled by a single label.

Now we construct a new path γ' as follows. Suppose that k_i is such that each edge in $\gamma_i \cap B_\Theta(1, \kappa^2 + 1)$ has label k_i. We put

$$\gamma_i' = \rho_{k_i}(\gamma_i)$$

Note that as ρ_{k_i} is a retraction onto Ω, and the endpoints of γ_i are in Ω, the path γ_i' has the same endpoints as γ_i. Put

$$\gamma' = \gamma_1' * \cdots * \gamma_m'$$

This is now a path joining x to y in Ω, and thus contains an edge

$$e \in B_\Omega(1, 1)$$

There exists an edge f in some γ_i, such that e lies in the image of f under the map ρ_{k_i} that we applied to γ_i. Since each ρ_k is an (κ, κ)-quasi-isometry, the edge f lies within $B_\Theta(1, \kappa^2 + 1)$. But then $\rho_{k_i}(f)$ is a path the edges of whom are never labelled by 0, and so in particular $e \notin E$, a contradiction.

We now apply Theorem 2.7, taking $\mathcal{P}$ to be the set of edges labelled by 0. Let C denote the cut we obtain, and let F denote its boundary.

To apply Theorem 2.8 we need to only show that for some $g \in \overline{A}$ we have $g.C$ properly contained in C or C^*. Since C^* is infinite, it contains an element $g \in \overline{A}$ such that $g.F \neq F$. Taking such a g, we see that either $g.C$ is properly contained in C^* (in which case we are done), or C is properly contained in $g.C$. In the latter case we have $g^{-1}.C \subset C$. We have thus verified all the hypotheses of Theorem 2.8.

Since the boundary F of the final cut C is labelled by 0, upon removal of the open edges in $\overline{A}.F$, the connected component containing 1_i contains the entire subgroup A_i, since vertices corresponding to elements of this subgroup are connected to 1_i by paths labelled by i. Thus A_i is a subgroup of a conjugate of a vertex group, and so it fixes a vertex in the associated action on a tree.

It remains to show the triviality of edge stabilisers in A. In fact we will show that no non-trivial subgroup $G < A$ fixes a narrow cut in Θ with boundary consisting only of edges labelled by 0. To this end, let C be such a cut, and F the set of edges forming the boundary of C.

We begin by considering the subgraph Ω. Let Γ be an infinite component of $\Omega - F$, and $h \in H$ be arbitrary. There are infinitely many vertices v in Γ such that no edge emanating from v lies in F (as the latter is finite). Take one such vertex, and consider an edge e in its star which corresponds to right multiplication with h. Since h normalises A, it in fact connects Ω to $h.\Omega$. On the other hand, there can be only a single component of $h.\Omega - F$ which is connected to Γ as the cut C is narrow: otherwise the components of $h.\Omega - F$ would lie in the same component of $\Theta - F$, and F would fail the definition of a boundary of a cut.

In summary, we have shown, that for each h, each infinite component Γ of $\Omega - F$ is connected (via an edge corresponding to right multiplication by h) to a unique infinite component of $h.\Omega - F$. In other words, infinite components of $\Omega - F$ and $h.\Omega - F$ are in bijection to each other, where the bijection identifies components which are connected in $\Theta - F$.

Now, we can think of Ω as the Bass-Serre tree for the splitting of A, whose vertices have been "blown up" to Cayley graphs of the subgroups A_i. In particular, each edge labelled by 0 disconnects Ω. This implies that $\Omega - F$, and hence each $h.\Omega - F$, has exactly two components, both of which are infinite. Namely, if $\Omega - F$

would have more than two infinite components, or just a single one, the same would be true for $\Theta - F$, violating narrowness of the cut F. It also implies that $F \cap h.\Omega$ consists of exactly one edge for each h. Since A acts freely on Ω, this implies the final claim of the theorem. □

3. Blow-ups

We make the convention that graphs of groups are always connected unless explicitly stated otherwise.

Proposition 3.1 (Blow-up with finite edge groups)**.** *Let G be a graph of groups with finite edge groups. For each vertex v suppose that the associated vertex group G_v acts on a connected space X_v in such a way that each finite subgroup of G_v fixes a point of X_v. Then there exists a connected space Y on which $\pi_1(G)$ acts, satisfying the following:*

1. *there is a $\pi_1(G)$-equivariant map $\pi\colon Y \to \widetilde{G}$;*
2. *if w is a vertex of $\widetilde{G}$ fixed by G_v, then $\pi^{-1}(w)$ is G_v-equivariantly isometric to X_v;*
3. *every finite subgroup of G fixes a point of Y.*

Moreover, when the spaces X_v are complete and CAT(0) then Y is a complete CAT(0) space.

Proof. Recall that the vertices of $\widetilde{G}$ are left cosets of the vertex groups G_v of G; for each vertex w we pick an element $z_w \in G$ to be a coset representative of such a coset.

We will build the space Y in two steps. First, we construct the preimage under π of the vertices of $\widetilde{G}$, and call it V. We define V to be the disjoint union of spaces X_w, where w runs over the vertices of $\widetilde{G}$, and X_w is an isometric copy of X_v, where v is the image of w under the quotient map $\widetilde{G} \to G$. We construct $\pi\colon V \to \widetilde{G}$ by declaring $\pi(X_w) = \{w\}$.

We now construct an action of $A = \pi_1(G)$ on V. Let us take $X_w \subset V$, and let $a \in A$. Let $u = a.w$, and note that its image in G is still v. The action of a on V will take X_w to X_u; using the identifications $X_w \simeq X_v \simeq X_u$ we only need to say how a is supposed to act on X_v, and here it acts as $z_w^{-1}a$.

We now construct the space Y by adding edges to V.

Let e be an edge of $\widetilde{G}$ with terminal endpoint w and initial endpoint u. Let X_e denote a copy of the unit interval. Now G_e is a finite subgroup of G_w, and so fixes a point in X_w seen as a subset of V. We glue the endpoint 1 of X_e to this point. Analogously, we glue the endpoint 0 to a point in X_u. Now, using the action of A, we equivariantly glue all the endpoints of the edges in the A-orbit of e. We proceed this way for all (geometric) edges. Note that this construction allows us to extend the definition of π.

When all the vertex spaces are complete CAT(0), it is clear that so is Y. □

Remark 3.2. Suppose that the spaces X_v in the above proposition are trees. Then the resulting space Y is a tree, and the quotient graph of groups is obtained from G by replacing v by the quotient graph of groups $X/\!\!/G_v$.

We will refer to the above construction as *blowing up* G by the spaces X_v. We warn the reader that our notion of a blow-up is not standard terminology (and has nothing to do with blow-ups in other fields).

When dealing with limit groups, we will need a more powerful version of a blow-up. We will use a method by Sam Brown, essentially following [**Bro2**, Theorem 3.1]; to this end let us start with a number of definitions and standard facts.

Definition 3.3. An *n-simplex of type M_κ* is the convex hull of $n+1$ points in general position lying in the n-dimensional model space M_κ of curvature κ, as defined in [**BH**].

An M_κ-simplicial complex K is a simplicial complex in which each simplex is endowed with the metric of a simplex of type M_κ, and the face inclusions are isometries.

Note that we will be interested in the case of $n = 2$ and negative κ, where the model space M_κ is just a suitably rescaled hyperbolic plane.

Definition 3.4. Let K be a M_κ-simplicial complex of dimension at most 2. The *link* of a vertex v is a metric graph whose vertices are edges of K incident at v, and edges are 2-simplices of K containing v. Inclusion of edges into simplices in X induces the inclusion of vertices into edges in the link. The length of an edge in the link is equal to the angle the edges corresponding to its endpoints make in the simplex.

Let us state a version of Gromov's link condition adapted to our setting.

Theorem 3.5 (Gromov's link condition [**BH**, Theorem II.5.2])**.** *Let K be a M_κ-simplicial complex of dimension at most 2, endowed with a cocompact simplicial isometric action. Then K is a locally* $\mathrm{CAT}(\kappa)$ *space if and only if the link of each vertex in K is* $\mathrm{CAT}(1)$.

Of course, for a graph being $\mathrm{CAT}(1)$ is equivalent to having no non-trivial simple loop of length less than 2π.

Lemma 3.6 ([**Bro2**, Lemma 2.29])**.** *For any $0 < \theta < \pi$ and any A, C with $C > A > 0$, there exists $k < 0$ and a locally* $\mathrm{CAT}(k)$ *M_k-simplicial annulus with one locally geodesic boundary component of length A, and one boundary component of length C which is locally geodesic everywhere except for one point where it subtends an angle greater than θ.*

Lemma 3.7. *Let Z be an infinite virtually cyclic group. Any two cocompact isometric actions on $\mathbb{R}$ have the same kernel, and the quotient of Z by the kernel is isomorphic to either $\mathbb{Z}$ or the infinite dihedral group D_∞.*

Proof. Clearly both actions on $\mathbb{R}$ can be made into actions on 2-regular trees with a single edge orbit and no edge inversions; each such action gives us a decomposition of $\mathbb{Z}$ into a graph of finite groups, where the kernel of the action is the unique edge group, and the quotient is as claimed. Let G_1 and G_2 denote the graphs of groups, and K_1 and K_2 denote the respective edge groups.

Suppose that one of the graphs, say G_1, has only one vertex. Then K_1 is also equal to the vertex group, and we have $K_2 \leqslant K_1$, since any finite group acting on a tree has a fixed point. If G_2 also has a single vertex than $K_1 \leqslant K_2$ by the same argument and we are done. Otherwise $Z/K_1 \simeq \mathbb{Z}$ is a quotient of $Z/K_2 \simeq D_\infty$, which is impossible.

Now suppose that both G_1 and G_2 have two vertices each. Let G_v be a vertex group of G_1. Arguing as before we see that it fixes a point in the action of Z on $\widetilde{G}_2$, and so some index 2 subgroup of G_v fixes an edge. Thus $K_1 \cap K_2$ is a subgroup of K_1 of index at most two. If the index is two, then the image of K_1 in $Z/K_2 \simeq D_\infty$ is a normal subgroup of cardinality 2. But D_∞ does not have such subgroups, and so $K_1 \leqslant K_2$. By symmetry $K_2 \leqslant K_1$ and we are done. □

Let us record the following standard fact.

Lemma 3.8. *Let Z be an infinite virtually cyclic group acting properly by semisimple isometries on a complete* $\mathrm{CAT}(0)$ *space X. Then Z fixes an image of a geodesic in X (called an* axis*).*

For the purpose of the next proposition, let us introduce some notation.

Definition 3.9. A CAT(-1) M_{-1}-simplicial complex of dimension at most 2 with finitely many isometry classes of simplices will be called *useful.*

Proposition 3.10 (Blow-up with virtually cyclic edge groups)**.** *Suppose that $\kappa \in \{0, -1\}$. Let G be a finite graph of groups with virtually cyclic edge groups. For each vertex v suppose that the associated vertex group G_v acts properly on a connected complete* CAT(κ) *simplicial complex X_v by semi-simple isometries. Suppose further that*

(A1) there exists an orientation of geometric edges of G such that the initial vertex of every edge e is useful*: is it is a vertex u with X_u useful; and*

(A2) when X_u is useful and $e_1, \dots, e_n$ are all the edges of G incident at u carrying an infinite edge group, then the axes preserved by $g^{-1}X_{e_i}g$ with $i \in \{1, \dots, n\}$ and $g_i \notin X_{e_i}$ can be taken to be simplicial and pairwise transverse.

Then there exists a connected complete CAT(κ) *space Y on which $\pi_1(G)$ acts, satisfying the following:*

(1) there is a $\pi_1(G)$-equivariant map $\pi \colon Y \to \widetilde{G}$;

(2) if w is a vertex of $\widetilde{G}$ fixed by G_v, then $\pi^{-1}(w)$ is G_v-equivariantly isometric to X_v.

Proof. We will proceed exactly as in the proof of Proposition 3.1, with two exceptions: firstly, we will rescale the spaces X_v before we start the construction; secondly, we will need to deal with infinite virtually cyclic edge groups. Let us first explain how to deal with the infinite edge groups, and then it will become apparent how we need to rescale the useful spaces.

Let e be an oriented edge of $\widetilde{G}$ with infinite stabiliser G_e (note that this is a slight abuse of notation, as we usually reserve G_e to be an edge group in G rather than a stabiliser in $\widetilde{G}$). The group is virtually cyclic, and so, by Lemma 3.8, fixes an axis in each of the vertex spaces corresponding to the endpoints of e (it could of course be two axes in a single space, if e is a loop). The actions on these axes are equivariant by Lemma 3.7, and the only difference is the length of the quotient of the axis by G_e; we will denote the two lengths by λ_e^+ and λ_e^-, where λ_e^+ is the amount by which G_e translates the axis corresponding to the terminus of e, and λ_e^- to the origin.

We claim that we can rescale the spaces X_v and orient the geometric edges so that for any edge e with infinite stabiliser we have the initial vertex of e useful and $\lambda_e^+ \leqslant \lambda_e^-$. Let us assume that we have already performed a suitable rescaling – we will come back to it at the end of the proof.

Let u denote the initial (useful) endpoint of e; let w denote the other endpoint of e. We replace each 2-dimensional simplex in X_u by the comparison simplex of type $M_{-\frac{1}{2}}$ – note that, in particular, this does not affect the metric on the 1-skeleton of X_u, and hence does not affect the constant λ_e^-. Let $\widehat{X}_u$ denote the resulting space.

In $\widehat{X}_u$ we have generated, in Brown's terminology, an *excess angle* δ (depending on u), that is in the link of any vertex x in $\widehat{X}_u$ the distance between any two points which were of distance at least π in the link of x in X_u is at least $\pi + 2\delta$ in the link in $\widehat{X}_u$. By possibly decreasing δ, we may assume that $\delta < \frac{\pi}{3}$, and that the distance between any two distinct vertices in a link of a vertex in $\widehat{X}_u$ is at least δ (this is possible since there are only finitely many different isometry types of simplices in X_u, and so in $\widehat{X}_u$). We still have G_u acting on $\widehat{X}_u$ simplicially and isometrically.

Suppose that $\lambda_e^+ = \lambda_e^-$. Then we take X_e to be a flat strip $[0,1] \times \mathbb{R}$ on which G_e acts by translating the $\mathbb{R}$ factor so that the quotient is isometric to $[0,1] \times \mathbb{R}/\lambda_e^+\mathbb{Z}$.

If $\lambda_e^+ \neq \lambda_e^-$ then we take X_e to be the universal cover of an annulus from Lemma 3.6 with boundary curves of length λ_e^+ and λ_e^-, and $\theta = \pi - \delta$. The space X_e is a $\mathrm{CAT}(k_e)$ M_{k_e}-simplicial complex for some $k_e < 0$.

We glue the preimage (in X_e) of each of the boundary curves to the corresponding axis of G_e, so that the gluing is an G_e-equivariant isometry. The gluing along the preimage of the shorter curve (or both curves if they are of equal length) proceeds along convex subspaces, and so if the vertex space was $\mathrm{CAT}(\mu)$ with $\mu \leqslant 0$, then the glued-up space is still locally $\mathrm{CAT}(\mu)$ along the axis of G_e.

The situation is different at the useful end: here we glue in along a non-convex curve. We claim that the resulting space is still locally $\mathrm{CAT}(k_e)$ along this geodesic. This follows from Gromov's link condition (Theorem 3.5), and the observation that in the link of any vertex of $\widehat{X}_u$ we introduced a single path (a *shortcut*) of length at least $\pi - \delta$ between vertices whose distance before the introduction of the shortcut was at least $\pi + 2\delta$. A simple closed curve which traverses both endpoints of the shortcut therefore had length at least $2\pi + 4\delta$ before introducing the shortcut, and thus still has length $\geqslant 2\pi + \delta$ afterwards. Thus there is still no non-trivial simple loop shorter than 2π.

We now use the action of $A = \pi_1(G)$ to equivariantly glue in copies of X_e for all edges in the orbit of e. We proceed in the same way for all the other (geometric) edges.

Now we need to look at the curvature. The useful spaces have all been altered to be $M_{-\frac{1}{2}}$-simplicial complexes, and so they are now $\mathrm{CAT}(-\frac{1}{2})$. If we had any $\mathrm{CAT}(0)$ vertex spaces, then they remain $\mathrm{CAT}(0)$. The universal covers X_e of annuli are $\mathrm{CAT}(k_e)$ with $k_e < 0$; the infinite strips are $\mathrm{CAT}(0)$. The gluing into the non-useful spaces did not disturb the curvature. A single gluing into a useful space did not disturb the curvature either, but the situation is more complicated when we glue more than one space X_e into a single $\widehat{X}_u$, since we could have introduced multiple shortcuts of length at least $\pi - \delta$ into a link of a single vertex. If a curve traverses one (or no) shortcut, then the argument given above shows that it has length at least 2π. If it traverses more more than 2, then (as $\delta < \pi/3$), it also has length $\geqslant 2\pi$. In the final case where it goes through exactly two, note that the endpoints of the shortcuts are all distinct by the transversality assumption (A2). Hence, by the choice of δ, any path connecting these endpoints has length $> \delta$, and so the total path has length $> 2(\pi - \delta) + 2\delta$ as well.

We conclude that our space Y is complete and $\mathrm{CAT}(k)$, where k is the maximum of the values k_e, κ and $-\frac{1}{2}$. When $\kappa = 0$ we have $k = 0$ and we are done. Otherwise, observing that we had only finitely many edges in G, we have $k < 0$, and so we can rescale Y to obtain a $\mathrm{CAT}(-1)$ space, as claimed.

We still need to explain how to rescale the vertex spaces. We order the vertices of the graph of groups G in some way, obtaining a list $v_1, \ldots, v_m$. The space X_{v_1} we do not rescale. Up to reorienting the geometric edges running from v_1 to itself we see that the constants λ_e^+ and λ_e^- for such edges satisfy $\lambda_e^+ \leqslant \lambda_e^+$.

We look at the full subgraph Γ of G spanned by the vertices $v_1, \ldots, v_i$. Inductively, we assume that the spaces corresponding to vertices in Γ have already been rescaled as required. Now we attach v_{i+1} to Γ, together with all edges connecting v_{i+1} to itself or Γ. If $X_{v_{i+1}}$ is not useful, then we have no edges of the latter type, and all edges connecting v_{i+1} to Γ are oriented towards v_{i+1}. Clearly we can rescale X_v to be sufficiently small so that the desired inequalities are satisfied (note that there are only finitely many edges to consider).

If $X_{v_{i+1}}$ is useful then we can reorient all edges connecting v_{i+1} to Γ so that they run away from v_{i+1}. Now we can make $X_{v_{i+1}}$ sufficiently big to satisfy the desired inequalities. We also reorient the edges connecting v_{i+1} to itself in a suitable manner. □

4. Relative Karrass–Pietrowski–Solitar theorem

The following theorem is a generalisation of a theorem of Karrass–Pietrowski–Solitar [**KPS**], which lies behind the Nielsen realisation theorem for free groups.

Theorem 4.1 (Relative Karrass–Pietrowski–Solitar theorem)**.** *Let*

$$\varphi\colon H \to \mathrm{Out}(A)$$

be a monomorphism with a finite domain, and let

$$A = A_1 * \cdots * A_n * B$$

be a decomposition preserved by H, with each A_i finitely generated, non-trivial, and B a (possibly trivial) finitely generated free group. Let $A_1, \ldots, A_m$ be the minimal factors. Then the associated extension $\overline{A}$ of A by H is isomorphic to the fundamental group of a finite graph of groups with finite edge groups, with m distinguished vertices $v_1, \ldots, v_m$, such that the vertex group associated to v_i is a conjugate of the extension $\overline{A_i}$ of A_i by $\mathrm{Stab}_H(i)$, and vertex groups associated to other vertices are finite.

Proof. The proof goes along precisely the same lines as the original proof of Karrass–Pietrowski–Solitar [**KPS**], with the exception that we use Relative Stallings' Theorem (Theorem 2.9) instead of the classical one.

We will prove the result by an induction on a *complexity* (n, f) where n is the number of factors A_i, and f is the rank of the free group B in the decomposition. We order the complexity lexicographically. The cases of complexity $(0, f)$ follow from the usual Nielsen realisation theorem for free groups (see Theorem 8.1).

Thus, for the inductive step, we assume a complexity (m, f) with $m > 0$. We begin by applying Theorem 2.9 to the finite extension $\overline{A}$. We obtain a graph of groups P with one edge and a finite edge group, such that each A_i lies up to conjugation in a vertex group, and no non-trivial subgroup of any factor A_i fixes an edge.

Let v be any vertex of $\widetilde{P}$. The group P_v is a finite extension of $A \cap P_v$ by a subgroup H_v of H. Let us look at the structure of $P_v \cap A$ more closely.

Consider the graph of groups associated to the product $A_1 * \ldots A_n * B$ and apply Kurosh's theorem [**Ser1**, Theorem I.14] to the subgroup $P_v \cap A$. We obtain that $P_v \cap A$ is a free product of groups of the form $P_v \cap xA_ix^{-1}$ for some $x \in A$, and a free group B'.

Let us suppose that the intersection $P_v \cap xA_ix^{-1}$ is nontrivial for some i and $x \in A$. This implies that a non-trivial subgroup G of A_i fixes the vertex $x^{-1}.v$. We also know that A_i fixes some vertex v_i in $\widetilde{P}$ by construction, and thus so does G. If $x^{-1}.v \neq v_i$, this would imply that G fixes an edge, which is impossible. Hence $v_i = x^{-1}.v$ and in particular we have that $xA_ix^{-1} \leqslant P_v$.

Now suppose that $P_v \cap yA_iy^{-1}$ is non-trivial for some other element $y \in A$. Then $x^{-1}.v = v_i = y^{-1}.v$, and so $xy^{-1} \in A \cap P_v$. This implies that the two free factors $P_v \cap xA_ix^{-1}$ and $P_v \cap yA_iy^{-1}$ of $P_v \cap A$ are conjugate inside the group, and so they must coincide.

We consider the action of A on the tree $\tilde{P}$, and conclude that A is equal to the fundamental group of the graph of groups $\tilde{P}/\!\!/A$. The discussion above shows that:

i) The stabilizer of a vertex $v \in \tilde{P}$ has the structure

$$P_v \cap A = x_{i(v,1)}A_{i(v,1)}x_{i(v,1)}^{-1} * \cdots * x_{i(v,k)}A_{i(v,k)}x_{i(v,k)}^{-1} * B'$$

where the indices $i(v,k)$ are all distinct, and B' is some free group.

ii) If a conjugate of A_i intersects some stabilizer of v non-trivially, then it stabilizes v.

iii) For each i there is exactly one vertex v so that a conjugate of A_i appears as $A_{i(v,l)}$ in the description above.

iv) The edge groups in $\tilde{P}/\!\!/A$ are trivial.

Since the splitting which P defines is non-trivial, the index of $P_v \cap A$ in $\overline{A}$ is infinite, and thus A is not a subgroup of P_v for any v.

Next, we aim to show that the complexity of each $P_v \cap A$ is strictly smaller than that of A. To begin, note that the only way that this could fail is if there is some vertex w so that

$$P_w \cap A = x_1 A_1 x_1^{-1} * \cdots * x_m A_m x_m^{-1} * B'$$

for B' a free group. Since all edge groups in $\tilde{P}/\!\!/A$ are trivial, A is obtained from $P_w \cap A$ by a free product with a free group. Such an operation cannot decrease the rank of B', and in fact increases it unless the free product is trivial. But in the latter case we would have $P_w \cap A = A$, which is impossible.

We have thus shown that each P_v is an extension

$$P_v \cap A \to P_v \to H_v$$

where H_v is a subgroup of H, the group $P_v \cap A$ decomposes in a way which is preserved by H_v, and its complexity is smaller than that of A. Therefore the group P_v satisfies the assumption of the inductive hypothesis.

We now use Proposition 3.1 (together with the remark following it) to construct a new graph of groups Q, by blowing P up at u by the result of the theorem applied to P_u, with u varying over some chosen lifts of the vertices of P.

By construction, Q is a finite graph of groups with finite edge groups, and the fundamental group of Q is indeed $\overline{A}$. Also, Q inherits distinguished vertices from the graphs of groups we blew up with. Thus, Q is as required in the assertion of our theorem, with two possible exceptions.

Firstly, it might have too many distinguished vertices. This would happen if for some i and j we have A_i and A_j both being subgroups of, say, P_v, which are conjugate in $\overline{A}$ but not in P_v. Let $h \in \overline{A}$ be an element such that $hA_ih^{-1} = A_j$. Since both A_i and A_j fix only one vertex, and this vertex is v, we must have $h \in \mathbb{P}_v$, and so A_i and A_j are conjugate inside P_v.

Secondly, it could be that the finite extensions of A_i we obtain as vertex groups are not extensions by $\mathrm{Stab}_H(i)$. This would happen if $\mathrm{Stab}_H(i)$ is not a subgroup of H_v. Let us take $h \in \overline{A}$ in the preimage of $\mathrm{Stab}_H(i)$, such that $hA_ih^{-1} = A_i$. Then in the action on $\widetilde{P}$ the element h takes a vertex fixed by A_i to another such; if these were different, then A_i would fix an edge, which is impossible. Thus h fixes the same vertex as A_i. This finishes the proof. □

5. Fixed points in the graph of relative free splittings

Consider a free product decomposition

$$A = A_1 * \cdots * A_n * B$$

with B a finitely generated free group. Handel and Mosher [**HM**] (see also the work of Horbez [**Hor**]) defined a *graph of relative free splittings* $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$ associated to such a decomposition. Its vertices are finite non-trivial graphs of groups with trivial edge groups, and such that each A_i is contained in a conjugate of a vertex group; two such graphs of groups define the same vertex when the associated universal covers are A-equivariantly isometric. Two vertices are connected by an edge if and only if the graphs of groups admit a common refinement.

In their article, Handel and Mosher prove that $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$ is connected and Gromov hyperbolic [**HM**, Theorem 1.1].

Observe that the subgroup $\mathrm{Out}(A, \{A_1, \ldots, A_n\})$ of $\mathrm{Out}(A)$ consisting of those outer automorphisms of A which preserve the decomposition

$$A = A_1 * \cdots * A_n * B$$

acts on this graph. We offer the following fixed point theorem for this action on $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$.

Corollary 5.1. *Let $H \leqslant \mathrm{Out}(A, \{A_1, \ldots, A_n\})$ be a finite subgroup, and suppose that the factors A_i are finitely generated. Then H fixes a point in the free-splitting graph $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$.*

Proof. Theorem 4.1 gives us an action of the extension $\overline{A}$ on a tree T; in particular A acts on this tree, and this action satisfies the definition of a vertex in $\mathcal{FS}(A, \{A_1, \ldots, A_n\})$. Since the whole of $\overline{A}$ acts on T, every outer automorphism in H fixes this vertex. □

6. Fixed points in the outer space of a free product

Take any finitely generated group A, and consider its *Grushko decomposition*, that is a free splitting

$$A = A_1 * \cdots * A_n * B$$

where B is a finitely generated free group, and each group A_i is finitely generated and freely indecomposible, that is it cannot act on a tree without a global fixed point (note that $\mathbb{Z}$ is not freely indecomposible in this sense).

Grushko's Theorem [**Gru**] tells us that such a decomposition is essentially unique; more precisely, if

$$A = A'_1 * \cdots * A'_m * B'$$

is another such decomposition, then $B \cong B'$, $m = n$, and there is a permutation β of the set $\{1, \ldots, n\}$ such that A_i is conjugate to $A'_{\beta(i)}$. In particular, this implies that the decomposition

$$A = A_1 * \cdots * A_n * B$$

is preserved in our sense by every outer automorphism of A.

In [**GL**] Guirardel and Levitt introduced $P\mathcal{O}$, the (projectivised) outer space of a free product. It is a simplicial complex whose vertices are equivalence classes of pairs (G, ι), where:

1. G is a finite graph of groups with trivial edge groups;
2. edges of G are given positive lengths;
3. for every $i \in \{1, \ldots, n\}$, there is a unique vertex v_i in G such that the vertex group G_{v_i} is conjugate to A_i;
4. all other vertices have trivial vertex groups;
5. every leaf of G is one of the vertices $\{v_1, \ldots, v_n\}$;
6. $\iota\colon \pi_1(G) \to A$ is an isomorphism.

The equivalence relation is given by postcomposing ι with an inner automorphism of A, and by multiplying the lengths of all edges of G by a positive constant. We also consider two pairs G, ι and G', ι' equivalent if there exists an isometry $\psi\colon G \to G'$ such that $\iota = \iota' \circ \psi$.

Because of the essential uniqueness of the Grushko decomposition, the group $\mathrm{Out}(A)$ acts on $P\mathcal{O}$ by postcomposing the marking ι. We offer the following result for this action.

Corollary 6.1. *Let A be a finitely generated group, and let $H \leqslant \mathrm{Out}(A)$ be a finite subgroup. Then H fixes a vertex in $\mathcal{PO}$.*

Proof. Theorem 4.1 gives us an action of the extension $\overline{A}$ on a tree T, and we may assume that this action is minimal; in particular A acts on this tree, and this action satisfies the definition of a vertex in $\mathcal{PO}$ (with all edge lengths equal to 1). Since the whole of $\overline{A}$ acts on T, every outer automorphism in H fixes this vertex. □

Note that $\mathcal{PO}$ has been shown in [**GL**, Theorem 4.2, Corollary 4.4] to be contractible.

7. Relative Nielsen realisation

In this section we use Theorem 4.1 to prove relative Nielsen Realisation for free products. To do this we need to formalise the notion of a marking of a space.

Definition 7.1. We say that a path-connected topological space X with a universal covering $\widetilde{X}$ is *marked* by a group A if and only if it comes equipped with an isomorphism between A and the group of deck transformations of $\widetilde{X}$.

Remark 7.2. Given a space X marked by a group A, we obtain an isomorphism $A \cong \pi_1(X,p)$ by choosing a basepoint $\widetilde{p} \in \widetilde{X}$ (where p denotes its projection in X).

Conversely, an isomorphism $A \cong \pi_1(X,p)$ together with a choice of a lift $\widetilde{p} \in \widetilde{X}$ of p determines the marking in the sense of the previous definition.

Definition 7.3. Suppose that we are given an embedding $\pi_1(X) \hookrightarrow \pi_1(Y)$ of fundamental groups of two path-connected spaces X and Y, both marked. A map $\iota\colon X \to Y$ is said to *respect the markings via the map* $\widetilde{\iota}$ if and only if $\widetilde{\iota}\colon \widetilde{X} \to \widetilde{Y}$ is $\pi_1(X)$-equivariant (with respect to the given embedding $\pi_1(X) \hookrightarrow \pi_1(Y)$), and satisfies the commutative diagram

$$\begin{array}{ccc} \widetilde{X} & \xrightarrow{\widetilde{\iota}} & \widetilde{Y} \\ \downarrow & & \downarrow \\ X & \xrightarrow{\iota} & Y \end{array}$$

We say that ι *respects the markings* if and only if such an $\widetilde{\iota}$ exists.

Suppose that we have a metric space X marked by a group A, and a group H acting on X. Of course such a setup yields the induced action $H \to \mathrm{Out}(A)$, but in fact it does more: it gives us an extension

$$1 \to A \to \overline{A} \to H \to 1$$

where $\overline{A}$ is the group of all lifts of elements of H to automorphisms of the universal covering $\widetilde{X}$ of X.

Definition 7.4. Suppose that we are given a group extension

$$A \to \overline{H} \to H$$

We say that an action $\varphi\colon H \to \mathrm{Isom}(X)$ of H on a metric space X *realises the extension* $\overline{H}$ if and only if X is marked by A, and the extension

$$\pi_1(X) \to G \to H$$

induced by φ fits into the commutative diagram

$$\begin{array}{ccccc} A & \longrightarrow & \overline{H} & \longrightarrow & H \\ \downarrow \simeq & & \downarrow \simeq & & \| \\ \pi_1(X) & \longrightarrow & G & \longrightarrow & H \end{array}$$

When A is centre-free, and we are given an embedding $H \leqslant \mathrm{Out}(A)$, we say that an action φ as before *realises the action* $H \to \mathrm{Out}(A)$ if and only if it realises the corresponding extension.

Now we are ready to state the relative Nielsen Realisation theorem for free products.

Theorem 7.5 (Relative Nielsen Realisation)**.** *Let $\varphi\colon H \to \mathrm{Out}(A)$ be a homomorphism with a finite domain, and let*

$$A = A_1 * \cdots * A_n * B$$

be a decomposition preserved by H, with each A_i finitely generated, and B a (possibly trivial) finitely generated free group. Let $A_1, \ldots, A_m$ be the minimal factors.

Suppose that for each $i \in \{1, \ldots, m\}$ we are given a complete NPC space X_i marked by A_i, on which $\mathrm{Stab}_i(H)$ acts in such a way that the associated extension of A_i by $\mathrm{Stab}_H(i)$ is isomorphic (as an extension) to the extension $\overline{A}_i$ coming from $\overline{A}$. Then there exists a complete NPC space X realising the action φ, and such that for each $i \in \{1, \ldots, m\}$ we have a $\mathrm{Stab}_H(i)$-equivariant embedding $\iota_i\colon X_i \to X$ which preserves the marking.

Moreover, the images of the spaces X_i are disjoint, and collapsing each X_i and its images under the action of H individually to a point yields a graph with fundamental group abstractly isomorphic to the free group B.

As outlined in the introduction, the proof is very similar to the classical proof of Nielsen realisation, with our new relative Stallings' and Karrass–Pietrowski–Solitar theorems in place of the classical ones.

Proof. When φ is injective we first apply Theorem 4.1 to obtain a graph of groups G, and then use Proposition 3.1 and blow up each vertex of $\widetilde{G}$ by the appropriate $\widetilde{X_i}$; we call the resulting space $\widetilde{X}$. The space X is obtained by taking the quotient of the action of A on $\widetilde{X}$.

If φ is not injective, then we consider the induced map

$$H/\ker\varphi \to \mathrm{Out}(A)$$

apply the previous paragraph, and declare H to act on the resulting space with $\ker\varphi$ in the kernel. □

Remark 7.6. In the above theorem the hypothesis on the spaces X_i being complete and NPC can be replaced by the condition that they are semi-locally simply connected, and any finite group acting on their universal covering fixes at least one point.

Remark 7.7. On the other hand, when we strengthen the hypothesis and require the spaces X_i to be NPC cube complexes (with the actions of our finite groups preserving the combinatorial structure), then we may arrange for X to also be a cube complex. When constructing the blow ups, we may always take the fixed points of the finite groups to be midpoints of cubes, and then X is naturally a cube complex, when we take the cubical barycentric subdivisions of the complexes X_i instead of the original cube complexes X_i.

Remark 7.8. In [**HOP**] Osajda, Przytycki and the first-named author develop a more topological approach to Nielsen realisation and the Karrass–Pietrowski–Solitar theorem. In that article, Nielsen realisation is shown first, using *dismantlability* of the sphere graph (or free splitting graph) of a free group, and the Karrass–Pietrowski–Solitar theorem then follows as a consequence.

The relative Nielsen realisation theorem with all free factors A_i being finitely generated free groups is a fairly quick consequence of the methods developed in

[**HOP**] – however, the more general version proved here cannot at the current time be shown using the methods of [**HOP**]: to the authors knowledge no analogue of the sphere graph exhibits suitable properties. It would be an interesting problem to find a "splitting graph" for free products which has dismantling properties analogous to the ones shown in [**HOP**] to hold for arc, sphere and disk graphs.

8. Nielsen realisation for limit groups

In the last section we are going to prove a Nielsen realisation statement for limit groups. It relies on the three classical Nielsen realisation theorems:

Theorem 8.1 ([**Cul, Khr1, Zim1**])**.** *Let H be a finite subgroup of* $\mathrm{Out}(F_n)$*, where F_n denotes the free group of rank n. There exists a finite graph X realising the given action $H < \mathrm{Out}(F_n)$.*

Theorem 8.2. *Let*

$$\mathbb{Z}^n \to \overline{H} \to H$$

be a finite extension of $\mathbb{Z}^n$. There exists a metric n-torus X realising this extension.

Theorem 8.3 (Kerckhoff [**Ker1, Ker2**])**.** *Let H be a finite subgroup of* $\mathrm{Out}(\pi_1(\Sigma))$ *where Σ is a closed surface of genus at least 2. There exists a hyperbolic metric on Σ such that Σ endowed with this metric realises the given action $H < \mathrm{Out}(\pi_1(\Sigma))$.*

Now we are ready to proceed with limit groups.

Definition 8.4. A group A is called *fully residually free* if and only if for any finite subset $\{a_1, \ldots, a_n\} \subseteq A \smallsetminus \{1\}$ there exists a free quotient $q\colon A \to F$ such that $q(a_i) \neq 1$ for each i.

A finitely generated fully residually free group is called a *limit group.*

Note that the definition immediately implies that limit groups are torsion free.

The crucial property of one-ended limit groups is that they admit JSJ-decompositions invariant under automorphisms.

Theorem 8.5 (Bumagin–Kharlampovich–Myasnikov [**BKM**, Theorem 3.13 and Lemma 3.16])**.** *Let A be a one-ended limit group. Then there exists a finite graph of groups G with all edge groups cyclic, each vertex group being finitely generated free, finitely generated free abelian, or the fundamental group of a closed surface, such that $\pi_1(G) = A$ and such that any automorphism φ of A induces an A-equivariant isometry ψ of $\widetilde{G}$ such that the following diagram commutes*

$$\begin{array}{ccc} A & \longrightarrow & \mathrm{Isom}(\widetilde{G}) \\ \Big\downarrow{\scriptstyle \varphi} & & \Big\downarrow{\scriptstyle c_\psi} \\ A & \longrightarrow & \mathrm{Isom}(\widetilde{G}) \end{array}$$

where c_ψ denotes conjugation by ψ.

Moreover, every maximal abelian subgroup of A is conjugate to a vertex group of G, and every edge in G connects a vertex carrying a maximal abelian subgroup to a vertex carrying a non-abelian free group or a surface group.

We will refer to the graph of groups G as the *canonical JSJ decomposition.*

Definition 8.6. Recall that a subgroup $G \leqslant A$ is *malnormal* if and only if $a^{-1}Ga \cap G \neq \{1\}$ implies that $a \in G$ for every $a \in A$.

Following Brown, we say that a family of subgroups $G_1, \ldots, G_n$ of A is *malnormal* if and only if for every $a \in A$ we have that $a^{-1}G_i a \cap G_j \neq \{1\}$ implies that $i = j$ and $a \in G_i$.

We will use another property of limit groups and their canonical JSJ decompositions.

Proposition 8.7 ([**BKM**, Theorem 3.1(3),(4)])**.** *Let A be a limit group. Every non-trivial abelian subgroup of A lies in a unique maximal abelian subgroup, and every maximal abelian subgroup is malnormal.*

Corollary 8.8. *Let G_v be a non-abelian vertex group in a canonical JSJ decomposition of a one-ended limit group A. Then the edge groups carried by edges incident at v form a malnormal family in G_v.*

Proof. Let Z_1 and Z_2 denote two edge groups carried by distinct edges, e and e' say, incident at v. Without loss of generality we may assume that each of these groups is infinite cyclic. Suppose that there exists $g \in G_v$ and a non-trivial $z \in g^{-1}Z_1g \cap Z_2$. Each Z_i lies in a unique maximal subgroup M_i of A. But then the abelian subgroup generated by z lies in both M_1 and M_2, which forces $M_1 = M_2$ by uniqueness. Now

$$g^{-1}M_1g \cap M_1 \neq \{1\}$$

which implies that $g \in M_1$ (since M_1 is malnormal), and so $g^{-1}Z_1g = Z_1$, which in turn implies that $z \in Z_1 \cap Z_2$.

The edges e and e' form a loop in G, and so there is the corresponding element t in $A = \pi_1(G)$. Observe that t commutes with z, and so the group $\langle t, z\rangle$ must lie in M_1. But this is a contradiction, as t does not fix any vertices in $\widetilde{G}$. □

We are now going to use [**Bro2**, Lemma 2.31]; we are however going to break the argument of this lemma in two parts.

Lemma 8.9 (Brown)**.** *Let X be a connected M_{-1}-simplicial complex of dimension at most 2. Let $A = \pi_1(X)$, and suppose that we are given a malnormal family $\{G_1, \ldots, G_n\}$ of infinite cyclic subgroups of A. Then, after possibly subdividing X, each group G_i fixes a (simplicial) axis a_i in the universal cover of X, and the images in X of axes a_i and a_j for $i \neq j$ are distinct.*

In the second part of [**Bro2**, Lemma 2.31] we need to introduce an extra component, namely a simplicial action of a finite group H on X, which permutes the groups G_i up to conjugation.

Lemma 8.10 (Brown)**.** *Let X be a locally* $\mathrm{CAT}(-1)$ *connected finite M_{-1}-simplicial complex of dimension at most 2. Let $A = \pi_1(X)$, and suppose that we are given a family $\{c_1, \ldots, c_n\}$ of locally geodesic simplicial closed curves with images pairwise distinct. Suppose that we have a finite group H acting simplicially on X in a way preserving the images of the curves $c_1, \ldots, c_n$ setwise. Then there exists a locally* $\mathrm{CAT}(k)$ *2-dimensional finite simplicial complex X' of curvature k, with $k < 0$, with a transverse family of locally geodesic simplicial closed curves $\{c'_1, \ldots, c'_n\}$, such that X' is H-equivariantly homotopic to X, and the homotopy takes c'_i to c_i for each i.*

Sketch of proof. The proof of [**Bro2**, Lemma 2.31] goes through verbatim, with a slight modification; to explain the modification let us first briefly recount Brown's proof.

We start by finding two local geodesics, say c_1 and c_2, which contain segments whose union is a tripod – one arm of the tripod is shared by both segments. We glue in a *fin*, that is a 2-dimensional M_k-simplex, so that one side of the simplex is glued to the shared segment of the tripod, and another side is glued to another arm (the intersection of the two sides goes to the central vertex of the tripod). This way one of the curves, say c_1, is no longer locally geodesic, and we replace it by a locally geodesic curve identical to c_1 except that instead of travelling along two sides of the fin, it goes along the third side.

The problem is that after the gluing of a fin our space will usually not be locally CAT(-1) (the third side of the fin introduces a shortcut in the link of the central vertex of the tripod). To deal with this, we first replace simplices in X by the corresponding M_k-simplices, and this creates an excess angle δ (compare also the proof of Proposition 3.10). Then gluing in the fin does not affect the property of being locally CAT(k).

We glue such fins multiple times, until all local geodesics intersect transversely; after each gluing we perform a replacement of simplices to generate the excess angle.

Now let us describe what changes in our argument. When gluing in a fin, we need to do it H-equivariantly in the following sense: a fin is glued along two consecutive edges, say (e, e'), and H acts on pairs of consecutive edges. We thus glue in one fin for each coset of the stabiliser of (e, e') in H. This way, when we introduce shortcuts in a link of a vertex, no two points are joined by more than one shortcut. Since we are gluing multiple fins simultaneously, we need to make the angle $\pi - \delta$ sufficiently close to π. ☐

Note that when we say that the family $\{c'_1, \ldots, c'_n\}$ is transverse, we mean that each curve c'_i intersects transversely with the other curves and itself.

Theorem 8.11. *Let A be a limit group, and let*

$$A \to \overline{H} \to H$$

be an extension of A by a finite group H. Then there exists a complete locally CAT(κ) space X realising the extension $\overline{H}$, where $\kappa = -1$ when A is hyperbolic, and $\kappa = 0$ otherwise.

Proof. We are first going to assume that A, and so $\overline{H}$, are one-ended. We apply Theorem 8.5 and obtain a connected graph of groups G with

$$\pi_1(G) = A$$

for which we can extend the natural action of A on $\widetilde{G}$ to an action of $\overline{H}$. Taking the quotient by $\overline{H}$ we obtain a new graph of groups Γ with

$$\pi_1(\Gamma) = \overline{H}$$

The edge groups of Γ are virtually cyclic, and vertices are finite extensions of finitely generated free or free-abelian groups, or finite extensions of fundamental groups of closed surfaces.

Using Theorems 8.1 to 8.3, for each vertex group Γ_v we construct a complete NPC space X_v marked by $A_v = A \cap \Gamma_v$, on which Γ_v / A_v acts in such a way that the induced extension is isomorphic to Γ_v. The space $\widetilde{X}_v$ is isometric either to a Euclidean space, the hyperbolic plane, or a tree, and the group A_v acts by deck transformations upon it.

When $\widetilde{X}_u$ is the hyperbolic space, we can triangulate it Γ_u-equivariantly, and so $\widetilde{X}_u$ and X_u have the structure of 2-dimensional finite M_{-1}-simplicial complexes. Moreover, we can triangulate it in such a way that each axis fixed by an infinite cyclic group carried by an edge incident at u is also simplicial. Observe that Γ_v / A_v permutes these axes, and so each of the corresponding edge groups in Γ preserves such an axis as well.

Now we apply Lemma 8.9 and conclude that distinct axes do not coincide. Thus we may use Lemma 8.10, and replace X_u by a new CAT(-1) M_{-1}-simplicial complex (after rescaling) of dimension at most 2, which has only finitely many isometry classes of simplices, and in which our axes intersects each other and themselves transversely.

We argue in the analogous manner for spaces X_u which are trees.

Observing that each infinite edge group preserves an axis in each of the relevant vertex spaces by Lemma 3.8, we apply Proposition 3.10, and take the resulting space to be X.

Let us now consider a limit group A which is not one-ended. In this case we apply the classical version of Stalling theorem to $\overline{H}$, and split it over a finite group. We will in fact apply the theorem multiple times, so that we obtain a finite graph of groups B with finite edge groups, with all vertex groups finitely generated and one-ended, and $\pi_1(G) = A$; the fact that we only have to apply the theorem finitely many times follows from finite presentability of A (see [**BKM**, Theorem 3.1(5)]) and Dunwoody's accessibility [**Dun2**].

The one-ended vertex groups are themselves finite extensions of limit groups, and so for each of them we have a connected metric space to act on by the first part of the current proof. We finish the argument by an application of Proposition 3.1 – the assumption on finite groups fixing points is satisfied since the vertex spaces are complete and CAT(0). □

CHAPTER V

On the smallest non-abelian quotient of $\mathrm{Aut}(F_n)$

This is joint work with Barbara Baumeister and Emilio Pierro.

We show that the smallest non-abelian quotient of $\mathrm{Aut}(F_n)$ is $\mathrm{PSL}_n(\mathbb{Z}/2\mathbb{Z}) = \mathrm{L}_n(2)$, thus confirming a conjecture of Mecchia–Zimmermann. In the course of the proof we give an exponential (in n) lower bound for the cardinality of a set on which $\mathrm{SAut}(F_n)$, the unique index 2 subgroup of $\mathrm{Aut}(F_n)$, can act non-trivially. We also offer new results on the representation theory of $\mathrm{SAut}(F_n)$ in small dimensions over small, positive characteristics, and on rigidity of maps from $\mathrm{SAut}(F_n)$ to finite groups of Lie type and algebraic groups in characteristic 2.

1. Introduction

Investigating finite quotients of mapping class groups and outer automorphism groups of free groups has a long history. The first fundamental result here is that groups in both classes are residually finite – this is due to Grossman [**Gro2**].

Once we know that the groups admit many finite quotients, we can start asking questions about the structure or size of such quotients. This is of course equivalent to studying normal subgroups of finite index in mapping class groups and $\mathrm{Out}(F_n)$.

When $n \geqslant 3$, the groups $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$ have unique subgroups of index 2, denoted respectively by $\mathrm{SOut}(F_n)$ and $\mathrm{SAut}(F_n)$. Both of these subgroups are perfect, and so the abelian quotients of $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$ are well understood. The situation for mapping class groups is very similar.

The simplest way of obtaining a non-abelian quotient of $\mathrm{Out}(F_n)$ or $\mathrm{Aut}(F_n)$ comes from observing that $\mathrm{Out}(F_n)$ acts on the abelianisation of F_n, that is $\mathbb{Z}^n$. In this way we obtain (surjective) maps

$$\mathrm{Aut}(F_n) \to \mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$$

The finite quotients of $\mathrm{GL}_n(\mathbb{Z})$ are controlled by the congruence subgroup property and are well understood. In particular, the smallest (in terms of cardinality) such quotient is $\mathrm{PSL}_n(\mathbb{Z}/2\mathbb{Z}) = \mathrm{L}_n(2)$, obtained by reducing $\mathbb{Z}$ modulo 2. According to a conjecture of Mecchia–Zimmermann [**MZ**], the group $\mathrm{L}_n(2)$ is the smallest non-abelian quotient of $\mathrm{Out}(F_n)$.

In [**MZ**] Mecchia and Zimmermann confirmed their conjecture for $n \in \{3, 4\}$. In this paper we prove it for all $n \geqslant 3$. In fact we prove more:

Theorem 9.1. *Let $n \geqslant 3$. Every non-trivial finite quotient of $\mathrm{SAut}(F_n)$ is either greater in cardinality than $\mathrm{L}_n(2)$, or isomorphic to $\mathrm{L}_n(2)$. Moreover, if the quotient is $\mathrm{L}_n(2)$, then the quotient map is the natural map postcomposed with an automorphism of $\mathrm{L}_n(2)$.*

The natural map $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ is obtained by acting on $\mathrm{H}_1(F_n; \mathbb{Z}/2\mathbb{Z})$.

Zimmermann [**Zim2**] also obtained a partial solution to the corresponding conjecture for mapping class groups, but in general the question of the cardinality of the smallest non-abelian quotients of mapping class groups is still open.

In order to determine the smallest non-trivial quotient of $\mathrm{SAut}(F_n)$ we can restrict our attention to the finite simple groups which, by the Classification of Finite Simple Groups (CFSG), fall into one of the following four families:

(1) the cyclic groups of prime order;
(2) the alternating groups A_n, for $n \geqslant 5$;
(3) the finite groups of Lie type, and;
(4) the 26 sporadic groups.

For the full statement of the CFSG we refer the reader to [**CCN**$^+$, Chapter 1] and for a more detailed exposition of the non-abelian finite simple groups to [**Wil**]. For the purpose of this paper, we further divide the finite groups of Lie type into the following two families:

(3C) the "classical groups": $\mathtt{A}_n$, ${}^2\mathtt{A}_n$, $\mathtt{B}_n$, $\mathtt{C}_n$, $\mathtt{D}_n$ and ${}^2\mathtt{D}_n$, and;
(3E) the "exceptional groups": ${}^2\mathtt{B}_2$, ${}^2\mathtt{G}_2$, ${}^2\mathtt{F}_4$, ${}^3\mathtt{D}_4$, ${}^2\mathtt{E}_6$, $\mathtt{G}_2$, $\mathtt{F}_4$, $\mathtt{E}_6$, $\mathtt{E}_7$ and $\mathtt{E}_8$.

We turn first to the alternating groups and prove the following.

Theorem 3.16. *Let $n \geqslant 3$. Any action of $\mathrm{SAut}(F_n)$ on a set with fewer than $k(n)$ elements is trivial, where*

$$k(n) = \begin{cases} 7 & n = 3 \\ 8 & n = 4 \\ 12 \quad \textit{if} & n = 5 \\ 14 & n = 6 \end{cases}$$

and $k(n) = \max_{r \leqslant \frac{n}{2}-3} \min\{2^{n-r-p(n)}, \binom{n}{r}\}$ for $n \geqslant 7$, where $p(n)$ equals 0 when n is odd and 1 when n is even.

The bound given above for $n \geqslant 7$ is somewhat mysterious; one can however easily see that (for large n) it is bounded below by $2^{\frac{n}{2}}$.

Note that, so far, no such result was available for $\mathrm{SAut}(F_n)$ (one could extract a bound of $2n$ from the work of Bridson–Vogtmann [**BV1**]). Clearly, the bounds given above give precisely the same bounds for $\mathrm{SOut}(F_n)$. In this context the best bound known so far was $\frac{1}{2}\binom{n+1}{2}$ (for $n \geqslant 6$) – see Corollary 2.29. It was obtained by an argument of representation theoretic flavour. The proof contained in the current paper is more direct.

The question of the smallest set on which $\mathrm{SAut}(F_n)$ or $\mathrm{SOut}(F_n)$ can act non-trivially remains open, but we do answer the question on the growth of the size of such a set with n – it is exponential. Note that the corresponding question for mapping class groups has been answered by Berrick–Gebhardt–Paris [**BGP**].

Let us remark here that $\mathrm{Out}(F_n)$ (and hence also $\mathrm{SAut}(F_n)$) has plenty of alternating quotients – indeed, it was shown by Gilman [**Gil**] that $\mathrm{Out}(F_n)$ is residually alternating.

We use the bounds above to improve on a previous result of Kielak on rigidity of outer actions of $\mathrm{Out}(F_n)$ on free groups (see Theorem 3.19 for details).

Following the alternating groups, we rule out the sporadic groups. It was observed by Bridson–Vogtman [**BV1**] that any quotient of $\mathrm{SAut}(F_n)$ which does not factor through $\mathrm{SL}_n(\mathbb{Z})$ must contain a subgroup isomorphic to

$$(\mathbb{Z}/2\mathbb{Z})^{n-1} \rtimes A_n = 2^{n-1} \rtimes A_n$$

(one can easily see this subgroup inside of $\mathrm{SAut}(F_n)$, as it acts on the n-rose, that is the bouquet of n circles). Thus, for large enough n, sporadic groups are never quotients of $\mathrm{SAut}(F_n)$, and therefore our proof (asymptotically) is not sensitive to whether the list of sporadic groups is really complete.

Finally we turn to the finite groups of Lie type. Our strategy differs depending on whether we are dealing with the classical or exceptional groups. The exceptional groups are handled in a similar fashion to the sporadic groups – this time we use an alternating subgroup A_{n+1} inside $\mathrm{SAut}(F_n)$, which rigidifies the group in a similar way as the subgroup $2^{n-1} \rtimes A_n$ did. The degrees of the largest alternating subgroups of exceptional groups of Lie type are known (and listed for example in [**LS**]); in particular this degree is bounded above by 17 across all such groups.

The most involved part of the paper deals with the classical groups. In characteristic 2 we use an inductive strategy, and prove

Theorem 6.9. *Let $n \geqslant 3$. Let K be a finite group of Lie type in characteristic 2 of twisted rank less than $n-1$, and let $\overline{K}$ be a reductive algebraic group over an algebraically closed field of characteristic 2 of rank less than $n-1$. Then any homomorphism $\mathrm{Aut}(F_n) \to K$ or $\mathrm{Aut}(F_n) \to \overline{K}$ has abelian image, and any homomorphism $\mathrm{SAut}(F_{n+1}) \to K$ or $\mathrm{SAut}(F_{n+1}) \to \overline{K}$ is trivial.*

Note that there are precisely two abelian quotients of $\mathrm{Aut}(F_n)$ (when $n \geqslant 3$), namely $\mathbb{Z}/2\mathbb{Z} = 2$ and the trivial group.

In odd characteristic we need to investigate the representation theory of $\mathrm{SAut}(F_n)$. We prove

Theorem 7.12. *Let $n \geqslant 8$. Every irreducible projective representation of $\mathrm{SAut}(F_n)$ of dimension less than $2n-4$ over a field of characteristic greater than 2 which does not factor through the natural map $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ has dimension $n+1$.*

Note that over a field of characteristic greater than $n+1$, every linear representation of $\mathrm{Out}(F_n)$ of dimension less than $\binom{n+1}{2}$ factors through the map $\mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$ mentioned above (see [**Kie1**, 3.13]). Representations of $\mathrm{SAut}(F_n)$ over characteristic other than 2 have also been studied by Varghese [**Var**].

The proof of the main result (Theorem 9.1) for $n \geqslant 8$ is uniform; the small values of n need special attention, and we deal with them at the end of the paper. We also need a number of computations comparing orders of various finite groups; these can be found in the appendix.

Acknowledgments: The authors are grateful to Alastair Litterick, Yuri Santos Rego and Stefan Witzel for many helpful conversations. The second- and third-named authors were supported by the SFB 701 of Bielefeld University.

2. Preliminaries

2.1. Notation. We recall a few conventions which we use frequently. When it is clear, the prime number p will denote the cyclic group of that order. Furthermore, the elementary abelian group of order p^n will be denoted as p^n. These conventions are standard in finite group theory. We also follow Artin's convention, and use $\mathrm{L}_n(q)$ to denote PSL_n over the field of cardinality q.

We conjugate on the right, and use the following commutator convention

$$[g,h] = ghg^{-1}h^{-1}$$

The abstract symmetric group of degree n is denoted by S_n. Given a set I, we define $\mathrm{Sym}(I)$ to be its symmetric group. We define A_n and $\mathrm{Alt}(I)$ in the analogous manner for the alternating groups.

We fix n and denote by N the set $\{1,\ldots,n\}$.

2.2. Some subgroups and elements of $\mathrm{Aut}(F_n)$. We start by fixing n and a free generating set $a_1,\ldots,a_n$ for the free group F_n. Recall that $N = \{1,\ldots,n\}$. We will abuse notation by writing $F_n = F(N)$, and given a subset $I \subseteq N$ we will write $F(I)$ for the subgroup of $F(N)$ generated by the elements a_i with $i \in I$.

For every $i, j \in N$, $i \neq j$, set

$$\rho_{ij}(a_k) = \begin{cases} a_i a_j & \text{if} \quad k = i \\ a_k & \text{otherwise} \end{cases}$$

$$\lambda_{ij}(a_k) = \begin{cases} a_j a_i & \text{if} \quad k = i \\ a_k & \text{otherwise} \end{cases}$$

$$\sigma_{ij}(a_k) = \begin{cases} a_j & \quad k = i \\ a_i & \text{if} \quad k = j \\ a_k & \quad k \notin \{i, j\} \end{cases}$$

$$\sigma_{i(n+1)}(a_k) = \begin{cases} {a_i}^{-1} & \text{if} \quad k = i \\ a_k {a_i}^{-1} & \text{otherwise} \end{cases}$$

$$\epsilon_i(a_k) = \begin{cases} {a_i}^{-1} & \text{if} \quad k = i \\ a_k & \text{otherwise} \end{cases}$$

$$\delta(a_k) = \quad {a_k}^{-1} \quad \text{for every} \quad k$$

All of the endomorphisms of F_n defined above are in fact elements of $\mathrm{Aut}(F_n)$. The elements ρ_{ij} are the *right transvections*, the elements λ_{ij} are the *left transvections*, and the set of all transvections generates $\mathrm{SAut}(F_n)$.

The involutions ϵ_i pairwise commute, and hence generate 2^n inside $\mathrm{Aut}(F_n)$. We have $2^{n-1} = 2^n \cap \mathrm{SAut}(F_n)$. When talking about 2^n or 2^{n-1} inside $\mathrm{Aut}(F_n)$, we will always mean these subgroups.

The elements σ_{ij} with $i, j \in N$ generate a symmetric group S_n. Each of the sets

$$\{\epsilon_i \mid i \in N\}, \{\rho_{ij} \mid i, j \in N\} \text{ and } \{\lambda_{ij} \mid i, j \in N\}$$

is preserved under conjugation by elements of S_n, and the left conjugation coincides with the natural action by S_n on the indices. We have

$$A_n = S_n \cap \mathrm{SAut}(F_n)$$

The elements σ_{ij} with $i, j \in N \cup \{n+1\}$ generate a symmetric group S_{n+1}. Again, we have $A_{n+1} = S_{n+1} \cap \mathrm{SAut}(F_n)$. Again, when we talk about A_n, S_n, A_{n+1} or S_{n+1} inside $\mathrm{Aut}(F_n)$, we mean these subgroups.

Since the symmetric group S_n acts on 2^n by permuting the indices of the elements ϵ_i, we have $2^n \rtimes S_n < \mathrm{Aut}(F_n)$ (note that this is the Coxeter group of type $\mathtt{B}_n$). As usual, we will refer to this specific subgroup as $2^n \rtimes S_n$. Clearly, $\mathrm{SAut}(F_n) \cap 2^n \rtimes S_n$ contains $2^{n-1} \rtimes A_n$. We will denote this subgroup as D_n' (since it is isomorphic to the derived subgroup of the Coxeter group of type $\mathtt{D}_n$ when $n \geqslant 5$; note that we have no such isomorphism for $n = 4$). Note that 2^{n-1} inside D_n' is generated by the elements $\epsilon_i \epsilon_j$ with $i \neq j$.

Lemma 2.1. *Let $n \geqslant 3$. Then the normal closure of A_n in D_n' is the whole of D_n'.*

Proof.

$$\epsilon_1 \epsilon_2 = [\epsilon_1 \epsilon_3, \sigma_{13}\sigma_{12}] \in \langle\!\langle A_n \rangle\!\rangle$$

and so every $\epsilon_i \epsilon_j$ lies in the normal closure of A_n as well, since A_n acts transitively on unordered pairs in N. □

The Nielsen Realisation theorem for free groups (proved independently by Culler [**Cul**], Khramtsov [**Khr1**] and Zimmermann [**Zim1**]) states that any finite subgroup of $\mathrm{Aut}(F_n)$ can be seen as a group of basepoint preserving automorphisms of a graph with fundamental group identified with F_n. From this point of view, the subgroup $2^n \rtimes S_n$ is the automorphism group of the n-rose (the bouquet of n circles), and the subgroup S_{n+1} is the basepoint preserving automorphism group of the $n+1$-cage graph (a graph with two vertices and $n+1$ edges connecting them).

Remark 2.2. Throughout, we are going to make extensive use of the Steinberg commutator relations in $\mathrm{Aut}(F_n)$, that is

$$\rho_{ij}{}^{-1} = [\rho_{ik}{}^{-1}, \rho_{kj}{}^{-1}]$$

and

$$\lambda_{ij}{}^{-1} = [\lambda_{ik}{}^{-1}, \lambda_{kj}{}^{-1}]$$

We will also use

$$\rho_{ij}{}^{\delta} = \lambda_{ij}$$

Another two types of relations which we will frequently encounter are already present in the proof of the following lemma (based on observations of Bridson–Vogtmann [**BV1**]).

Lemma 2.3. *For $n \geqslant 3$, all automorphisms $\rho_{ij}{}^{\pm 1}$ and $\lambda_{ij}{}^{\pm 1}$ (with $i \neq j$) are conjugate inside* $\mathrm{SAut}(F_n)$.

Proof. Observe that

$$\rho_{ij}{}^{\epsilon_i \epsilon_j} = \lambda_{ij}$$

and that

$$\rho_{ij}{}^{\epsilon_j \epsilon_k} = \rho_{ij}{}^{-1}$$

where $k \notin \{i, j\}$.

When $n \geqslant 4$, the subgroup A_n acts transitively on ordered pairs in N, and so we are done.

Let us suppose that $n = 3$. In this case, using the 3-cycle $\sigma_{12}\sigma_{23}$, we immediately see that ρ_{12}, ρ_{23} and ρ_{31} are all conjugate. We also have

$$\rho_{12}{}^{\sigma_{12}\epsilon_3} = \rho_{21}$$

and thus all right transvections are conjugate, and we are done. □

We also note the following useful fact, following from Gersten's presentation of $\mathrm{SAut}(F_n)$ [**Ger**].

Proposition 2.4. *For every $n \geqslant 3$, the group* $\mathrm{SAut}(F_n)$ *is perfect.*

2.3. Linear quotients. Observe that abelianising the free group F_n gives us a map

$$\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$$

This homomorphism is in fact surjective, since each elementary matrix in $\mathrm{SL}_n(\mathbb{Z})$ has a transvection in its preimage. We will refer to this map as the natural homomorphism $\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$.

The finite quotients of $\mathrm{SL}_n(\mathbb{Z})$ (for $n \geqslant 3$) are controlled by the congruence subgroup property as proven by Mennicke [**Men**]). In particular, noting that $\mathrm{SAut}(F_n)$ is perfect (Proposition 2.4), we conclude that the non-trivial simple quotients of $\mathrm{SL}_n(\mathbb{Z})$ are the groups $\mathrm{L}_n(p)$ where p ranges over all primes. The smallest one is clearly $\mathrm{L}_n(2)$.

We will refer to the compositions of the natural map $\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$ and the quotient maps $\mathrm{SL}_n(\mathbb{Z}) \to \mathrm{L}_n(p)$ as natural maps as well.

We will find the following observations (due to Bridson–Vogtmann [**BV1**]) most useful.

Lemma 2.5. *Let $n \geqslant 3$, and let φ be a homomorphism with domain* $\mathrm{SAut}(F_n)$.

(1) *If n is even, and $\varphi(\delta)$ is central in $\operatorname{im}\varphi$, then φ factors through the natural map* $\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$.
(2) *For any n, if there exists $\xi \in 2^{n-1} \smallsetminus \{1, \delta\}$ such that $\varphi(\xi)$ is central in $\operatorname{im}\varphi$, then φ factors through the natural map* $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$.

(3) For any n, if there exists $\xi \in D'_n \smallsetminus 2^{n-1}$ such that $\varphi(\xi)$ is central in $\operatorname{im}\varphi$, then φ is trivial.

Proof. (1) We have

$$\delta \rho_{ij} \delta = \lambda_{ij}$$

for every i, j. Thus, φ factors through the group obtained by augmenting Gersten's presentation [**Ger**] of $\mathrm{SAut}(F_n)$ by the additional relations $\rho_{ij} = \lambda_{ij}$. But this is equivalent to Steinberg's presentation of $\mathrm{SL}_n(\mathbb{Z})$.

(2) We claim that there exists $\tau \in A_n$ such that

$$[\xi, \tau] = \epsilon_i \epsilon_j$$

We have $\xi = \prod_{i \in I} \epsilon_i$ for some $I \subset N$. Take $i \in I$ and $j \notin I$. Suppose first that there exist distinct α, β either in $I \smallsetminus \{i\}$ or in $N \smallsetminus (I \cup \{j\})$. Then $\tau = \sigma_{ij}\sigma_{\alpha\beta}$ is as claimed.

If no such α and β exist, then $n \leqslant 4$ and $\xi = \epsilon_i \epsilon_{i'}$ for some $i' \neq i$. Thus we may take $\tau = \sigma_{ii'}\sigma_{ji}$. This proves the claim

Now $\varphi([\xi, \tau]) = 1$ since $\varphi(\xi)$ is central. Using the action of A_n on 2^{n-1} we immediately conclude that $2^{n-1} \leqslant \ker \varphi$. Thus we have

$$\varphi(\rho_{ij}) = \varphi(\rho_{ij})^{\varphi(\epsilon_i \epsilon_j)} = \varphi({\rho_{ij}}^{\epsilon_i \epsilon_j}) = \varphi(\lambda_{ij})$$

Now Gersten's presentation tells us that φ factors through the natural map

$$\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$$

Moreover,

$$\varphi(\rho_{ij}) = \varphi(\rho_{ij})^{\varphi(\epsilon_j \epsilon_k)} = \varphi({\rho_{ij}}^{\epsilon_j \epsilon_k}) = \varphi({\rho_{ij}}^{-1})$$

where $k \notin \{i, j\}$. The result of Mennicke [**Men**] tells us that in this case φ factors further through $\mathrm{SL}_n(Z) \to \mathrm{L}_n(2)$.

(3) We write $\xi = \xi'\tau$, where $\xi' \in 2^{n-1}$ and $\tau \in A_n$.

Suppose first that τ is not a product of commuting transpositions. Then, without loss of generality, we have

$${\rho_{12}}^{\xi} = {x_{23}}^{\pm 1}$$

where $x \in \{\rho, \lambda\}$. Now

$$\varphi({\rho_{13}}^{-1}) = \varphi([{\rho_{12}}^{-1}, {\rho_{23}}^{-1}]) = [\varphi(x_{23})^{\mp 1}, \varphi(\rho_{23})^{-1}] = 1$$

as ${x_{23}}^{\pm 1}$ commutes with ρ_{23}. This trivialises φ, since $\mathrm{SAut}(F_n)$ is generated by transvections, and every two transvections are conjugate.

Now suppose that τ is a product of commuting transpositions. Then $n \geqslant 4$, and without loss of generality

$${\rho_{12}}^{\xi} = {x_{34}}^{\pm 1}$$

where $x \in \{\rho, \lambda\}$. Now

$$\varphi({\rho_{14}}^{-1}) = \varphi([{\rho_{12}}^{-1}, {\rho_{24}}^{-1}]) = [\varphi(x_{34})^{\mp 1}, \varphi(\rho_{24})^{-1}] = 1$$

as ${x_{34}}^{\pm 1}$ commutes with ρ_{24}. This trivialises φ as before. □

Remark 2.6. In (2), we can draw the same conclusion if we assume that

$$\varphi(\rho_{ij}) = \varphi(\lambda_{ij}) = \varphi(\rho_{ij})^{-1}$$

Corollary 2.7. *Let $\varphi\colon \mathrm{SAut}(F_n) \to K$ be a homomorphism. If K is finite and $\varphi|_{D'_n}$ is not injective, then*

(1) φ is trivial; or

(2) $|K| > |\mathrm{L}_n(2)|$; or

(3) $K \cong \mathrm{L}_n(2)$ and φ is the natural map up to postcomposition with an automorphism of $\mathrm{L}_n(2)$.

Many parts of the current paper are inductive in nature, and they are all based on the following observation.

Lemma 2.8. *For any $k \leqslant n+1$, the group* $\mathrm{SAut}(F_n)$ *contains an element* ξ *of order* k *whose centraliser contains* $\mathrm{SAut}(F_{n-k})$*. When* k *is odd then the centraliser contains* $\mathrm{SAut}(F_{n-k+1})$*. Moreover, when* $k \geqslant 5$ *then the normal closure of* ξ *is the whole of* $\mathrm{SAut}(F_n)$*.*

Proof. Let Γ be a graph with vertex set $\{u, v\}$ and edge set consisting of k distinct edges connecting u to v and $n-k+1$ distinct edges running from v to itself (see Figure 2.1). Let K denote the set of the former k edges. We identify the fundamental group of Γ with F_n. Elements of $\mathrm{SAut}(F_n)$ then correspond to based homotopy equivalences of Γ.

Let ξ denote the automorphism of Γ which cyclically permutes the edges of K. When k is odd the action of ξ on the remaining edges is trivial; when k is even, then ξ acts on $n-k$ of the remaining edges trivially, but it flips the last edge. It is clear that ξ defines an element of $\mathrm{SAut}(F_n)$ of order k. Also, ξ fixes pointwise a free factor of F_n of rank $n-k$ when k is even or $n-k+1$ when k is odd.

Observe that we have a copy of the alternating group A_k permuting the edges in K. Suppose that $k \geqslant 5$, and let τ denote some 3-cycle in A_k. It is obvious that $[\xi, \tau]$ is a non-trivial element of A_k; it is also clear that A_k is simple. These two facts imply that A_k lies in the normal closure of ξ. But this A_k contains an A_{k-1} contained in the standard $A_n < \mathrm{SAut}(F_n)$, and the result follows by Lemma 2.5(3). $\square$

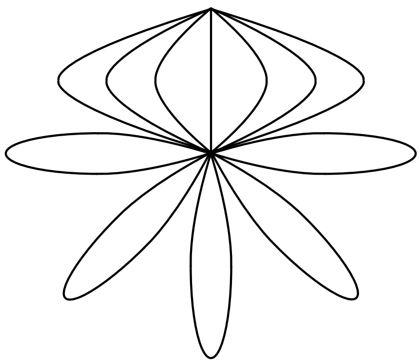

FIGURE 2.1. The graph Γ with $(n, k) = (11, 7)$

One can easily give an algebraic description of the element ξ above; in fact we will do this for $k = 3$ when we deal with classical groups in characteristic 3 in Section 7.1.

3. Alternating groups

In this section we will give lower bounds on the cardinality of a set on which the groups $\mathrm{SAut}(F_n)$ can act non-trivially. The cases $n \in \{3, \dots, 8\}$ are done in a somewhat ad-hoc manner, and we begin with these, developing the necessary tools along the way. We will conclude the section with a general result for $n \geqslant 9$.

As a corollary, we obtain that alternating groups are never the smallest quotients of $\mathrm{SAut}(F_n)$ (for $n \geqslant 3$), with a curious exception for $n = 4$, since in this case the (a fortiori smallest) quotient $\mathrm{L}_4(2)$ is isomorphic to the alternating group A_8.

Lemma 3.1 ($n = 3$)**.** *Any action of* $\mathrm{SAut}(F_3)$ *on a set* X *with fewer than* 7 *elements is trivial.*

This result can be easily verified using GAP. For this reason, we offer only a sketch proof.

SKETCH OF PROOF. The action gives us a homomorphism $\varphi\colon \mathrm{SAut}(F_3) \to S_6$. Since $\mathrm{SAut}(F_3)$ is perfect, the image lies in A_6.

Consider the set of transvections

$$T = \{{\rho_{ij}}^{\pm 1}, {\lambda_{ij}}^{\pm 1}\}$$

By Lemma 2.3 all elements in T are conjugate in $\mathrm{SAut}(F_3)$, and hence also in the image of φ.

Consider the equivalence relation on T where elements $x, y \in T$ are equivalent if $\varphi(x) = \varphi(y)$. It is clear that the equivalence classes are equal in cardinality. We first show that if any (hence each) of these equivalence classes has cardinality at least 3, then φ is trivial. Let Φ be such a class. Without loss of generality we assume that $\rho_{12} \in \Phi$.

Suppose that

$$|\Phi \cap \{{\rho_{12}}^{\pm 1}, {\lambda_{12}}^{\pm 1}\}| \geqslant 3$$

Independently of which 3 of the 4 elements lie in Φ, their image under φ is centralised by $\varphi(\epsilon_1\epsilon_2)$, and so in fact all four of these elements lie in Φ.

Now, we have

$$\varphi(\rho_{12}) = \varphi({\rho_{12}}^{-1}) = \varphi(\lambda_{12}) = \varphi({\lambda_{12}}^{-1})$$

In this case we conclude from Remark 2.6 that φ factors through $\mathrm{L}_3(2)$. But $\mathrm{L}_3(2)$ is a simple group containing an element of order 7, and so φ is trivial.

If $|\Phi \cap \{{\rho_{12}}^{\pm 1}, {\lambda_{12}}^{\pm 1}\}| < 3$ then there exists an element ${x_{ij}}^{\pm 1} \in \Phi$ with x being either ρ or λ, and with $(i, j) \neq (1, 2)$. If $(i, j) = (1, 3)$ then

$$\varphi({\rho_{13}}^{-1}) = \varphi([{\rho_{12}}^{-1}, {\rho_{23}}^{-1}]) = [\varphi(x_{13})^{\mp 1}, \varphi(\rho_{23})^{-1}] = 1$$

which implies that φ is trivial. We proceed in a similar fashion for all other values of (i, j). This way we verify our claim that if φ is non-trivial then $|\Phi| \leqslant 2$.

There are exactly 10 elements in T which commute with ρ_{12}, namely

$$C_T(\rho_{12}) = \{{\rho_{12}}^{\pm 1}, {\lambda_{12}}^{\pm 1}, {\lambda_{13}}^{\pm 1}, {\rho_{32}}^{\pm 1}, {\lambda_{32}}^{\pm 1}\}$$

Using what we have learned above about the cardinality of Φ, we see that φ is not trivial only if the conjugacy class of $\varphi(\rho_{12})$ in A_6 contains at least 5 elements commuting with $\varphi(\rho_{12})$. We also know that this conjugacy class has to contain $\varphi(\rho_{12})^{-1}$. By inspection we see that the only such conjugacy class in A_6 is that of $\tau = (12)(34)$.

The conjugacy class of τ has exactly 5 elements, say $\{\tau, \tau_1, \tau_2, \tau_1', \tau_2'\}$, and the elements τ_i and τ_j' do not commute for any $i, j \in \{1, 2\}$. Hence the maximal subset of the conjugacy class of τ in which all elements pairwise commute is of cardinality 3. But in $C_T(\rho_{12})$ we have 8 such elements, namely

$$\{{\rho_{12}}^{\pm 1}, {\lambda_{12}}^{\pm 1}, {\rho_{32}}^{\pm 1}, {\lambda_{32}}^{\pm 1}\}$$

This implies that $|\Phi| > 2$, which forces φ to be trivial. □

Note that $\mathrm{L}_3(2)$ acts non-trivially on the set of non-zero vectors in 2^3 (thought of as a vector space) which has cardinality 7, and so $\mathrm{SAut}(F_3)$ has a non-trivial action on a set of 7 elements. Thus the result above is sharp.

Lemma 3.2. *Let* $n \geqslant 4$. *Suppose that* $\mathrm{SAut}(F_n)$ *acts non-trivially on a set* X *so that* $\epsilon_1\epsilon_2$ *acts trivially. Then* $|X| \geqslant 2^{n-1}$.

Proof. When $\epsilon_1\epsilon_2$ acts trivially then the action factors through the natural map $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ by Lemma 2.5. Now, $\mathrm{L}_n(2)$ is simple and cannot act on a set smaller than 2^{n-1} non-trivially, by [**KL**, Theorem 5.2.2]. □

Lemma 3.3 ($n = 4$)**.** *Any action of* $\mathrm{SAut}(F_4)$ *on a set* X *with fewer than* 8 *elements is trivial.*

Proof. The group D_4' cannot act faithfully on fewer than 8 points, and the group $D_4'/\langle\delta\rangle$ cannot act faithfully on fewer than 12 points (both of these facts can be easily checked). Thus we are done by Lemma 2.5, since $\mathrm{L}_4(2) \cong A_8$ cannot act on fewer than 8 points. □

Note that in this case the result is also sharp, since we have the epimorphism

$$\mathrm{SAut}(F_4) \to \mathrm{L}_4(2) \cong A_8$$

Lemma 3.4. *Let* $n \geqslant 1$. *Let* D_n' *act on a set of cardinality less than* $2^{n-1-p(n)}$ *where* $p(n) = 0$ *for* n *odd, and* $p(n) = 1$ *for* n *even. Then* 2^{n-1} *acts trivially on every point fixed by* A_n.

Proof. Let x be a point fixed by A_n, and consider the 2^{n-1}-orbit thereof. Such an orbit corresponds to a subgroup of 2^{n-1} normalised by A_n. It cannot correspond to the trivial subgroup, as then the orbit would be too large. For the same reason it cannot correspond to the subgroup generated by δ (which is a subgroup when n is even). It is easy to see that the only remaining subgroup is the whole of 2^{n-1}, and so the action on x is trivial. □

Lemma 3.5. *Let* $n \geqslant 4$. *Suppose that* $\mathrm{SAut}(F_n)$ *acts on a set* X *of cardinality less than* $2^{n-1-p(n)}$, *where* $p(n)$ *is as above, in such a way that* A_n *has at most one non-trivial orbit. Then* $\mathrm{SAut}(F_n)$ *acts trivially on* X.

Proof. By Lemma 3.4, every point fixed by A_n is also fixed by 2^{n-1}. This implies that every A_n-orbit is already a D_n'-orbit, and so every point in X is fixed by 2^{n-1}, and thus the action of $\mathrm{SAut}(F_n)$ on X is trivial by Lemma 3.2. □

For higher values of n we will use the fact that actions of alternating groups on small sets are well-understood. Recall that $A_n = \mathrm{Alt}(N)$ denotes the group of even permutations of the set $N = \{1, \ldots, n\}$.

Definition 3.6. We say that a transitive action of $A_n = \mathrm{Alt}(N)$ on a set X is *associated to* k (with $k \in \mathbb{N}$) if for every (or equivalently any) $x \in X$ the stabiliser of x in A_n contains $\mathrm{Alt}(J)$ with $J \subseteq N$ and $|J| = k$, and does not contain $\mathrm{Alt}(J')$ for any $J' \subseteq N$ of larger cardinality. Notice that $k \geqslant 2$.

In the following lemma we write A_{n-1} for the subgroup $\mathrm{Alt}(N \smallsetminus \{n\})$ of $A_n = \mathrm{Alt}(N)$.

Lemma 3.7. *Let* $n \geqslant 6$. *Suppose that we are given a transitive action* π *of* A_n *on a set* X *which is associated to* k. *Then*

(1) *the action of* A_{n-1} *on each of its orbits is associated to* k *or* $k-1$;
(2) *if* $k > \frac{n}{2}$ *then there is exactly one orbit of* A_{n-1} *of the latter kind;*
(3) *if* $k > \frac{n}{2}$ *then any other transitive action* π' *of* A_n *on a set* X' *isomorphic to* $\pi|_{A_{n-1}}$ *when restricted to* A_{n-1} *is isomorphic to* π.

Proof. Before we start, let us make an observation: let I and J be two subsets of N of cardinality at least 3 each, and such that $I \cap J \neq \emptyset$. Then

$$\langle \mathrm{Alt}(I), \mathrm{Alt}(J) \rangle = \mathrm{Alt}(I \cup J)$$

There are at least two quick ways of seeing it: the subgroup $\langle \mathrm{Alt}(I), \mathrm{Alt}(J)\rangle$ clearly contains all 3-cycles; the subgroup $\langle \mathrm{Alt}(I), \mathrm{Alt}(J)\rangle$ acts 2-transitively, and so primitively, on $I \cup J$, and contains a 3-cycle, which allows us to use Jordan's theorem.

(1) Let x be an element in an A_{n-1}-orbit O. If the stabiliser S of x in A_n contains $\mathrm{Alt}(J)$ with $|J| = k$ and $n \notin J$, then $\mathrm{Alt}(J) \subseteq A_{n-1}$ and so the action of A_{n-1} on O is associated to at least k. It is clear that the action cannot be associated to any integer greater than k.

If n is in J for every $J \subseteq N$ of size k with $\mathrm{Alt}(J) \subseteq S$, then $\mathrm{Alt}(J \smallsetminus \{n\})$ is contained in $S \cap A_{n-1}$ and the action of A_{n-1} on O is associated with at least $k-1$. It is clear that this action cannot be associated to any integer greater than $k-1$.

(2) Clearly, there exists $x \in X$ such that its stabiliser S in A_n contains $\mathrm{Alt}(J)$ with $|J| = k$ and $n \in J$. Note that J is unique – if there were another subset $I \subseteq N$ with $|I| = k$ and $\mathrm{Alt}(I) \leqslant S$, then $I \cap J$ would need to intersect non-trivially (as $k > \frac{n}{2}$), and so we would have $\mathrm{Alt}(I \cup J) \leqslant S$. Hence we may conclude from the proof of (1) that the action of A_{n-1} on O, its orbit of x, is associated to $k-1$.

Now suppose that there exists a point $x' \in X$ with A_{n-1}-orbit O', stabiliser S' in A_n, and subset $J' \subseteq N$ of cardinality k with $n \in J'$ and $\mathrm{Alt}(J') \leqslant S'$. There exists $\tau \in A_n$ such that

$$x = \tau.x'$$

and so $S' = S^{\tau}$. Thus $\tau(J') = J$, and therefore there exists $\sigma \in \mathrm{Alt}(J)$ such that $\sigma\tau(n) = n$. But then also $\sigma^{-1}.x = x$ and so $x = \sigma\tau.x'$. But $\sigma\tau \in A_{n-1}$, and therefore $O = O'$.

(3) Let us start by looking at π'. By assumption, this action is associated to at least $k-1$, since it is when restricted to A_{n-1}. It also cannot be associated to any integer larger than k, since then (2) would forbid the existence of an A_{n-1}-orbit in X' associated to $k-1$, and we know that such an orbit exists.

If $k - 1 > \frac{n}{2}$ then (2) implies that $\pi'|_{A_{n-1}}$ has an orbit that is associated to $k-1$, but clearly none associated to $k-2$. Thus, by (2), π' is associated to k.

If $k - 1 \leqslant \frac{n}{2}$ then in particular $k \neq n$, and so there is an A_{n-1}-orbit in X associated to k, and therefore π' cannot be associated to $k-1$. We conclude that π' is associated to k.

Pick an $x \in X$ so that the A_{n-1}-action on the A_{n-1} orbit O of x is associated to $k-1$. Let $\theta\colon X \to X'$ be a A_{n-1}-equivariant bijection (which exists by assumption). Let $x' = \theta(x)$. Let S denote the stabiliser of x in A_n, and S' the stabiliser of x' in A_n.

We have $\mathrm{Alt}(J) \times G = H \leqslant S$, where $|J| = k$, the index $|S : H|$ is at most 2, and $G \leqslant \mathrm{Alt}(N \smallsetminus J)$ – we need to observe that J is unique, as proven above. Since the A_{n-1}-orbit of x is associated to $k-1$, we see that $n \in J$.

Now the stabiliser of x' in A_{n-1} is equal to $S \cap A_{n-1}$, and S' contains $S \cap A_{n-1}$ and some $\mathrm{Alt}(J')$ with $|J'| = k$. Since $k > \frac{n}{2}$, the subsets J and J' intersect, and so $\mathrm{Alt}(J \cup J') \leqslant S'$ as before. But the A_{n-1}-action on its orbit of x' is associated to $k-1$, and so we must have $J = J'$. This implies that $S' = S$, since the index of H in S is equal to the index of $H \cap A_{n-1}$ in $S \cap A_{n-1}$. □

Theorem 3.8 (Dixon–Mortimer [**DM**, Theorem 5.2A])**.** *Let $n \geqslant 5$, and let $r \leqslant n/2$ be an integer. Let $H < A_n = \mathrm{Alt}(N)$ be a proper subgroup of index less than $\binom{n}{r}$. Then one of the following holds:*

(1) The subgroup H contains a subgroup A_{n-r+1} of A_n fixing $r-1$ points in N.

(2) We have $n = 2m$ and $|A_n : H| = \frac{1}{2}\binom{n}{m}$. Moreover, H contains the product $A_m \times A_m$.

(3) *The pair* $(n, |A_n : H|)$ *is one of the six exceptional cases:*

$$(5,6),(6,6),(6,15),(7,15),(8,15),(9,120)$$

Note that the original theorem contains more information in each of the cases; for our purposes however, the above version will suffice.

We can rephrase (1) by saying that the action $A_n \curvearrowright A_n/H$ is associated to at least $n-r+1$.

Corollary 3.9 ($n=5$)**.** *Every action of* $\mathrm{SAut}(F_5)$ *on a set* X *of cardinality less than* 12 *is trivial.*

Proof. By Theorem 3.8, the only orbits of A_6 in X are of cardinality 1, 6 or 10. There can be at most one orbit of size greater than 1, and it contains at most two non-trivial orbits of A_5, since such orbits have cardinality at least 5. If there is at most one non-trivial A_5-orbit, we invoke Lemma 3.5. Otherwise, let x be a point on which A_6 acts non-trivially; its A_6-orbit consists of 10 points. We know from case (2) of Theorem 3.8 that it is fixed by two commuting 3-cycles. The A_5-orbit of x has cardinality 5, and so it is the natural A_5-orbit (by Theorem 3.8 again). Thus x is fixed by some standard A_4. But now any standard A_4 together with any two commuting 3-cycles generates A_6, and so x is fixed by A_6, which is a contradiction. □

Corollary 3.10 ($n=6$)**.** *Every action of* $\mathrm{SAut}(F_6)$ *on a set* X *of cardinality less than* 14 *is trivial.*

Proof. By Theorem 3.8, the only orbits of A_7 in X are either trivial or the natural orbits of size 7. There can be at most one such natural orbit, and so A_6 has at most one non-trivial orbit. We now invoke Lemma 3.5. □

We now begin the preparations towards the main tool in this section.

Lemma 3.11. *Let* $n \geqslant 2$, *and let* $r \leqslant n/2$ *be a positive integer. Let* D'_n *act on a set* X *of cardinality less than* $\binom{n}{r}$, *and let* x *be a point stabilised by*

$$\langle \{\epsilon_i\epsilon_j \mid i,j \in I\}\rangle$$

where I *is a subset of* N. *Then* I *can be taken to have cardinality at least* $n-r+1$, *provided that*

(1) I *contains more than half of the points of* N*; or*
(2) I *contains exactly half of the points, and* x *is not fixed by some* $\epsilon_i\epsilon_j$ *with* $i,j \notin I$.

Proof. Let S denote the stabiliser of x in 2^{n-1}. Consider a maximal (with respect to inclusion) subset J of N such that for all $i,j \in J$ we have $\epsilon_i\epsilon_j \in S$. We call such a subset a *block*. It is immediate that blocks are pairwise disjoint, and one of them, say J_0, contains I.

If the block J_0 contains more than half of the points in N (which is guaranteed to happen in the case of assumption (1)), then it is the unique largest block of S. In the case of assumption (2), the block J_0 may contain exactly half of the points, but all of the other blocks contain strictly fewer elements. Thus, again, J_0 is the unique largest block.

Since J_0 is unique, it is clear that any element $\tau \in A_n$ which does not preserve J_0 gives $S^\tau \neq S$, and so in particular $\tau.x \neq x$. Using this argument we see that X has to contain at least $\binom{n}{n-|J_0|}$ elements. But $|X| < \binom{n}{r}$, and so $|J_0| > n-r$, and we are done. □

Definition 3.12. Let $\mathrm{SAut}(F_n)$ act on a set X. For each point $x \in X$ we define

(1) I_x to be a subset of N such that

$$\langle \{\epsilon_i \epsilon_j \mid i,j \in I_x\} \rangle$$

fixes x, and I_x has maximal cardinality among such subsets.

(2) J_x to be a subset of N such that x is fixed by

$$\mathrm{Alt}(J_x) \leqslant \mathrm{Alt}(N) = A_n$$

and J_x is of maximal cardinality among such subsets.

The following is the main technical tool of this part of the paper.

Lemma 3.13. *Let $n \geqslant 5$. Suppose that $\mathrm{SAut}(F_n)$ acts transitively on a set X in such a way that*

(1) there exists a point $x_0 \in X$ with I_{x_0} containing more than half of the points in N; and

(2) for every $x \in X$ we have $|J_x| \geqslant \frac{n+3}{2}$.

Then every point x is fixed by $\mathrm{SAut}(F(J_x))$, provided that $|X| < \min\{2^{n-r}, \binom{n}{r}\}$ for some positive integer $r < \frac{n}{2} - 1$.

Proof. Lemma 3.11 tells us immediately that I_{x_0} contains at least $\nu = n - r + 1$ points. We claim that in fact every I_x contains at least ν points.

Since the action of $\mathrm{SAut}(F_n)$ on X is transitive, and $\mathrm{SAut}(F_n)$ is generated by transvections, it is enough to prove that for every point x with I_x of size at least ν, and every transvection, the image y of x under the transvection has $|I_y| \geqslant \nu$. For concreteness, let us assume that the transvection in question is ρ_{ij} (the situation is analogous for the left transvections). Since ρ_{ij} commutes with every involution $\epsilon_\alpha \epsilon_\beta$ with $\alpha, \beta \in I_x \smallsetminus \{i,j\}$, we see that y is fixed by $\epsilon_\alpha \epsilon_\beta$. But

$$|I_x| - 2 \geqslant \nu - 2 = n - r - 1 > \frac{n}{2}$$

and so $I_y \supseteq I_x \smallsetminus \{i,j\}$ (here we use the fact that I_y is defined to be the largest block). Therefore $|I_y| \geqslant \nu$ by Lemma 3.11. We have thus established that I_x contains at least ν points for every $x \in X$.

Note that the sets I_x form a poset under inclusion. Pick an element $z \in X$ so that I_z is minimal in this poset. Let Z denote the subset of X consisting of points w with $I_w = I_z$. Now for every $i,j \in I_z$ and every $w \in Z$ we have

$$I_z \smallsetminus \{i,j\} \subseteq I_{\rho_{ij}.w}$$

since ρ_{ij} commutes with involutions $\epsilon_\alpha \epsilon_\beta$ with $\alpha, \beta \in I_z \smallsetminus \{i,j\}$ as before.

Assume there exists $k \in I_{\rho_{ij}.w} \smallsetminus I_z$. By definition of $I_{\rho_{ij}.w}$ and using the above inclusion, there exists $l \in I_z \smallsetminus \{i,j\}$ such that $\epsilon_l \epsilon_k$ fixes $\rho_{ij}.w$. But $\epsilon_l \epsilon_k$ commutes with ρ_{ij}, and hence fixes w, which forces $k \in I_z$, a contradiction. Thus

$$I_{\rho_{ij}.w} \subseteq I_z$$

Since I_z is minimal, we conclude that $\rho_{ij}.w \in Z$. An analogous argument applies to left transvections, and so Z is preserved by

$$\mathrm{SAut}(F(I_z)) = \langle \{\rho_{ij}, \lambda_{ij} \mid i,j \in I_z\} \rangle \leqslant \mathrm{SAut}(F_n)$$

But in the action of $\mathrm{SAut}(F(I_z))$ on Z the involutions $\epsilon_\alpha \epsilon_\beta$ with $\alpha, \beta \in I_z$ act trivially. Therefore this action is trivial by Lemma 3.2, since X has fewer than 2^{n-r} points and $n - r + 1 > \frac{n}{2} + 2 \geqslant 3\frac{1}{2}$. In particular, we have $I_w \subseteq J_w$ for every $w \in Z$.

Every $w \in Z$ is fixed by $\mathrm{SAut}(F(I_w))$, but also by $\mathrm{Alt}(J_w)$ by assumption. Thus, it is fixed by the subgroup of $\mathrm{SAut}(F_n)$ generated by the two subgroups. It is clear that this is $\mathrm{SAut}(F(J_w \cup I_w))$ and so we have finished the proof for points in

Z. Now we also see that in fact $I_w = J_w$, since the subgroup of 2^{n-1} corresponding to J_w lies in $\mathrm{SAut}(F(J_w))$ and hence fixes w.

Let $x \in X$ be any point. Since the action of $\mathrm{SAut}(F_n)$ is transitive, there exists a finite sequence $z = x_0, x_1, \ldots, x_{m-1}, x_m = x$ such that for every i there exists a transvection τ_i with $\tau_i.x_i = x_{i+1}$ (we assume as well that the elements of the sequence are pairwise disjoint). We claim that every x_i is fixed by $\mathrm{SAut}(F(J_{x_i}))$. Let i be the smallest index so that our claim is not true for x_i. As usual, for concreteness, let us assume that $\tau_{i-1} = \rho_{\alpha\beta}$. Note that we cannot have both α and β in $J_{x_{i-1}}$, since then the action of $\rho_{\alpha\beta}$ on x_{i-1} would be trivial.

Consider the intersection $(J_{x_{i-1}} \cap J_{x_i}) \smallsetminus \{\alpha, \beta\}$. By assumption, the intersection $J_{x_{i-1}} \cap J_{x_i}$ contains at least 3 points, and at most one of these points lies in $\{\alpha, \beta\}$. Thus there exist $\alpha', \beta' \in J_{x_{i-1}} \cap J_{x_i}$ such that $\rho_{\alpha'\beta'}$ commutes with $\rho_{\alpha\beta}$. The action of $\rho_{\alpha'\beta'}$ on x_{i-1} is trivial, and thus it must also be trivial on $x_i = \rho_{\alpha\beta}.x_{i-1}$. We also know that $\mathrm{Alt}(J_{x_i})$ acts trivially on x_i, and so every right transvection with indices in J_{x_i} acts trivially on x_i. This implies that x_i is fixed by $\mathrm{SAut}(F(J_{x_i}))$, which contradicts the minimality of x_i, and so proves the claim, and therefore the result. □

Proposition 3.14 ($n \geqslant 7$)**.** *Let $n \geqslant 7$. Every action of* $\mathrm{SAut}(F_n)$ *on a set of cardinality less than*

$$\max_{r \leqslant \frac{n}{2}-3} \min \left\{ 2^{n-r-p(n)}, \binom{n}{r} \right\}$$

is trivial, where $p(n)$ equals 0 *when n is odd and* 1 *when n is even.*

Proof. Let X denote the set on which we are acting. Without loss of generality we will assume that $\mathrm{SAut}(F_n)$ acts on X transitively, and that X is non-empty.

Let R denote a value of r for which

$$\max_{r \leqslant \frac{n}{2}-3} \min \left\{ 2^{n-r-p(n)}, \binom{n}{r} \right\}$$

is attained. Note that $R > 1$ by Lemma 10.1 for $n \geqslant 8$; a direct computation shows that $R = 2$ for $n \in \{7, 8\}$.

Let us first look at the action of A_{n+1}. Since $|X| < \binom{n}{R} < \binom{n+1}{R}$, Theorem 3.8 tells us that each orbit of A_{n+1} is

(1) associated to at least $n - R + 1$; or
(2) as described in case (2) of the theorem − this is immediately ruled out, since X would have to be too large by Lemma 10.2 for $n \geqslant 12$, and by direct computation for $n \in \{8, 10\}$; or
(3) one of the two exceptional actions $(8, 15)$ or $(9, 120)$ as in case (3) of the theorem.

For now let us assume that we are in case (1). Thus $|J_x| \geqslant n - R$ for each $x \in X$, and so the action of $\mathrm{SAut}(F_n)$ on X satisfies assumption (2) of Lemma 3.13. Also, by Lemma 3.7, there is at least one point $y \in X$ with $|J_y| \geqslant n - R + 1$.

Let J_0 denote a largest (with respect to cardinality) subset of N such that $\mathrm{Alt}(J_0)$ has a fixed point in X. Let x_0 be such a fixed point. Note that J_0 has at least $n - R + 1$ elements, and $|X| < 2^{n-R-p(n)}$, which implies by Lemma 3.4 that I_{x_0} contains at least $n - R + 1$ elements. As $R < \frac{n}{2}$, we conclude that the action of $\mathrm{SAut}(F_n)$ on X satisfies assumption (1) of Lemma 3.13.

We are now in position to apply Lemma 3.13. We conclude that x_0 is fixed by $\mathrm{SAut}(F(J_{x_0}))$. Let us consider the graph Γ from Figure 2.1 with $k = |J_{x_0}| + 1$ and fundamental group isomorphic to F_n. We can choose such an isomorphism so that $\mathrm{Alt}(J_{x_0})$ acts on Γ by permuting (in a natural way) all but one of the edges which are not loops. But it is clear that we also have a supergroup G of $\mathrm{Alt}(J_{x_0})$, which is isomorphic to an alternating group of rank $|J_{x_0}|+1$, and acts by permuting all such edges. By construction, $G < \mathrm{SAut}(F(J_{x_0}))$ and so $G.x_0 = x_0$. But now consider the action of G on X – by Lemma 3.7, it has to agree with the action of $\mathrm{Alt}(J')$, where J' is a superset of J_{x_0} with a single new element. By assumption, $\mathrm{Alt}(J')$ does not fix any point in X. However G does, and this is a contradiction. This implies that there is no superset J', but then we must have $J_{x_0} = N$, and so $\mathrm{SAut}(F_n)$ fixes a point in X. But the action is transitive, and so X is a single point. This proves the result.

Now let us investigate the exceptional cases. The first one occurs when $n = 7$, and the A_8-orbit of x has cardinality 15. We have $R = 2$ in this case, and so X has fewer than $\binom{7}{2} = 21$ elements. In this case, there are at most 5 points in $X \smallsetminus A_8.x$, and so the action of A_8 on each of these is trivial. Thus A_7 also fixes these points.

Now consider the action of A_7 on $A_8.x$. Since A_7 cannot fix any point here, and the smallest orbit of A_7 of size other than 1 and 7 has to be of size 15 by Theorem 3.8, we conclude, noting that $15 = 2 \cdot 7 + 1$, that A_7 acts transitively on $A_8.x$. Therefore the action of A_7 on X has exactly 1 non-trivial orbit, and so we may apply Lemma 3.5 – note that X has fewer than 2^6 points.

The remaining case occurs for $n = 8$; we have $R = 2$ and so X has fewer than $\binom{8}{2} = 28$ elements. But then we cannot have an orbit of size 120, and thus this exceptional case does not occur. □

Remark 3.15. In particular, we can put $r = \lfloor \frac{n}{2} \rfloor - 3$ in the above result; we see that (asymptotically) $\binom{n}{r}$ grows much faster than 2^{n-r} (in fact it grows like $n^{-1/2}2^n$), and so we obtain an exponential bound on the size of a set on which we can act non-trivially. The smallest set with a non-trivial action of $\mathrm{SAut}(F_n)$ known is also exponential in size – coming from the action of $\mathrm{L}_n(2)$ on the cosets of its largest maximal subgroup (see [**KL**, Table 5.2.A]). Hence the result above answers the question about the asymptotic size of such a set.

Theorem 3.16. *Let $n \geqslant 3$. Any action of $\mathrm{SAut}(F_n)$ on a set with fewer than $k(n)$ elements is trivial, where*

$$k(n) = \begin{cases} 7 & & n = 3 \\ 8 & & n = 4 \\ 12 & \textit{if} & n = 5 \\ 14 & & n = 6 \end{cases}$$

and $k(n) = \max_{r \leqslant \frac{n}{2} - 3} \min\{2^{n-r-p(n)}, \binom{n}{r}\}$ for $n \geqslant 7$, where $p(n)$ equals 0 when n is odd and 1 when n is even.

Proof. This follows from Lemmata 3.1 and 3.3, Corollaries 3.9 and 3.10, and Proposition 3.14. □

As commented after the proofs of Lemmata 3.1 and 3.3, the bounds are sharp when $n \in \{3, 4\}$.

Corollary 3.18. *Let $n \geqslant 3$ and K be a quotient of $\mathrm{SAut}(F_n)$ with $|K| \leqslant \mathrm{L}_n(2)$. If K is isomorphic to an alternating group, then $n = 4$ and $K \cong A_8 \cong L_4(2)$.*

Proof. The proof consists of two parts. Firstly, Lemma 10.3 tells us that

$$|A_{\binom{n}{2}}| > |\,\mathrm{L}_n(2)|$$

for $n \geqslant 7$. In view of the bounds in Theorem 3.16, this proves the result for $n \geqslant 7$ – for $n \in \{7, 8\}$ we have computed above that $r = 2$; for larger values of n we have $2^{n-3} > \binom{n}{2}$ by Lemma 10.1.

	Order
A_5	60
$\mathrm{L}_3(2)$	168
A_6	360
A_7	2520
$\mathrm{L}_4(2) \cong A_8$	21060
A_9	181440
A_{10}	1814400
$\mathrm{L}_5(2)$	9999360
A_{11}	19958400
A_{12}	239500800
A_{13}	3113510400
$\mathrm{L}_6(2)$	20158709760

TABLE 3.17. Small alternating groups

Secondly, for $3 \leqslant n \leqslant 6$, Table 3.17 lists all alternating groups of degree at least 5 smaller or equal (in cardinality) than $\mathrm{L}_6(2)$. The table also lists the groups $\mathrm{L}_n(2)$ in the relevant range. All these groups are listed in increasing order. The result follows from inspecting the table and comparing it to the bounds in Theorem 3.16. □

3.1. An application. We record here a further application of the bounds established in Theorem 3.16.

Theorem 3.19. *Let $n \geqslant 12$ be an even integer, and let $m \neq n$ satisfy $m < \binom{n+1}{2}$. Then every homomorphism*

$$\varphi\colon \mathrm{Out}(F_n) \to \mathrm{Out}(F_m)$$

has image of cardinality at most 2.

Proof. [**Kie1**, Theorems 6.8 and 6.10] tell us that φ has a finite image. Every finite subgroup of $\mathrm{Out}(F_m)$ can be realised by a faithful action on a finite connected graph Γ with $\delta(\Gamma) \geqslant 3$ and Euler characteristic $1-m$. These two facts immediately imply that Γ has fewer than $2m$ vertices. But now

$$2 \cdot \binom{n+1}{2} < \binom{n}{2} + 2 \cdot \binom{n+1}{2} + \binom{n+2}{2} = \binom{n}{4}$$

and

$$2 \cdot \binom{n+1}{2} < 2^{n-5}$$

for $n \geqslant 14$ by an argument analogous to Lemma 10.1. Therefore, for $n \geqslant 14$ we have

$$2 \cdot \binom{n+1}{2} < \min\left\{\binom{n}{r}, 2^{n-r-1}\right\}$$

with $r = 4$ (and such an r satisfies $r \leqslant \frac{n}{2} - 3$).

For $n = 12$ we take $r = 3$ and compute directly that the inequality also holds.

In any case, the action of $\mathrm{SOut}(F_n)$ (via φ) on the vertices of Γ is trivial. Now each vertex has at most $2m-2$ edges emanating from it, and so again the action of $\mathrm{SOut}(F_n)$ on these edges is trivial. Thus $\varphi(\mathrm{SOut}(F_n))$ is trivial, which proves the result. □

4. Sporadic groups

In this section we show that sporadic groups are never the smallest quotients of $\mathrm{SAut}(F_n)$. The proof relies on determining for each sporadic group its *D'-rank*, that is the largest n such that the group contains D'_n. This information can be extracted from the lists of maximal subgroups contained in [**CCN$^+$**] or in [**Wil**]; the lists are complete with the exception of the Monster group, in which case the list of possible maximal subgroups is known. The upper bound for the D'-rank of each sporadic group is recorded in Table 4.1 (which also lists the groups $\mathrm{L}_n(2)$ for comparison).

If K is a sporadic group of D'-rank smaller than n, then K is not the smallest quotient of $\mathrm{SAut}(F_n)$ by Lemma 2.5 (observing that K is not the smallest quotient of $\mathrm{SL}_n(\mathbb{Z})$). This observation allows us to rule out all but one sporadic group; the Deucalion is Fi_{22}, and we deal with it by other means.

Lemma 4.2. *Every homomorphism $\varphi\colon \mathrm{SAut}(F_7) \to \mathrm{Fi}_{22}$ is trivial.*

Proof. In the ATLAS [**CCN$^+$**] we see that there is a single conjugacy class of elements of order 5 in Fi_{22} denoted $5A$; moreover, the centraliser of an element $x \in 5A$ is of cardinality 600. We also see that Fi_{22} contains a copy of S_{10}, and we may without loss of generality assume that x is a 5-cycle in S_{10}. But then the centraliser of x inside S_{10} is $5 \times S_5$, which is already of order 600, and thus coincides with the centraliser of x in Fi_{22}.

Let τ be the element of order 5 given by Lemma 2.8; since its normal closure is $\mathrm{SAut}(F_7)$, its image in Fi_{22} is not trivial. Looking at the centraliser of τ, we obtain a homomorphism

$$\psi\colon \mathrm{SAut}(F_3) \to 5 \times S_5$$

Since $\mathrm{SAut}(F_3)$ is perfect (Proposition 2.4), the image of ψ must lie within A_5. Lemma 3.1 tells us that then ψ is trivial. But ψ is a restriction of φ, and so φ trivialises a transvection, say ρ_{67}, and thus φ is trivial. $\square$

Proposition 4.3. *Let $n \geqslant 3$ and K be a sporadic simple group. Then K is not the smallest finite quotient of $\mathrm{SAut}(F_n)$.*

Proof. Let K be a sporadic group, and suppose that it is a smallest finite quotient of $\mathrm{SAut}(F_n)$. We must have

$$|K| \leqslant |\,\mathrm{L}_n(2)|$$

In fact, the inequality is strict, since for each n the group $\mathrm{L}_n(2)$ is not isomorphic to a sporadic group (this is visible in Table 4.1). Thus, by Lemma 2.5, we see that the epimorphism $\varphi\colon \mathrm{SAut}(F_n) \to K$ has to be injective on D'_n. Inspection of Table 4.1 shows that this is only possible for $K = \mathrm{Fi}_{22}$, in which case $n \geqslant 7$. But this is ruled out by Lemma 4.2. $\square$

5. Algebraic groups and groups of Lie type

In this section we review the necessary information about algebraic groups over fields of positive characteristic, and the (closely related) finite groups of Lie type.

5.1. Algebraic groups. We begin by discussing connected algebraic groups. Following [**GLS**], we will denote such groups by $\overline{K}$. We review here only the facts that will be useful to us, focusing on simple, semi-simple, and reductive algebraic groups.

Let r be a prime, and let $\overline{\mathbb{F}}$ be an algebraically closed field of characteristic r. The simple algebraic groups over $\overline{\mathbb{F}}$ are classified by the Dynkin diagrams $\mathtt{A}_n$ (for each n), $\mathtt{B}_n$ (for $n \geqslant 3$), $\mathtt{C}_n$ (for $n \geqslant 2$), $\mathtt{D}_n$ (for $n \geqslant 4$), $\mathtt{E}_n$ (for $n \in \{6,7,8\}$), $\mathtt{F}_4$, and $\mathtt{G}_2$. The index of the diagram is defined to be the *rank* of the associated group.

K	Bound for D'-rank	Order of K
M_{11}	3	7920
$L_4(2)$		21060
M_{12}	3	95040
J_1	4	175560
M_{22}	3	443520
J_2	4	604800
$L_5(2)$		9999360
M_{23}	3	10200960
HS	4	44352000
J_3	4	50232960
M_{24}	3	244823040
M^cL	4	898128000
He	4	4030387200
$L_6(2)$		20158709760
Ru	5	145926144000
Suz	5	448345497600
O′N	4	460815505920
Co_3	5	495766656000
Co_2	6	42305421312000
Fi_{22}	7	64561751654400
$L_7(2)$		163849992929280
HN	6	273030912000000
Ly	5	51765179004000000
Th	6	90745943887872000
Fi_{23}	7	4089470473293004800
Co_1	6	4157776806543360000
$L_8(2)$		5348063769211699200
J_4	7	86775571046077562880
$L_9(2)$		699612310033197642547200
Fi'_{24}	9	1255205709190661721292800
$L_{10}(2)$		366440137299948128422802227200
B	10	4154781481226426191177580544000000
$L_{11}(2)$		768105432118265670534631586896281600
$L_{12}(2)$		6441762292785762141878919881400879415296000
$L_{13}(2)$		2161232893550926958761174333380796550786643394560000
M	12	808017424794512875886459904961710757005754368000000000

TABLE 4.1. Upper bounds for the D'-ranks of the sporadic groups.

To each Dynkin diagram we associate a finite number of simple algebraic groups; two such groups associated to the same diagram are called *versions*; they become isomorphic upon dividing them by their respective finite centres. Two versions are particularly important: the *universal* (or *simply-connected*) one, which maps onto every other version with a finite central kernel, and the *adjoint* version, which is a quotient of every other version with a finite central kernel.

Every semi-simple algebraic group over $\overline{\mathbb{F}}$ is a central product of finitely many simple algebraic groups over $\overline{\mathbb{F}}$. The rank of such a group is defined to be the sum of the ranks of the simple factors, and is well-defined.

Every reductive algebraic group is a product of an abelian group and a semi-simple group. Its rank is defined to be the rank of the semi-simple factor, and again it is well-defined.

Given an algebraic group $\overline{K}$, a maximal with respect to inclusion closed connected solvable subgroup of $\overline{K}$ will be referred to as a *Borel subgroup*. The Borel subgroups always exist, and are conjugate, and hence one can talk about the Borel subgroup (up to conjugation). When $\overline{K}$ is reductive, any closed subgroup thereof containing a Borel subgroup is called *parabolic*. Let us now state the main tool in our approach towards algebraic groups and groups of Lie type.

Theorem 5.1 (Borel–Tits [**GLS**, Theorem 3.1.1(a)])**.** *Let $\overline{K}$ be a reductive algebraic group over an algebraically closed field, let $\overline{X}$ be a closed unipotent subgroup, and let $\overline{N}$ denote the normaliser of $\overline{X}$ in $\overline{K}$. Then there exists a parabolic subgroup $\overline{P} \leqslant \overline{K}$ such that $\overline{X}$ lies in the unipotent radical of $\overline{P}$, and $\overline{N} \leqslant \overline{P}$.*

We will not discuss the other various terms appearing above beyond what is strictly necessary. For our purpose we only need to observe the following.

Remark 5.2. (1) The unipotent radical of a reductive group is trivial.
(2) If $\overline{K}$ is defined in characteristic r, then every finite r-group in $\overline{K}$ is a closed unipotent subgroup.

Theorem 5.3 (Levi decomposition [**GLS**, Theorem 1.13.2, Proposition 1.13.3])**.** *Let $\overline{P}$ be a proper parabolic subgroup in a reductive algebraic group $\overline{K}$.*
(1) Let $\overline{U}$ denote the unipotent radical of $\overline{P}$ (note that $\overline{U}$ is nilpotent). There exists a subgroup $\overline{L} \leqslant \overline{P}$, such that $P = \overline{U} \rtimes \overline{L}$.
(2) The subgroup $\overline{L}$ (the Levi factor*) is a reductive algebraic group of rank smaller than $\overline{K}$.*

5.2. Finite groups of Lie type. Let r be a prime, and q a power thereof.

Any finite group of Lie type K is obtained as a fixed point set of a Steinberg endomorphism of a connected simple algebraic group $\overline{K}$ defined over an algebraically closed field of characteristic r. Such groups have a *type*, which is related to the Dynkin diagram of $\overline{K}$, and an associated *twisted rank*. As stated in Section 1, the finite groups of Lie type fall into two families: the types $\mathtt{A}_n, {}^2\mathtt{A}_n, \mathtt{B}_n, \mathtt{C}_n, \mathtt{D}_n$ and ${}^2\mathtt{D}_n$ are called *classical*, and the types ${}^2\mathtt{B}_2$, ${}^3\mathtt{D}_4$, $\mathtt{E}_6$, ${}^2\mathtt{E}_6$, $\mathtt{E}_7$, $\mathtt{E}_8$, $\mathtt{F}_4$, ${}^2\mathtt{F}_4$, $\mathtt{G}_2$ and ${}^2\mathtt{G}_2$ are called *exceptional*.

The types of the classical groups and their twisted ranks are listed in Table 5.4. Note that $\lceil . \rceil$ denotes the ceiling function. In the case of the exceptional groups, for the groups of types $\mathtt{G}_2$, $\mathtt{F}_4$, $\mathtt{E}_6$, $\mathtt{E}_7$ and $\mathtt{E}_8$ the twisted rank is equal to the rank. Groups of type ${}^2\mathtt{B}_2$ or ${}^2\mathtt{G}_2$ have twisted rank 1, those of type ${}^3\mathtt{D}_4$ or ${}^2\mathtt{F}_4$ have twisted rank 2 and groups of type ${}^2\mathtt{E}_6$ have twisted rank 4. Groups of types ${}^2\mathtt{B}_2$ and ${}^2\mathtt{F}_4$ are defined only over fields of order 2^{2m+1} while groups of type ${}^2\mathtt{G}_2$ are defined only over fields of order 3^{2m+1}. All groups of all other types are defined in all characteristics.

As was the case with algebraic groups, each type corresponds to a finite number of finite groups (the *versions*), and two such are related by dividing by the centre as before. The smallest version (in cardinality, say) is called *adjoint* as before; the adjoint version is a simple group with the following exceptions [**CCN**$^+$, Chapter 3.5]

$$\mathtt{A}_1(2) \cong S_3, \quad \mathtt{A}_1(3) \cong A_4, \quad \mathtt{C}_2(2) \cong S_6, \quad {}^2\mathtt{A}_2(2) \cong 3^2 \rtimes Q_8,$$
$$\mathtt{G}_2(2) \cong {}^2\mathtt{A}_2(3) \rtimes 2, \quad {}^2\mathtt{B}_2(2) \cong 5 \rtimes 4, \quad {}^2\mathtt{G}_2(3) \cong \mathtt{A}_1(8) \rtimes 3, \quad {}^2\mathtt{F}_4(2)$$

where Q_8 denotes the quaternion group of order 8, and the index 2 derived subgroup of ${}^2\mathtt{F}_4(2)$ is simple, known as the Tits group. For the purpose of this paper, we treat T as a finite group of Lie type.

Type	Conditions	Twisted rank	Dimension	Classical isomorphism
$\mathtt{A}_n(q)$	$n \geqslant 1$	n	$n+1$	$\mathrm{L}_{n+1}(q)$
${}^2\mathtt{A}_n(q)$	$n \geqslant 2$	$\lceil \frac{n}{2} \rceil$	$n+1$	$\mathrm{U}_{n+1}(q)$
$\mathtt{B}_n(q)$	$n \geqslant 2$	n	$2n+1$	$\mathrm{O}_{2n+1}(q)$
$\mathtt{C}_n(q)$	$n \geqslant 3$	n	$2n$	$\mathrm{S}_{2n}(q)$
$\mathtt{D}_n(q)$	$n \geqslant 4$	n	$2n$	$\mathrm{O}^+_{2n}(q)$
${}^2\mathtt{D}_n(q)$	$n \geqslant 4$	$n-1$	$2n$	$\mathrm{O}^-_{2n}(q)$

TABLE 5.4. The classical groups of Lie type

For reference, we also recall the following additional exceptional isomorphisms

$$\mathtt{A}_1(4) \cong \mathtt{A}_1(5) \cong A_5, \quad \mathtt{A}_1(9) \cong A_6, \quad \mathtt{A}_1(7) \cong \mathtt{A}_2(2), \quad \mathtt{A}_3(2) \cong A_8, \quad {}^2\mathtt{A}_3(2) \cong \mathtt{C}_2(3)$$

In addition, $\mathtt{B}_n(2^e) \cong \mathtt{C}_n(2^e)$ for all $n \geqslant 3$ and $e \geqslant 1$. We will sometimes abuse the notation, and denote the adjoint version of some type by the type itself.

The adjoint version of a classical group over q comes with a natural projective module over an algebraically closed field in characteristic r; the dimensions of these modules are taken from [**KL**, Table 5.4.C] and listed in Table 5.4. Note that these projective modules are irreducible.

A *parabolic* subgroup of K is any subgroup containing the Borel subgroup of K, which is obtained by taking an α-invariant Borel subgroup in $\overline{K}$ (where α denotes a Steinberg endomorphism), and intersecting it with K. Note that such a Borel subgroup always exists – in fact, its intersection with K is equal to the normaliser of some Sylow r-subgroup of K.

Theorem 5.5 (Borel–Tits [**GLS**, Theorem 3.1.3(a)])**.** *Let K be a finite group of Lie type in characteristic r, and let R be a non-trivial r-subgroup of K. Then there exists a proper parabolic subgroup $P \leqslant K$ such that R lies in the normal r-core of P, and $N_K(R) \leqslant P$.*

Note that the normal r-core of K is trivial.

Theorem 5.6 (Levi decomposition [**GLS**, Theorem 2.6.5(e,f,g), Proposition 2.6.2(a,b)])**.** *Let P be a proper parabolic in a finite group K of Lie type in characteristic r.*

(1) *Let U denote the normal r-core of P (note that U is nilpotent). There exists a subgroup $L \leqslant P$, such that $L \cap U = \{1\}$ and $LU = P$.*
(2) *The subgroup L (the* Levi factor*) contains a normal subgroup M such that L/M is abelian of order coprime to r.*
(3) *The subgroup M is isomorphic to a central product of finite groups of Lie type (the* simple factors *of L) in characteristic r such that the sum of the twisted ranks of these groups is lower than the twisted rank of K.*
(4) *When K is of classical type other than ${}^2\mathtt{D}$ or $\mathtt{B}$, then each simple factor of L is either of the same type as K, or of type $\mathtt{A}$. For type ${}^2\mathtt{D}$ we also get factors of type ${}^2\mathtt{A}_3$; for type $\mathtt{B}$ we also get factors of type $\mathtt{C}_2$.*
(5) *When K is of classical type other than ${}^2\mathtt{A}$ or ${}^2\mathtt{D}$, then the simple factors of L are defined over the same field. The groups ${}^2\mathtt{A}(q)$ admit simple factors of L of type $\mathtt{A}(q^2)$, and the groups ${}^2\mathtt{D}(q)$ admit a simple factor of L of type $\mathtt{A}_1(q^2)$.*

6. Groups of Lie type in characteristic 2

Because of the special role the involutions $\epsilon_1, \ldots, \epsilon_n$ play in the structure of $\mathrm{Aut}(F_n)$, groups of Lie type in characteristic 2 require a different approach than groups in odd characteristic. The strategy is to look at the centraliser of ϵ_n in

$\mathrm{Aut}(F_n)$, note that it contains $\mathrm{Aut}(F_{n-1})$, and then use the Borel-Tits theorem (Theorem 5.5) for its image. The same strategy works for reductive algebraic groups in characteristic 2.

Before we proceed to the main part of this section, we will investigate maps $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$, with the aim of showing that there is essentially only one such non-trivial map.

Lemma 6.1 ([**KL**, Proposition 5.3.7]). *Let A_n be the alternating group of degree n where $3 \leqslant n \leqslant 8$. The degree $R_p(A_n)$ of the smallest nontrivial irreducible projective representation of A_n over a field of characteristic p is as given in Table 6.2. If $n \geqslant 9$, then the degree of the smallest nontrivial projective representation of A_n is $n-2$.*

n	$R_2(A_n)$	$R_3(A_n)$	$R_5(A_n)$	$R_7(A_n)$
5	2	2	2	2
6	3	2	3	3
7	4	4	3	4
8	4	7	7	7

TABLE 6.2.

Remark 6.3. Moreover, the result of Wagner [**Wag**] tells us the following: Assume that the characteristic is 2 and that $n \geqslant 9$. Then the smallest non-trivial A_n-module appears in dimension $n-1$ when n is odd and $n-2$ when n is even, and is unique. This module appears as an irreducible module in our group $D_n' = 2^{n-1} \rtimes A_n$, where for n even we take a quotient by $\langle \delta \rangle$.

Lemma 6.4. *Let $n \geqslant 4$. Let $\varphi\colon D_n' \to \mathrm{L}_m(2) = \mathrm{GL}(V)$ be a homomorphism.*

(1) *When $m < n-1$ then $\varphi(2^{n-1})$ is trivial.*
(2) *Suppose that $m = n-1$ and $\varphi(2^{n-1})$ is non-trivial. Then n is even, $\varphi(\delta) = 1$ and we can choose a basis of V in such a way that either for $i < n-1$ the element $\varphi(\epsilon_i\epsilon_{i+1})$ is given by the elementary matrix E_{1i}, that is the matrix equal to the identity except at the position $(1i)$, or each element $\varphi(\epsilon_i\epsilon_{i+1})$ is given by the elementary matrix E_{i1}.*
(3) *When $m = n$ and we additionally assume that φ is injective and that when $n = 8$ the representation $\varphi|_{A_8}$ is the 8-dimensional permutation representation, then we can choose a basis of V in such a way that either for each $i < n-2$ the element $\varphi(\epsilon_i\epsilon_{i+1})$ is given by the elementary matrix E_{1i}, or each element $\varphi(\epsilon_i\epsilon_{i+1})$ is given by the elementary matrix E_{i1}.*

Proof. Fix n, and proceed by induction on m. Clearly $m > 1$.

Consider the subgroup $V \rtimes \varphi(2^{n-1}) < V \rtimes \mathrm{GL}(V)$. It is a 2-group, hence it is nilpotent, and therefore it has a non-trivial centre Z. Since $\varphi(2^{n-1})$ acts faithfully, we have $Z \leqslant V$ as a subgroup, and hence also as a 2-vector subspace. Clearly,

$$Z = \{v \in V \mid \varphi(\xi)(v) = v \text{ for all } \xi \in 2^{n-1}\}$$

and therefore Z is preserved setwise by $\varphi(D_n')$, as 2^{n-1} is a normal subgroup of D_n'.

Suppose that $\dim Z \leqslant \dim V/Z = m - \dim Z$. If $n \geqslant 5$, this implies that Z is a trivial A_n-module: for $n \geqslant 9$ and $n = 7$ this follows from Lemma 6.1, for $n \in \{5, 6\}$ we observe that A_n is simple and larger in cardinality than $\mathrm{L}_{n-3}(2)$; and $n = 8$ we see that $A_8 \cong \mathrm{L}_4(2)$ is larger than $\mathrm{L}_3(2)$, which is enough for (1) and (2); for (3) we us the additional hypothesis on the A_8-representation.

When $n = 4$ we could have $m = n$ and $\dim Z = 2$, in which case Z does not have to be a trivial A_4-module. But in this case we have

$$\mathrm{GL}(Z) \cong \mathrm{GL}(V/Z) \cong \mathrm{L}_2(2) \cong S_3$$

and every homomorphism $D'_4 \to S_3$ has $2^3 \rtimes V_4$ in its kernel, where V_4 denotes the Klein four-group. But then φ takes $2^3 \rtimes V_4$ to an abelian group of matrices which differ from the identity only in the top-right 2×2 corner. Thus $\varphi(\delta) = \varphi([\epsilon_1\epsilon_2, \sigma_{13}\sigma_{24}]) = 1$, contradicting the injectivity of φ.

We may therefore assume that Z is a trivial A_n module even when $n = 4$.

Suppose that $m < n - 1$. Then, by the inductive hypothesis, we know that the action of 2^{n-1} on V/Z is trivial; it is also trivial on Z by construction. Hence $\varphi(D'_n)$ is a subgroup of

$$2^{\dim Z(m-\dim Z)} \rtimes \mathrm{GL}(V/Z)$$

and φ takes 2^{n-1} into the $2^{\dim Z(m-\dim Z)}$ part. But this subgroup cannot contain 2^{n-1} as an A_n-module, since as a $\mathrm{GL}(V/Z)$-module it is a direct sum of $(m - \dim Z)$-dimensional modules, and $\dim Z \geqslant 1$. This shows that φ is not injective on 2^{n-1}. But the only subgroup of 2^{n-1} which can lie in $\ker \varphi$ is $\langle \delta \rangle$, and therefore if $\varphi(2^{n-1})$ is not trivial, then we need to be able to fit a $(n-2)$-dimensional module into $2^{\dim Z(m-\dim Z)}$. This is impossible when $m < n - 1$, and so (1) follows.

All of the above was conducted under the assumption that $\dim Z \leqslant m - \dim Z$. If this is not true, then we take the transpose inverse of φ; for this representation the inequality is true, and the kernel of this representation coincides with the kernel of φ.

When $m = n - 1$ then we have just proven (2) – it is clear that we can change the basis of V if necessary to have each $\varphi(\epsilon_i\epsilon_{i+1})$ as required.

In case (3) we immediately see that $\dim Z = 1$. If 2^{n-1} does not act trivially on V/Z then we apply (2). Since now φ is injective, it must take 2^{n-1} to the subgroup of $\mathrm{L}_n(2)$ generated by E_{ji} with $j \in \{1, 2\}$ and $i > j$. Suppose that for some i we have

$$\varphi(\epsilon_i\epsilon_{i+1}) = E_{12} + E_{2(i+1)} + M$$

where $M \in \langle \{E_{1i} \mid i > 2\} \rangle$. Then $\varphi(\epsilon_i\epsilon_{i+1})$ is of order 4, which is impossible. So

$$\varphi(2^{n-1}) \leqslant \langle \{E_{ji} \mid j \in \{1, 2\}, i > 2\} \rangle \cong 2^{n-2} \oplus 2^{n-2}$$

as an A_n-module, which contradicts injectivity of φ. Therefore 2^{n-1} acts trivially on V/Z, and the result follows as before. □

Remark 6.5. In fact, for $n = 3$ we can obtain identical conclusions, with the exception that in (3) we may need to postcompose φ with an outer automorphism of $\mathrm{L}_3(2)$. To see this note that in (1) we have $\mathrm{L}_m(2) = \mathrm{L}_1(2) = \{1\}$; in (2) we have $\mathrm{L}_m(2) = \mathrm{L}_2(2) \cong S_3$, and every map from $\mathcal{D}'_3 \cong A_4$ to S_3 has $2^2 \cong V_4$ in the kernel. For (3) we see that $\mathrm{L}_m(2) = \mathrm{L}_3(2)$ contains exactly two conjugacy classes of $A_4 \cong D'_3$, and these are related by an outer automorphism of $\mathrm{L}_3(2)$ (see [**CCN**$^+$]). Thus, up to postcomposing φ with an outer automorphism of $\mathrm{L}_3(2)$, we may assume that φ maps the involutions $\epsilon_i\epsilon_j$ in the desired manner.

Proposition 6.6. *Let $n \geqslant 3$ and $m \leqslant n$ be integers. If*

$$\varphi\colon \mathrm{SAut}(F_n) \to \mathrm{L}_m(2)$$

is a non-trivial homomorphism, then $m = n$ and φ is equal to the natural map $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ postcomposed with an automorphism of $\mathrm{L}_n(2)$.

Proof. Assume that $n \geqslant 3$. Observe that if φ is not injective on D_n' then we are done by Lemma 2.5 – one has to note that when $\varphi(\delta) = 1$ then we know that φ factors through $\mathrm{SL}_n(\mathbb{Z})$, and so we know using the congruence subgroup property that every non-trivial map $\mathrm{SL}_n(\mathbb{Z}) \to \mathrm{L}_n(2)$ factors through the natural such map. We will assume that φ is injective on D_n'.

We apply Lemma 6.4 (and Remark 6.5 when $n = 3$); for $n = 8$ we consider $\varphi(A_9)$ – by Remark 6.3, φ must be the unique 8-dimensional representation, and so as an A_8-module V is the natural permutation representation. Up to possibly taking the transpose inverse of φ, we see that $m = n$, and

$$\varphi(\epsilon_i \epsilon_{i+1}) = E_{1i}$$

Let Z denote the subspace of V generated by the first basis vector, note that Z is precisely the centraliser in V of $\varphi(2^{n-1})$ and coincides with the commutator $[V, \varphi(\xi)]$ for every $\xi \in 2^{n-1} \smallsetminus \{1\}$. A direct computation shows that if a matrix in $\mathrm{L}_n(2)$ commutes with some E_{1i}, then it preserves Z. In fact this remains true for any non-zero sum of matrices E_{1i}.

The group $\mathrm{SAut}(F_n)$ is generated by transvections, and each of them commutes with some $\epsilon_i \epsilon_j$ as $n \geqslant 4$, and so $\mathrm{SAut}(F_n)$ preserves Z. Thus we have a representation

$$\mathrm{SAut}(F_n) \to \mathrm{GL}(V/Z) \cong \mathrm{L}_{n-1}(2)$$

and such a representation is trivial, or n is even and the representation has δ in its kernel by Lemma 6.4. But then it factors through $\mathrm{SL}_n(\mathbb{Z})$, and therefore must be trivial, since the smallest quotient of $\mathrm{SL}_n(\mathbb{Z})$ is $\mathrm{L}_n(2)$.

Therefore φ take $\mathrm{SAut}(F_n)$ to 2^{n-1}, and hence must be trivial. □

We now proceed to the main discussion. We start by looking at small values of n. These considerations will form the base of our induction.

Lemma 6.7. *Let $\overline{\mathbb{F}}$ be an algebraically closed field of characteristic 2. Every homomorphism $\varphi\colon \mathrm{SAut}(F_3) \to \mathrm{L}_2(\overline{\mathbb{F}}) = \mathrm{PSL}_2(\overline{\mathbb{F}})$ is trivial.*

Proof. We start by observing that $\mathrm{PSL}_2(\overline{\mathbb{F}}) = \mathrm{SL}_2(\overline{\mathbb{F}})$, since the only element in $\overline{\mathbb{F}}$ which squares to 1 is 1 itself.

Suppose first that $\epsilon_1 \epsilon_2$ lies in the kernel of φ. Then φ descends to a map

$$\mathrm{L}_3(2) \to \mathrm{L}_2(\overline{\mathbb{F}})$$

by Lemma 2.5. Since $\mathrm{L}_3(2)$ is simple, this map is either faithful or trivial. But it cannot be faithful, since the upper triangular matrices in $\mathrm{L}_3(2)$ form a 2-group (the dihedral group of order 8) which is nilpotent of class 2, whereas every non-trivial 2-subgroup of $\mathrm{L}_2(\overline{\mathbb{F}})$ is abelian. (Alternatively, one can use the fact that $\mathrm{L}_3(2)$ has no non-trivial projective representations in dimension 2 in characteristic 2, as can be seen from the 2-modular Brauer table which exists in GAP.)

Hence we may assume that $\varphi(\epsilon_1 \epsilon_2) \neq 1$. Consider the 2-subgroup of $\mathrm{L}_2(\overline{\mathbb{F}})$ generated by $\varphi(\epsilon_i \epsilon_j)$ with $1 \leqslant i, j \leqslant 3$ (it is isomorphic to 2^2). As before, up to conjugation, this subgroup lies within the unipotent subgroup of upper triangular matrices with ones on the diagonal. Now a direct computation shows that the matrices which commute with

$$\begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}$$

with $x \neq 0$ are precisely the matrices of the form

$$\begin{pmatrix} * & * \\ 0 & * \end{pmatrix}$$

In particular, this implies that if an element of $\mathrm{L}_2(\overline{\mathbb{F}})$ commutes with $\varphi(\epsilon_1\epsilon_2)$, then it also commutes with $\varphi(\epsilon_2\epsilon_3)$. This applies to $\varphi(\sigma_{12}\epsilon_3)$, and so

$$\varphi(\epsilon_2\epsilon_3) = \varphi(\epsilon_2\epsilon_3)^{\varphi(\sigma_{12}\epsilon_3)} = \varphi(\epsilon_1\epsilon_3)$$

and so $\varphi(\epsilon_1\epsilon_2) = 1$, contradicting our assumption. □

Lemma 6.8. *Every homomorphism φ from $\mathrm{Aut}(F_3)$ to a finite group of Lie type in characteristic 2 of twisted rank 1 has abelian image.*

Proof. The groups we have to consider as targets here are the versions of $\mathtt{A}_1(q)$, ${}^2\mathtt{A}_2(q)$, and ${}^2\mathtt{B}_2(q)$, where q is a power of 2 (where the exponent is odd in type ${}^2\mathtt{B}_2$). Since $\mathrm{SAut}(F_3)$ is perfect, and we claim that it has to be contained in the kernel of our homomorphism, we need only look at the adjoint versions.

For type $\mathtt{A}_1$ the result follows from Lemma 6.7, since $\mathrm{L}_2(q) \leqslant \mathrm{L}_2(\overline{\mathbb{F}})$ with $\overline{\mathbb{F}}$ algebraically closed and of characteristic 2. The simple group of type ${}^2\mathtt{B}_2(q)$ has no elements of order 3 (this can easily be seen from the order of the group), and so $\mathrm{SAut}(F_3)$ lies in the kernel of the homomorphism by Lemma 2.5.

We are left with the type ${}^2\mathtt{A}_2(q)$. In this case we observe that, up to conjugation, there are only two parabolic subgroups of $K = {}^2\mathtt{A}_2(q)$, namely K itself and a Borel subgroup B.

Suppose first that ϵ_3 has a non-trivial image in K. By Theorem 5.5, the image of the centraliser of ϵ_3 in $\mathrm{Aut}(F_n)$ lies in B. The Borel subgroup B is a semi-direct product of the unipotent subgroup by the torus. The torus contains no elements of order 2. Moreover, the only elements of order 2 in the unipotent subgroup lie in its centre – this can be verified by a direct computation with matrices.

The centraliser of ϵ_3 in $\mathrm{Aut}(F_3)$ contains $\mathrm{Aut}(F_2)$, which is generated by involutions $\epsilon_1, \epsilon_2, \rho_{12}\epsilon_2$ and $\rho_{21}\epsilon_1$. Thus the image of $\mathrm{Aut}(F_2)$ lies in the centre of the unipotent subgroup of B, which is abelian. Therefore, we have

$$\varphi(\rho_{12}) = \varphi({\rho_{12}}^{\epsilon_1\epsilon_2\sigma_{12}}) = \varphi(\lambda_{21})$$

and therefore

$$\varphi(\rho_{13})^{-1} = [\varphi(\rho_{12})^{-1}, \varphi(\rho_{23})^{-1}] = [\varphi(\lambda_{21})^{-1}, \varphi(\rho_{23})^{-1}] = 1$$

This trivialises the subgroup $\mathrm{SAut}(F_3)$ as claimed.

Recall that we have assumed that ϵ_3 is not in the kernel of φ; when it is, then the homomorphism factors through $\mathrm{L}_3(2)$, which is simple and not a subgroup of ${}^2\mathtt{A}_2(q)$ whenever q is a power of 2 – this can be seen by inspecting the maximal subgroups of ${}^2\mathtt{A}_2(q)$ [**BHRD**]. □

Theorem 6.9. *Let $n \geqslant 3$. Let K be a finite group of Lie type in characteristic 2 of twisted rank less than $n-1$, and let $\overline{K}$ be a reductive algebraic group over an algebraically closed field of characteristic 2 of rank less than $n-1$. Then any homomorphism $\mathrm{Aut}(F_n) \to K$ or $\mathrm{Aut}(F_n) \to \overline{K}$ has abelian image, and any homomorphism $\mathrm{SAut}(F_{n+1}) \to K$ or $\mathrm{SAut}(F_{n+1}) \to \overline{K}$ is trivial.*

Proof. We start by looking at the finite group K, and a homomorphism

$$\varphi\colon \mathrm{Aut}(F_n) \to K$$

Since $\mathrm{SAut}(F_n)$ is perfect and of index 2 in $\mathrm{Aut}(F_n)$, we may without loss of generality divide K by its centre; we may also assume that K is not solvable.

Our proof is an induction on n. The base case ($n = 3$) is covered by Lemma 6.8. In what follows let us assume that $n > 3$.

We claim that $\varphi(\mathrm{SAut}(F_{n-1}))$ lies in a proper parabolic subgroup P of K. If $\varphi(\epsilon_n)$ is central then $\varphi(\epsilon_{n-1}\epsilon_n)$ is trivial, since ϵ_n and ϵ_{n-1} are conjugate. Thus φ factors through

$$\mathrm{Aut}(F_n) \to \mathrm{L}_n(2)$$

by Lemma 2.5. Let $\eta\colon \mathrm{L}_n(2) \to K$ denote the induced homomorphism.

The group $\mathrm{L}_n(2)$ contains $\mathrm{L}_{n-1}(2)$ inside a proper parabolic subgroup which normalises a non-trivial 2-group G. This 2-group contains an elementary matrix, and so if $\eta(G)$ is trivial, then so is every elementary matrix in $\mathrm{L}_n(2)$, and therefore η is trivial (as $\mathrm{L}_n(2)$ is generated by elementary matrices). This proves the claim.

Now let us assume that G has a non-trivial image in K. Thus $\eta(G)$ does not lie in the normal 2-core of K, and therefore, by Theorem 5.5, the normaliser of G in $\mathrm{L}_n(2)$ is mapped by η into a proper parabolic subgroup P. Clearly, we may choose G so that it is normalised by the image of $\mathrm{Aut}(F_{n-1})$ in $\mathrm{L}_n(2)$. This way we have shown that $\varphi(\mathrm{Aut}(F_{n-1}))$ lies in P.

Now assume that $\varphi(\epsilon_n)$ is not central, and so in particular not trivial. We conclude, using Theorem 5.5, that $\varphi(\mathrm{Aut}(F_{n-1}))$ lies in a parabolic P inside K such that $P \neq K$. Hence we have

$$\varphi(\mathrm{Aut}(F_{n-1})) \leqslant P < K$$

irrespectively of what happens to ϵ_n, which proves the claim.

Consider the induced map $\psi\colon \mathrm{Aut}(F_{n-1}) \to P/U \cong L$ (using the notation of Theorem 5.6). Note that in fact the image of ψ lies in M, since L/M is abelian and contains no element of order 2. Now M is a central product of finite groups of Lie type in characteristic 2, where the sum of the twisted ranks is lower than that of K. Thus, using the projections, we get maps from $\mathrm{Aut}(F_{n-1})$ to finite groups of Lie type of twisted rank less than $n-2$. By the inductive assumption all such maps have abelian image, and so the image of $\mathrm{Aut}(F_{n-1})$ in M is abelian. This forces φ to contain $\mathrm{SAut}(F_{n-1})$ in its kernel, and the result follows, since U is nilpotent and $\mathrm{SAut}(F_{n-1})$ is perfect as $n \geqslant 4$.

Now let us look at a homomorphism $\varphi\colon \mathrm{Aut}(F_n) \to \overline{K}$. We proceed as above; the base case ($n = 3$) is covered by Lemma 6.7.

We claim that, as before, $\varphi(\mathrm{Aut}(F_{n-1}))$ is contained in a proper parabolic subgroup $\overline{P}$ of $\overline{K}$. This is proved exactly as before using Theorem 5.1, except that now we use the fact that every finite 2-group in $\overline{K}$ is a closed unipotent subgroup. Note that $\overline{P}$ is a proper subgroup, since $\overline{K}$ is reductive, and thus its unipotent radical is trivial.

Again as before we look at the induced map $\psi\colon \mathrm{Aut}(F_{n-1}) \to \overline{P}/\overline{U} \cong \overline{L}$. By Theorem 5.3, the group L is reductive of lower rank, and so ψ has abelian image by induction. But then $\varphi|_{\mathrm{Aut}(F_{n-1})}$ has solvable image, and so $\varphi(\mathrm{SAut}(F_{n-1})) = \{1\}$. Therefore

$$\varphi(\mathrm{SAut}(F_n)) = \{1\}$$

as well, and the image of φ is abelian.

The statements for $\mathrm{SAut}(F_{n+1})$ follow from observing that the natural embedding $\mathrm{SAut}(F_n) \hookrightarrow \mathrm{SAut}(F_{n+1})$ extends to an embedding $\mathrm{Aut}(F_n) \hookrightarrow \mathrm{SAut}(F_{n+1})$, where we map an element $x \in \mathrm{Aut}(F_n)$ of determinant -1 to $x\epsilon_{n+1}$. When this copy of $\mathrm{Aut}(F_n)$ has an abelian image under a homomorphism, then the homomorphism is trivial on $\mathrm{SAut}(F_n)$, and hence on the whole of $\mathrm{SAut}(F_{n+1})$. □

Theorem 6.10. *Let $n \geqslant 8$. Let K be a finite simple group of Lie type in characteristic 2 which is a quotient of $\mathrm{SAut}(F_n)$. Then either $|K| > |\mathrm{L}_n(2)|$, or $K = \mathrm{L}_n(2)$ and φ is obtained by postcomposing the natural map $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ by an automorphism of $\mathrm{L}_n(2)$.*

Proof. By Theorem 6.9, K is of twisted rank at least $n-2$.

Since $n \geqslant 8$, by Lemma 10.4 we see that all the finite simple groups of Lie type in characteristic 2 and twisted rank at least $n-2$ are larger than $\mathrm{L}_n(2)$, with

TABLE 7.2. Some conjugacy classes in $\mathtt{C}_2(3)$

Class	$2A$	$2B$	$3C$	$3D$	$4A$	$4B$	$5A$	$6E$	$6F$
$\lvert x^G \cap C_G(x)\rvert$	13	22	6	12	8	4	4	2	2

the exception of $\mathtt{A}_{n-2}(2)$ and $\mathtt{A}_{n-1}(2)$. Proposition 6.6 immediately tells us that $K = \mathrm{L}_n(2)$ and φ is as claimed. □

7. Classical groups in odd characteristic

7.1. Field of 3 elements. In this subsection we use Borel–Tits in characteristic 3. To do this, we need to find suitable elements of order 3 in $\mathrm{SAut}(F_n)$.

Let $\gamma = \epsilon_{n-1}\epsilon_n{\lambda_{(n-1)n}}^{-1}\rho_{n(n-1)}$. A direct computation immediately shows that γ is of order 3. Also, the centraliser of γ in $\mathrm{SAut}(F_n)$ contains $\mathrm{SAut}(F_{n-2})$. In fact, γ is the element constructed in Lemma 2.8. We define it here algebraically, since it allows us to easily show the following.

Lemma 7.1. *Let $n \geqslant 4$. The normal closure of γ inside $\mathrm{SAut}(F_n)$ is the whole of $\mathrm{SAut}(F_n)$.*

Proof. Let C denote the normal closure. Then

$$\begin{aligned}{\rho_{n1}}^{-1}C &= [{\rho_{n(n-1)}}^{-1}, {\rho_{(n-1)1}}^{-1}]C\\ &= [\epsilon_{n-1}\epsilon_n{\lambda_{(n-1)n}}^{-1}, {\rho_{(n-1)1}}^{-1}]C\\ &= \lambda_{(n-1)1}\rho_{(n-1)1}C\end{aligned}$$

where the last equality follows by expanding the commutator. Now

$$\begin{aligned}{\rho_{21}}^{-1}C &= [{\rho_{2n}}^{-1}, {\rho_{n1}}^{-1}]C\\ &= [{\rho_{2n}}^{-1}, \lambda_{(n-1)1}\rho_{(n-1)1}]C\\ &= C\end{aligned}$$

and we are done. □

Lemma 7.3. *Every homomorphism $\varphi\colon \mathrm{SAut}(F_4) \to K$, where K is a finite group of Lie type of type $\mathtt{A}_2(3)$ or $\mathtt{C}_2(3)$, is trivial.*

Proof. Since $\mathrm{SAut}(F_4)$ is perfect, we may assume that K is simple. If K is of type $\mathtt{A}$, then $K \cong \mathrm{L}_3(3)$ which has no element of order 5 – this follows immediately from the order of the group. But then φ trivialises the five cycle in A_5, and so φ is trivial by Lemma 2.5.

Suppose that K is of type $\mathtt{B}$. We are now going to argue as in the proof of Lemma 3.1. Consider the set of transvections

$$T = \{{\rho_{ij}}^{\pm 1}, {\lambda_{ij}}^{\pm 1}\}$$

Recall that any two elements in T are conjugate in $\mathrm{SAut}(F_4)$ (by Lemma 2.3). Let $C_T(\rho_{12})$ denote the set of elements in T which commute with ρ_{12}. There are exactly 24 elements in $C_T(\rho_{12})$, namely

$$\{{\rho_{12}}^{\pm1}, {\lambda_{12}}^{\pm1}, {\lambda_{13}}^{\pm1}, {\lambda_{14}}^{\pm1}, {\rho_{32}}^{\pm1}, {\lambda_{32}}^{\pm1}, {\rho_{42}}^{\pm1}, {\lambda_{42}}^{\pm1}, {\rho_{34}}^{\pm1}, {\lambda_{34}}^{\pm1}, {\rho_{43}}^{\pm1}, {\lambda_{43}}^{\pm1}\}$$

Table 7.2 lists every conjugacy class in K conjugate to its own inverse, as can be computed in GAP; it also lists the number of elements in the conjugacy class which commute with a fixed representative of the class.

Note that if $\varphi(\rho_{12})$ is an involution, then a direct computation with GAP reveals that φ factors through

$$\mathrm{SAut}(F_4)/\langle\langle {\rho_{12}}^2\rangle\rangle \cong 2^4 \rtimes \mathrm{L}_4(2)$$

(Note that an analogous statement is true for $n = 3$, but for large enough n the quotient is infinite, as shown in [**BV1**].) The group K is simple and non abelian, and so φ factors through $\mathrm{L}_4(2)$. Hence φ is trivial, as $\mathrm{L}_4(2)$ is simple and not isomorphic to K.

We may thus assume that $\varphi(\rho_{12})$ is not an involution. Inspecting Table 7.2 we see that there are at most 12 elements in the conjugacy class of $\varphi(\rho_{12})$ which commute with $\varphi(\rho_{12})$. Thus there exist two elements in $C_T(\rho_{12})$ which get identified under φ. Without loss of generality we may assume that we have

$$\varphi(\rho_{12}) = \varphi({x_{ij}}^{\pm 1})$$

where x is either ρ or λ, and ${x_{ij}}^{\pm 1} \notin \{\rho_{12}, {\rho_{12}}^{-1}\}$.

If $j > 2$, take k such that $k \in \{2, 3, 4\} \smallsetminus \{i, j\}$. Now

$$\varphi({x_{ik}}^{-1}) = \varphi([{x_{ij}}^{-1}, {x_{jk}}^{-1}]) = [\varphi(\rho_{12})^{\mp 1}, \varphi(x_{jk})^{-1}] = 1$$

and so φ is trivial. Let us assume that $j \leqslant 2$.

Similarly, if $i > 1$, take $k \in \{3, 4\} \smallsetminus \{i\}$. Now

$$\varphi({x_{kj}}^{-1}) = \varphi([{x_{ki}}^{-1}, {x_{ij}}^{-1}]) = [\varphi(x_{ki})^{-1}, \varphi(\rho_{12})^{\mp 1}] = 1$$

and so φ is trivial.

We are left with the case $(i, j) = (1, 2)$ and $x = \lambda$. If ${x_{ij}}^{\pm 1} = \lambda_{12}$ then φ factors through $\mathrm{SL}_4(\mathbb{Z})$, since adding the relation $\rho_{12}{\lambda_{12}}^{-1}$ takes the Gersten's presentation of $\mathrm{SAut}(F_n)$ to the Steinberg's presentation of $\mathrm{SL}_n(\mathbb{Z})$. But we know all the finite simple quotients of $\mathrm{SL}_4(\mathbb{Z})$, and K is not one of them. Hence φ is trivial.

We are left with the case ${x_{ij}}^{\pm 1} = {\lambda_{12}}^{-1}$. Gersten's presentation contains the relation

$$(\rho_{12}{\rho_{21}}^{-1}\lambda_{12})^4$$

Using the relation $\rho_{12}\lambda_{12}$ gives

$$(\rho_{12}{\rho_{21}}^{-1}{\rho_{12}}^{-1})^4$$

which is equivalent to ${\rho_{21}}^4$. Thus $\varphi(\rho_{21})$, and hence also $\varphi(\rho_{12})$, has order 4. Inspecting Table 7.2 again we see that in fact we have at least three elements in $C_T(\rho_{12})$ which coincide under φ, and so, without loss of generality, there exists ${x_{ij}}^{\pm 1} \notin \{\rho_{12}, {\rho_{12}}^{-1}, {\lambda_{12}}^{-1}\}$ such that

$$\varphi(\rho_{12}) = \varphi({x_{ij}}^{\pm 1})$$

Thus we are in one of the cases already considered. ☐

Lemma 7.4. *Every homomorphism $\varphi\colon \mathrm{SAut}(F_5) \to K$ where K is a finite group of Lie type of type $\mathtt{A}_3(3)$, ${}^2\mathtt{A}_3(3)$ or $\mathtt{C}_3(3)$ is trivial.*

Proof. As always, we assume that K is simple. The simple group $\mathtt{A}_3(3) \cong \mathrm{L}_4(3)$ contains two conjugacy classes of involutions [**CCN**$^+$] where they are denoted $2A$ and $2B$. The ATLAS also gives the order of their centralisers. The centraliser of an involution in class $2B$ has order $1152 = 2^7 3^2$ and hence is solvable by Burnside's $p^a q^b$-Theorem. The structure of the centraliser of an involution in class $2A$ is given in [**CCN**$^+$] and is isomorphic to

$$(4 \times A_6) \rtimes 2$$

Consider $\epsilon_4\epsilon_5 \in \mathrm{SAut}(F_5)$. If $\varphi(\epsilon_4\epsilon_5) = 1$ then φ factors through $\mathrm{L}_5(2)$ (by Lemma 2.5), which is simple and non-isomorphic to $\mathtt{A}_2(3)$. This trivialises φ.

If $\varphi(\epsilon_4\epsilon_5) \neq 1$ then φ maps $\mathrm{SAut}(F_3)$ (which centralises $\epsilon_4\epsilon_5$) to either a solvable group, or to $(4 \times A_6) \rtimes 2$. In both cases we have $\mathrm{SAut}(F_3) \leqslant \ker\varphi$, as $\mathrm{SAut}(F_3)$ is perfect and has no non-trivial homomorphisms to A_6 by Lemma 3.1. This trivialises φ.

The simple group ${}^2\mathtt{A}_3(3)$ has a single conjugacy class of involutions, denoted $2A$ in [**CCN**$^+$], and the centraliser of an involution in this class again has order 1152, hence it is solvable, and so we argue as before.

The conjugacy classes of maximal subgroups of the simple group $\mathtt{C}_3(3)$ are known [**CCN**$^+$, pg.113]. By inspection we see that it does not contain D_5' and so φ factors through $\mathrm{L}_5(2)$ by Lemma 2.5. But $\mathrm{L}_5(2)$ contains an element of order 31, whereas $\mathtt{C}_3(3)$ does not. Thus φ is trivial. □

Lemma 7.5. *Let $n \geqslant 4$. Every homomorphism $\varphi\colon \mathrm{SAut}(F_n) \to K$ is trivial, where*

(1) *n is even, and K is the of type $\mathtt{A}_k(3)$ or $\mathtt{B}_k(3)$ or $\mathtt{C}_2(3)$ with $k \leqslant \frac{n}{2}$; or*
(2) *n is odd, and K is of type $\mathtt{A}_k(3), {}^2\mathtt{A}_3(3), \mathtt{C}_k(3), \mathtt{D}_k(3)$ or ${}^2\mathtt{D}_k(3)$ with $k \leqslant \frac{n+1}{2}$.*

Proof. As usual, since $\mathrm{SAut}(F_n)$ is perfect, we may assume that we are dealing with adjoint versions; therefore we will use type to denote its adjoint version.

The proof is an induction; the base case when n is even is covered by Lemma 7.3, upon noting that for $k = 1$ we only have to consider $\mathtt{A}_1(3)$, which is solvable.

When n is odd, the base case consists of the groups $\mathtt{A}_1(3)$, $\mathtt{A}_2(3)$, $\mathtt{A}_3(3)$, $\mathtt{C}_2(3)$, $\mathtt{C}_3(3)$, and ${}^2\mathtt{A}_3(3)$. The first two are subgroups of the third, which is covered by Lemma 7.4, and so is $\mathtt{C}_3(3)$. The group $\mathtt{C}_2(3)$ is covered by Lemma 7.3. The remaining group ${}^2\mathtt{A}_3(3)$ is again covered by Lemma 7.4.

Now suppose that $n > 4$. Consider $\varphi(\gamma)$. If this is trivial, then we are done by Lemma 7.1. Otherwise, Theorem 5.5 tells us that φ maps $\mathrm{SAut}(F_{n-2})$ to a parabolic subgroup P of K. We will now use the notation of Theorem 5.6.

Let
$$\psi\colon \mathrm{SAut}(F_{n-2}) \to L$$
be the map induced by taking the quotient $P \to P/U \cong L$. Since L/M is abelian, and $\mathrm{SAut}(F_{n-2})$ is perfect, we immediately see that $\operatorname{im}\psi \leqslant M$.

Suppose that n is even. Then M admits projections onto groups of type $\mathtt{A}_l(3)$ or $\mathtt{B}_l(3)$ with $l < k$ or $\mathtt{C}_2(3)$. The inductive hypothesis shows that $\psi(\mathrm{SAut}(F_{n-2}))$ lies in the intersection of the kernels of such projections. But M is a central product of the images of these projections, and so ψ is trivial. But then φ trivialises $\mathrm{SAut}(F_{n-2})$, and the result follows.

When n is odd the situation is similar: the group M admits projections to groups of type $\mathtt{A}_l(3)$, ${}^2\mathtt{A}_3(3)$, $\mathtt{C}_l(3), \mathtt{D}_l(3), {}^2\mathtt{D}_l(3)$ or $\mathtt{A}_1(9)$. The last group is isomorphic to A_6, and every homomorphism from $\mathrm{SAut}(F_3)$ to A_6 is trivial by Lemma 3.1. The other groups are covered by the inductive hypothesis, and we conclude as before. □

Remark 7.6. In fact the groups of type $\mathtt{A}_k(3)$ are not quotients of $\mathrm{SAut}(F_n)$ when $k \leqslant n-2$ which will become clear in the following section.

7.2. Representations of D_n'.

Our aim now is to control projective representations of $\mathrm{SAut}(F_n)$ in small dimensions over fields of odd characteristic. To do this we will first develop some representation theory of the subgroup D_n'.

Definition 7.7. The action of D_n' on $\mathbb{Z}^n$ obtained by abelianising F_n is the *standard action*. Tensoring $\mathbb{Z}^n$ with a field $\mathbb{F}$ gives us the *standard D_n'-module* $\mathbb{F}^n$, and the image of the generators $a_1, \dots, a_n$ in $\mathbb{F}^n$ is the *standard basis*.

Definition 7.8. Let π be a representation of 2^{n-1}. We set
$$E_I = \{v \in V \mid \pi(\epsilon_i\epsilon_j)(v) = (-1)^{\chi_I(i)+\chi_I(j)}v\}$$
with χ_I standing for the characteristic function of $I \subseteq N$.

Note that $E_I = E_{N \smallsetminus I}$, but otherwise these subspaces intersect trivially.

Lemma 7.9. *Let $n \geqslant 7$. Let $\pi \colon D_n' \to \mathrm{GL}(V)$ be a linear representation of D_n' over a field of characteristic other than 2 in dimension $k < 2n$, such that there is no vector fixed by all elements $\pi(\epsilon_i \epsilon_j)$. Then $k = n$ and π is the standard representation.*

Proof. The elements $\pi(\epsilon_i \epsilon_j) \in \mathrm{GL}(V)$ are all commuting involutions, and so we can simultaneously diagonalise them (since the characteristic of the ground field is not 2). This implies that

$$V = \bigoplus_{|I| \leqslant \frac{n}{2}} E_I$$

Note that for each $m \leqslant \frac{n}{2}$, the subgroup A_n acts on $\bigoplus_{|I|=m} E_I$; such a subspace is also preserved by the subgroup 2^{n-1}, and so by the whole of D_n'. Since there are no vectors fixed by each $\pi(\epsilon_i \epsilon_j)$, we have

$$V = \bigoplus_{|I| > 0} E_I$$

The action of A_n permutes the subspaces E_I according to the natural action of A_n on the subsets of N. Hence for any $k < \frac{n}{2}$ we have

$$\dim \bigoplus_{|I|=k} E_I = \binom{n}{k} \dim E_I$$

for any $I \subseteq N$ with $|I| = k$, and for $k = \frac{n}{2}$ (assuming that n is even) we have

$$\dim \bigoplus_{|I|=k} E_I = \frac{1}{2}\binom{n}{k} \dim E_I$$

Since $\dim V < 2n$, we conclude that

$$V = \bigoplus_{i \in N} E_{\{i\}}$$

and each $E_{\{i\}}$ is 1-dimensional. Let us pick a non-zero vector in $E_{\{i\}}$ for each i; these vectors form a basis of V.

It is immediate that with respect to this basis, the action of 2^{n-1} agrees with that of the restriction of the standard representation of D_n' to 2^{n-1}; moreover, it also shows that for each $\tau \in A_n$ the matrix $\pi(\tau)$ is a monomial matrix obtainable from the matrix given by the standard representation of D_n' by multiplication by a diagonal matrix.

Since $n \geqslant 6$, the setwise stabiliser in A_n of any $E_{\{i\}}$ is simple (as it is isomorphic to A_{n-1}), and so we can rescale each vector in our basis so that $\pi(\tau)$ becomes a permutation matrix for each $\tau \in A_n$, and this concludes our proof. ☐

Recall that $R_p(A_n)$ (occurring in the statement of the following result) denotes the minimal dimension of a faithful projective representation of A_n as in [**KL**]

Proposition 7.10. *Let $n \geqslant 8$ be even. Let $\pi \colon D_n' \to \mathrm{PGL}(V)$ be a faithful projective representation of dimension less than $n + R_p(A_n)$ over an algebraically closed field $\overline{\mathbb{F}}$ of characteristic $p > 2$. Then the projective representation lifts to a representation $\overline{\pi} \colon D_n' \to \mathrm{GL}(V)$, and the module V splits as $W \oplus U$ where W is a sum of trivial modules, and U is the standard module of D_n'.*

Proof. Let $d \in \mathrm{GL}(V)$ be a lift of $\pi(\delta)$. Since δ is an involution, d^2 is central, and so the characteristic polynomial of d is $x^2 - \lambda$ for some $\lambda \in \mathbb{F}^\times$. Since the field $\mathbb{F}$ is algebraically closed and not of characteristic 2, this characteristic polynomial has two distinct roots, and so d is diagonalisable. Upon multiplying d by a central matrix we may assume that at least one of the entries in the diagonal matrix of d is 1. Thus all the entries are ± 1, and in particular d is also an involution.

For any $\xi \in D'_n$, let $\overline{\xi} \in \mathrm{GL}(V)$ denote a lift of $\pi(\xi)$. Since δ is central in D'_n, every $\overline{\xi}$ either preserves the eigenspaces of d, or permutes them. This way we obtain a homomorphism $D'_n \to \mathbb{Z}/2\mathbb{Z}$, which has to be trivial by Lemma 2.1. Thus every $\overline{\xi}$ preserves the eigenspaces of d.

Since $\pi(\delta)$ is not trivial (as π is faithful), the involution d has a non-trivial eigenspace for each eigenvalue, and the same is true for any other involution lifting $\pi(\delta)$.

Take an eigenspace of d of dimension less than n. By [**KL**, Corollary 5.5.4], the projective module obtained by restricting to this eigenspace is not faithful – in fact, the action of D'_n on W has the whole of 2^{n-1} in its kernel. Therefore, if both eigenspaces of d are of dimension less than n, then the kernel of π contains an index two subgroup of 2^{n-1}, and therefore is not trivial. This contradicts the assumption on faithfulness of π.

We conclude that one of the eigenspaces of d, say U, has dimension at least n. But then the other eigenspace W, has dimension less than $R_p(A_n)$, and so the restricted projective A_n-module W is trivial. Hence it is also a trivial projective D'_n-module. The abelianisation of D'_n is trivial, and so for each ξ we may choose $\overline{\xi}$ so that its restriction to W is the identity matrix. In this way we obtain a homomorphism $\overline{\pi}\colon D'_n \to \mathrm{GL}(V)$ by declaring $\overline{\pi}(\xi) = \overline{\xi}$. Note that W is a sum of trivial submodules of this representation, and so in particular it is the (+1)-eigenspace of $\overline{\delta}$.

It is easy to see that in fact $V = U \oplus W$ as a D'_n-module, since $\overline{\xi}$ preserves the eigenspaces of $\overline{\delta}$ for every $\xi \in D'_n$, as remarked above.

Suppose that there is a non-zero vector in U fixed by each $\overline{\epsilon_i\epsilon_j}$. Then it is also fixed by $\overline{\delta}$, as $\delta = \epsilon_1 \cdots \epsilon_n$ and n is even. But $\overline{\delta}$ acts as minus the identity on U, which is a contradiction. Hence we may apply Lemma 7.9 to U and finish the proof. □

Corollary 7.11. *Let $n \geqslant 8$. Let $\pi\colon D'_n \to \mathrm{PGL}(V)$ be a faithful projective representation over an algebraically closed field $\overline{\mathbb{F}}$ of characteristic $p > 2$ of dimension less than $2 \cdot R_p(A_n)$ when n is even or less than $n + R_p(A_{n-1}) - 1$ when n is odd. Then the representation lifts to a representation $\overline{\pi}\colon D'_n \to \mathrm{GL}(V)$, and the module V splits as $W \oplus U$ where W is a sum of trivial and U is the standard module of D'_n.*

Proof. When n is even the result is covered by Proposition 7.10; let us assume that n is odd.

We apply Proposition 7.10 to two subgroups P_1 and P_2 of D'_n isomorphic to D'_{n-1}, where P_i is the stabiliser of a_i in D'_n.

If $\dim V < n-1$ then we immediately learn that V is sum of trivial P_1 modules, and thus it is also a sum of trivial D'_n-modules, as D'_n is the closure of A_n which is simple and has a non-trivial intersection with P_1. Let as assume that

$$\dim V \geqslant n - 1 \geqslant 5$$

We obtain a lift of the projective representations of P_1 and P_2 into $\mathrm{GL}(V)$; it is immediate that the two lifts agree on each $\epsilon_i\epsilon_j$ with $i, j > 2$, since each of the lifts of such an element is an involution with (-1)-eigenspace of dimension 2 and (+1)-eigenspace of dimension $\dim V - 2 > 2$, lifting $\pi(\epsilon_i\epsilon_j)$. Similarly, the lifts of the elements $\sigma_{ij}\sigma_{kl}$ (with $i, j, k, l > 2$ all distinct) also agree. It follows that the lifts agree on $P_1 \cap P_2 \cong D'_{n-2}$.

We now repeat the argument for any two stabilisers P_i and P_j. This way we have defined a map from generators of D'_n to $\mathrm{GL}(V)$, which respects all relations supported by some P_i. But it is easy to see that such relations are sufficient for defining the group, and so the map induces a homomorphism $\overline{\pi}\colon D'_n \to \mathrm{GL}(V)$ as required.

Let U_i denote the standard P_i-module, and W_i its complement which is a sum of trivial P_i-modules. Let $U = \sum U_i$. We claim that U is D'_n-invariant: take a generator ξ of D'_n lying in, say, $P_1 \smallsetminus P_2$. Let $x \in U_2$. Then $x = y + z$ with $y \in U_1$ and $z \in W_1$, and so

$$\overline{\pi}(\xi)(x) = y' + z = x - (y - y') \in U_2 + U_1$$

Similar computations for arbitrary indices prove the claim. Now Lemma 7.9 implies that U is the standard representation of D'_n.

Consider V as a 2^{n-1}-representation. Since V is a vector space over a field of characteristic $p > 2$, this representation is semi-simple, and so U has a complement W. It is clear that W is is a sum of trivial 2^{n-1}-representations, since all the non-trivial modules of elements $\epsilon_i\epsilon_j$ are contained in some U_l, and thus in U. For the same reason it is clear that W is a sum of trivial A_n-modules – for this we look at elements $\sigma_{ij}\sigma_{kl}$. We conclude that W is a sum of trivial D'_n-modules. □

7.3. Projective representations of $\mathrm{SAut}(F_n)$**.** Now we use the rigidity of D'_n-representations developed above in the context of projective representations of $\mathrm{SAut}(F_n)$.

Theorem 7.12. *Let $n \geqslant 8$. Let $\pi\colon \mathrm{SAut}(F_n) \to \mathrm{PGL}(V)$ be a projective representation of dimension k with $k < 2n - 4$ over an algebraically closed field $\overline{\mathbb{F}}$ of characteristic other than 2. If π does not factor through the natural map $\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$, then $k \geqslant n+1$ and the projective module V contains a trivial projective module of dimension $k - n - 1$.*

Proof. Since $\overline{\mathbb{F}}$ is algebraically closed, $\mathrm{PGL}(V) = \mathrm{PSL}(V)$. As π does not factor through the natural map $\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$, Lemma 2.5 tells us that π restricted to D'_n is injective.

By Corollary 7.11 we see that there is a lifting $\overline{\pi}$ of the projective representation of D'_n to a linear representation on V such that $V = W \oplus U$ as a D'_n-module, where U is standard and W is a sum of trivial modules. Let $u_1, \dots, u_n$ denote the standard basis for U. For notational convenience we will write

$$\overline{\xi} = \overline{\pi}(\xi)$$

for $\xi \in D'_n$.

Note that Corollary 7.11 implies that

$$U = \bigoplus_{i \in N} E_{\{i\}}$$

where $E_{\{i\}}$ is spanned by u_i.

Let us pick a lift of $\pi(\rho_{12})$ acting linearly on V; we will call it $\overline{\rho_{12}}$. Since ρ_{12} commutes with $\epsilon_i\epsilon_j$ with $i, j > 2$, the element $\overline{\rho_{12}}$ permutes the eigenspaces of $\overline{\epsilon_i\epsilon_j}$. But for a given pair (i, j), the eigenspaces of $\overline{\epsilon_i\epsilon_j}$ have dimensions 2 and $\dim V - 2 \geqslant n - 2 > 2$. Thus $\overline{\rho_{12}}$ preserves each eigenspace of $\overline{\epsilon_i\epsilon_j}$. It follows that

$$\overline{\rho_{12}}(u_i) \in \langle u_i \rangle = E_{\{i\}}$$

for each $i > 2$.

Let us choose lifts $\overline{\rho_{ij}}$ of $\pi(\rho_{ij})$ for each pair (i, j). By a discussion identical to the one above we see that $\overline{\rho_{ij}}$ preserves $E_{\{l\}}$ for $l \notin \{i, j\}$.

We may choose $\overline{\rho_{12}}$ so that it fixes u_3. We have

$$[\overline{\rho_{14}}^{-1}, \overline{\rho_{42}}^{-1}] = \lambda \cdot \overline{\rho_{12}}^{-1}$$

for some $\lambda \in \overline{\mathbb{F}} \smallsetminus \{0\}$. But clearly $[\overline{\rho_{14}}^{-1}, \overline{\rho_{42}}^{-1}](u_3) = u_3$, since both $\overline{\rho_{14}}^{-1}$ and $\overline{\rho_{42}}^{-1}$ preserve u_3 up to homothety. Therefore $\lambda = 1$, and thus

$$\overline{\rho_{12}}^{-1}(u_i) = u_i$$

for all $i > 4$. Replacing 4 by another number greater than 3 in the calculation above yields the same result for any $i > 3$. Using analogous argument we may choose each $\overline{\rho_{ij}}$ so that it fixes u_l for all $l \not\in \{i, j\}$. It follows that conjugating $\overline{\rho_{12}}$ by an element $\overline{\xi}$ (with $\xi \in A_n$) yields an appropriate element $\overline{\rho_{ij}}$, and not just $\overline{\rho_{ij}}$ up to homothety.

We also see that $\overline{\rho_{12}}$ preserves $Z = W \oplus \langle u_1, u_2 \rangle$, as this is the centraliser of

$$\langle \{\overline{\epsilon_i \epsilon_j}\} \mid i, j > 2\} \rangle$$

Note that W is a subspace of Z of codimension 2; therefore $W' = \overline{\rho_{12}}^{-1}(W) \cap W$ is a subspace of W of codimension at most 2, and so of dimension at least $k - n - 2$. Let $x \in W'$ be any vector. Now $\overline{\rho_{12}}(x)$ lies in W, and so

$$\overline{\rho_{12}}(x) = \overline{\sigma_{12}\sigma_{13}\epsilon_3\epsilon_2}\,\overline{\rho_{12}}(x)$$

Thus

$$\overline{\rho_{31}}\,\overline{\rho_{12}}(x) = \overline{\rho_{31}}\,\overline{\sigma_{12}\sigma_{13}\epsilon_3\epsilon_2}\,\overline{\rho_{12}}(x) = \overline{\sigma_{12}\sigma_{13}\epsilon_3\epsilon_2}\,\overline{\rho_{12}}^{-1}\overline{\rho_{12}}(x) = \overline{\sigma_{12}\sigma_{13}\epsilon_3\epsilon_2}(x) = x$$

where the last equality follows from the fact that $x \in W$. Observe that

$$\overline{\rho_{12}}.x = \overline{\epsilon_2\epsilon_3}\,\overline{\rho_{12}}.x = \overline{\rho_{12}}^{-1}\,\overline{\epsilon_2\epsilon_3}.x = \overline{\rho_{12}}^{-1}.x$$

Using a similar argument we show that

$$\overline{\rho_{31}}^{-1}\overline{\rho_{12}}^{-1}(x) = x$$

and so

$$[\overline{\rho_{31}}^{-1}, \overline{\rho_{12}}^{-1}](x) = x$$

But $[\overline{\rho_{31}}^{-1}, \overline{\rho_{12}}^{-1}] = \overline{\rho_{32}}^{-1}$, and so $\overline{\rho_{32}}(x) = x$. Conjugating by elements $\overline{\xi}$ with $\xi \in D_n'$ we conclude that

$$\overline{\rho_{ij}}(x) = x = \overline{\lambda_{ij}}(x)$$

for every i and j. This implies that W' is preserved by $\mathrm{SAut}(F_n)$, and the restricted projective module is trivial.

If $\dim W' = k - n - 1$ then we are done. Let us assume that this is not the case, that is that W' is of codimension 2 in W. Consider the involution $\rho_{12}\epsilon_2\epsilon_3$. We set

$$\overline{\rho_{12}\epsilon_2\epsilon_3} = \overline{\rho_{12}}\,\overline{\epsilon_2\epsilon_3}$$

Note that this element satisfies

$$\overline{\rho_{12}\epsilon_2\epsilon_3}^2 = \nu I$$

for some $\nu \in \mathbb{F}^\times$. But $\overline{\rho_{12}\epsilon_2\epsilon_3}(u_3) = -u_3$ and so $\nu = 1$. Therefore $\overline{\rho_{12}\epsilon_2\epsilon_3}$ is an involution.

Let $Y = Z/W'$. Note that $\overline{\epsilon_2\epsilon_3}$ acts on Y, and its (-1)-eigenspace of dimension exactly 1, and the $(+1)$-eigenspace of dimension 3. We also have an action of the involution $\overline{\rho_{12}\epsilon_2\epsilon_3}$ on Y.

Since $\overline{\rho_{12}\epsilon_2\epsilon_3}$ acts trivially on W' and its (-1)-eigenspace in the complement of Z in V is of dimension 1, the dimension of its (-1)-eigenspace in Y must be odd (here we use the fact that π is a map to $\mathrm{PSL}(\mathbb{F})$). Thus there are at least two linearly independent vectors v_1 and v_2 lying in the intersection of the $(+1)$-eigenspace of $\overline{\epsilon_2\epsilon_3}$ and some eigenspace of $\overline{\rho_{12}\epsilon_2\epsilon_3}$.

Since $\overline{\rho_{12}\epsilon_2\epsilon_3}$ is an involution and the characteristic of $\mathbb{F}$ is odd, there exists a complement of W' in Z on which $\overline{\rho_{12}\epsilon_2\epsilon_3}$ acts as on Y. Thus, we have the two vectors corresponding to v_1 and v_2; we will abuse the notation by calling them v_1 and v_2 as well.

Since W' lies in the $(+1)$-eigenspace of $\overline{\epsilon_2\epsilon_3}$, so do v_1 and v_2. Thus there is a non-zero linear combination v_3 of v_1 and v_2 which lies in W, since the codimension of W in the $(+1)$-eigenspace of $\overline{\epsilon_2\epsilon_3}$ is 1. Also, $\overline{\epsilon_2\epsilon_3}$ act trivially on this vector, and so we have found another vector in W which is mapped to W by $\overline{\rho_{12}}$. Arguing exactly as before we show that $\langle v_3\rangle$ is $\mathrm{SAut}(F_n)$ invariant and trivial as a projective module. Hence $W' \oplus \langle v_3\rangle$ is also $\mathrm{SAut}(F_n)$ invariant, and is trivial as a projective module since $\mathrm{SAut}(F_n)$ is perfect (Proposition 2.4). □

Let us remark here that there do exist representations of $\mathrm{SAut}(F_n)$ in dimension $n+1$ over any field (over $\mathbb{Z}$ in fact) which do not factor through the natural map $\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$ – see [**BV3**, Proposition 3.2].

Theorem 7.13. *Let $n \geqslant 10$. Then every finite simple classical group of Lie type in odd characteristic which is a quotient of* $\mathrm{SAut}(F_n)$ *is larger in order than* $\mathrm{L}_n(2)$.

Proof. Let K be such a quotient, and suppose that $|K| \leqslant |\mathrm{L}_n(2)|$. Let k denote the rank of K.

If K is of type $\mathtt{A}$ or ${}^2\mathtt{A}$, then Lemma 10.5 tells us that $k \leqslant 2n-8$. If K is of any other classical type, then Lemma 10.6 tells us that $k \leqslant n-4$.

Let V be the natural projective module of K, and let m denote its dimension. Note that V is an irreducible projective K-module, and $m < 2n-6$. Thus Theorem 7.12 implies that either the representation

$$\mathrm{SAut}(F_n) \to K \to \mathrm{PGL}(V)$$

factors through the natural map $\mathrm{SAut}(F_n) \to \mathrm{SL}_n(\mathbb{Z})$, or $m = n+1$. In the former case we must have $K \cong \mathrm{L}_n(p)$ for some prime p, as K is simple. But $\mathrm{L}_n(p)$ is larger than $\mathrm{L}_n(2)$ for $p \geqslant 3$.

We may thus assume that $m = n+1$. If K is of type $\mathtt{A}$ or ${}^2\mathtt{A}$ then it is immediate that it is too big.

When n is even this means that K is the simple group $\mathtt{B}_{\frac{n}{2}}(q)$. This is larger (in cardinality) than $\mathrm{L}_n(2)$ for every $q > 3$ by Lemma 10.7, and so we may assume that $q = 3$. But this is impossible by Lemma 7.5.

When n is odd, K is one of the simple groups $\mathtt{C}_{\frac{n+1}{2}}(q)$, $\mathtt{D}_{\frac{n+1}{2}}(q)$ or ${}^2\mathtt{D}_{\frac{n+1}{2}}(q)$. Lemma 10.7 immediately rules out all values of q except for $q = 3$, and again we are done by applying Lemma 7.5. □

8. The exceptional groups of Lie type

In this section we focus on exceptional groups of Lie type. These are

(1) the Suzuki-Ree groups ${}^2\mathtt{B}_2(2^{2m+1})$, ${}^2\mathtt{G}_2(3^{2m+1})$, ${}^2\mathtt{F}_4(2^{2m+1})$ and ${}^2F_4(2)'$,
(2) the Steinberg groups ${}^3\mathtt{D}_4(q)$, ${}^2\mathtt{E}_6(q)$ and
(3) the exceptional Chevalley groups $\mathtt{G}_2(q)$, $\mathtt{F}_4(q)$, $\mathtt{E}_6(q)$, $\mathtt{E}_7(q)$ and $\mathtt{E}_8(q)$.

They are defined for all $q \geqslant 2$, $m \geqslant 0$ and are all simple with the following exceptions: the group $\mathrm{Sz}(2) \cong 5 \rtimes 4$ which is visibly solvable; the group ${}^2\mathtt{G}_2(3)$ whose index 3 derived subgroup is isomorphic to $\mathtt{A}_1(8)$; the group $\mathtt{G}_2(2)$ whose index 2 derived subgroup is isomorphic to ${}^2\mathtt{A}_2(3)$; and the group ${}^2\mathtt{F}_4(2)$ whose derived subgroup ${}^2\mathtt{F}_4(2)'$ is simple.

For simplicity, in this section we always use the type symbols to denote the adjoint versions.

We now introduce the following notation: for a group K the *A-rank* of K is the largest n such that K contains a copy of the alternating group A_n. In particular, we will make use of the bounds on the A-rank of the exceptional groups given in [**LS**, Table 10.1].

Generally, to show that a group K is not the smallest quotient of some $\mathrm{SAut}(F_n)$ we argue as follows: let $n(K)$ be the smallest integer such that

$$|\,\mathrm{L}_{n(K)-1}(2)| < |K| \leqslant |\,\mathrm{L}_{n(K)}(2)|$$

and assume that we have an epimorphism $\varphi\colon\ \mathrm{SAut}(F_n) \to K$ with $n \geqslant n(K)$. Now we compare n to the 2-rank, the D'-rank, and A-rank of K. If the 2-rank is smaller than $n-1$ the we use Lemma 2.5 applied to the subgroup 2^{n-1} and conclude that K is in fact a quotient of $\mathrm{SL}_n(\mathbb{Z})$. But the smallest such quotient is $\mathrm{L}_n(2)$. If the A-rank of K is smaller than $n+1$, then we use Lemma 2.5 applied to the subgroup A_n (we observe that if A_{n+1} is not mapped injectively, then neither is A_n for any $n \geqslant 3$). Similarly for the D'-rank.

If the 2-rank and A-rank arguments fail, we look at centralisers. If $n \geqslant 5$ and the simple non-abelian factors of every involution in K have already been shown not to be quotients of $\mathrm{SAut}(F_{n-2})$, then we look at $\varphi(\epsilon_1\epsilon_2)$. If this is trivial then we are done by Lemma 2.5; otherwise we obtain a map from $\mathrm{SAut}(F_{n-2})$ (which centralises $\epsilon_1\epsilon_2$) to a group whose simple composition factors are not quotients of $\mathrm{SAut}(F_{n-2})$ (note that the abelian factors are ruled out by the fact that $\mathrm{SAut}(F_{n-2})$ is perfect). Thus $\mathrm{SAut}(F_{n-2})$ lies in the kernel of φ, and so in particular φ trivialises some transvection. But then it trivialises every transvection since they are all conjugate, and thus φ is trivial.

If $n \geqslant 5$ we may argue analogously using the element γ of order 3 from Lemma 7.1; if $n \geqslant 2+k$ for $k \geqslant 5$ odd we may argue in an analogous manner using Lemma 2.8.

Lemma 8.1. *Let K be a finite simple group belonging to one of the following families:*

(1) the Suzuki groups ${}^2\mathtt{B}_2(2^{2m+1})$,
(2) the small Ree groups ${}^2\mathtt{G}_2(3^{2m+1})$,
(3) the large Ree groups ${}^2\mathtt{F}_4(2^{2m+1})$, or,
(4) the Tits group ${}^2\mathtt{F}_4(2)'$,

where $m \geqslant 1$ is an integer. Then K is not the smallest finite non-trivial quotient of $\mathrm{SAut}(F_n)$.

Proof. The smallest K among the families considered is isomorphic to $\mathrm{Sz}(8) = {}^2\mathtt{B}_2(8)$, and has order greater than $|\,\mathrm{L}_4(2)|$; thus $n \geqslant 5$. The order of ${}^2\mathtt{B}_2(2^{2m+1})$ is coprime to 3, and the order of ${}^2\mathtt{G}_2(3^{2m+1})$ is coprime to 5. Hence it is clear that the simple Suzuki and small Ree groups cannot be quotients of $\mathrm{SAut}(F_n)$ for $n \geqslant 4$, since the alternating group A_5 cannot be mapped injectively, and so we may use Lemma 2.5.

Now assume that K is ${}^2\mathtt{F}_4(2^{2m+1})$ or the Tits group; observe that the smallest member of this family, the Tits group ${}^2\mathtt{F}_4(2)'$, has order greater than $\mathrm{L}_5(2)$ and so $n \geqslant 6$. But, by [**Mal2**, Proposition 2.2] we see that the A-rank of K is 6. □

Lemma 8.2. *Let K be a finite simple group belonging to one of the following families:*

(1) the exceptional groups of type $\mathtt{G}_2(q)$, where $q \geqslant 3$, or
(2) the exceptional groups of type ${}^3\mathtt{D}_4(q)$, where $q \geqslant 2$.

Then, K is not the smallest finite non-trivial quotient of $\mathrm{SAut}(F_n)$.

Proof. First, let $K \cong \mathtt{G}_2(q)$. We divide the proof into the case that q is either odd or even. When q is odd the 2-rank of K is 3 by [**Kle1**, Lemma 2.4], but

$$|K| \geqslant |G_2(3)| > |\,\mathrm{L}_4(2)|$$

and so $n \geqslant 5$.

When $q \geqslant 4$ is even, $|K| > |\,\mathrm{L}_5(2)|$ but from inspection of the list of maximal subgroups of K (see [**Coo**]) we see that the A-rank of K is at most 5.

For $K \cong {}^3\mathtt{D}_4(q)$ note that the smallest member of this family is ${}^3\mathtt{D}_4(2)$ and has order greater than $|\,\mathrm{L}_5(2)|$. The maximal subgroups of K are known (see [**Kle2**]) and we see that the A-rank of K is 5. □

For the remaining groups we again split into the odd and even characteristic case.

Lemma 8.3. *Let K be a finite simple exceptional group of type $\mathtt{F}_4$, $\mathtt{E}_6$, ${}^2\mathtt{E}_6$, $\mathtt{E}_7$ or $\mathtt{E}_8$ in odd characteristic. Then K is not the smallest non-trivial finite quotient of $\mathrm{SAut}(F_n)$.*

Proof. It is easy to see that if K belongs to any of these families, then the order of $|K|$ is bounded below when $q = 3$. If $K \cong \mathtt{F}_4(q)$, $\mathtt{E}_6(q)$ or ${}^2\mathtt{E}_6(q)$, then the A-rank of K is at most 7 [**LS**, Table 10.1] but the order of K is bounded below by the order of $\mathtt{F}_4(3)$ which has order greater than $\mathrm{L}_9(2)$.

If $K \cong \mathtt{E}_7(q)$, then then the A-rank of K is at most 10, but the smallest member of this family $\mathtt{E}_7(3)$ has order greater than $|\,\mathrm{L}_{14}(2)|$.

Finally, if $K \cong \mathtt{E}_8(q)$, then the A-rank of K is at most 11, but the smallest member of this family $\mathtt{E}_8(3)$ has order greater than $|\,\mathrm{L}_{19}(2)|$. □

Lemma 8.4. *Let K be a finite simple exceptional group of type $\mathtt{F}_4$, $\mathtt{E}_6$, ${}^2\mathtt{E}_6$, $\mathtt{E}_7$ or $\mathtt{E}_8$ defined over a finite field of order $q = 2^m \geqslant 4$. Then K is not the smallest non-trivial finite quotient of $\mathrm{SAut}(F_n)$.*

Proof. It is easy to see that if K belongs to any of these families, then the order of $|K|$ is bounded below when $q = 4$. The degree of the largest alternating group in each of these groups can be found in [**LS**, Table 10.1]. If $K \cong \mathtt{F}_4(q)$, then the A-rank of K is 10, but K has order greater than $\mathrm{L}_{10}(2)$. If $K \cong \mathtt{E}_6(q)$ or ${}^2\mathtt{E}_6(q)$, then the A-rank is bounded above by 12, but the smallest such group $\mathtt{E}_6(4)$ has order greater than $\mathrm{L}_{12}(2)$. Finally, if $K \cong \mathtt{E}_7(q)$ or $\mathtt{E}_8(q)$, then the A-rank of K is at most 17, but the smallest member of this family $\mathtt{E}_7(4)$ has order greater than $\mathrm{L}_{16}(2)$. □

In order to dispose of the remaining five cases, we state the following result whose proof can be found in [**AS**, Sections 15-17].

Lemma 8.5.
1. *Any non-abelian composition factor of an involution centraliser in $\mathtt{E}_6(2)$ is isomorphic to one of $\mathtt{A}_2(2)$, $\mathtt{A}_5(2)$ or $\mathtt{B}_3(2)$.*
2. *Any non-abelian composition factor of an involution centraliser in $\mathtt{E}_7(2)$ is isomorphic to one of $\mathtt{B}_3(2)$, $\mathtt{B}_4(2)$, $\mathtt{D}_6(2)$ or $\mathtt{F}_4(2)$.*
3. *Any non-abelian composition factor of an involution centraliser in $\mathtt{E}_8(2)$ is isomorphic to one of $\mathtt{B}_4(2)$, $\mathtt{B}_6(2)$, $\mathtt{F}_4(2)$ or $\mathtt{E}_7(2)$.*

We are now in a position to prove the following.

Lemma 8.6. *Let K be a finite simple exceptional group of type $\mathtt{F}_4(2)$, $\mathtt{E}_6(2)$, ${}^2\mathtt{E}_6(2)$, $\mathtt{E}_7(2)$ or $\mathtt{E}_8(2)$. Then K is not the smallest non-trivial finite quotient of $\mathrm{SAut}(F_n)$.*

Proof. If $K \cong \mathtt{F}_4(2)$, then $n \geqslant 8$, but from the comparison of the character tables of K [**CCN**$^+$] and of D'_8 which can be performed in GAP, we see that D'_8 is not a subgroup of K, and we use Lemma 2.5. We eliminate the case $K \cong {}^2\mathtt{E}_6(2)$ in the same way, except that here $n \geqslant 9$.

If $K \cong \mathtt{E}_6(2)$, then $|K| > |\,\mathrm{L}_8(2)|$. By the preceding lemma, it remains to show that any homomorphism from $\mathrm{SAut}(F_n)$ with $n \geqslant 7$ to $\mathtt{A}_2(2)$, $\mathtt{A}_5(2)$ or $\mathtt{B}_3(2)$ is trivial. It can easily be checked in GAP that none of these groups contains a subgroup isomorphic to D_7', hence the result follows from Lemma 2.5. (Also, we will revisit $\mathrm{SAut}(F_7)$ in the next section.)

If $K \cong \mathtt{E}_7(2)$, then $|K| > |\mathrm{L}_{11}(2)|$. By the preceding lemma, it remains to show that any homomorphism from $\mathrm{SAut}(F_n)$ with $n \geqslant 10$ to $\mathtt{B}_3(2)$, $\mathtt{B}_4(2)$, $\mathtt{D}_6(2)$ or $\mathtt{F}_4(2)$ is trivial. The groups $\mathtt{B}_3(2)$, $\mathtt{B}_4(2)$ and $\mathtt{D}_6(2)$ are of classical type in even characteristic and smaller in cardinality than $\mathrm{L}_{10}(2)$, hence we can apply Theorem 6.10. The maximal subgroups of $\mathtt{F}_4(2)$ are known and can be found in [**CCN**$^+$]; it is clear by inspection that the A-rank of $\mathtt{F}_4(2)$ is 10.

Finally, if $K \cong \mathtt{E}_8(2)$, then $|K| > |\mathrm{L}_{15}(2)|$. By the preceding lemma, it remains to show that any homomorphism from $\mathrm{SAut}(F_n)$ with $n \geqslant 14$ to $\mathtt{B}_4(2)$, $\mathtt{B}_6(2)$, $\mathtt{F}_4(2)$ or $\mathtt{E}_7(2)$ is trivial. As before, Theorem 6.10 takes care of $\mathtt{B}_4(2)$ and $\mathtt{B}_6(2)$ since they are smaller in cardinality than $\mathrm{L}_{14}(2)$, whereas the A-rank of $\mathtt{F}_4(2)$ and $\mathtt{E}_7(2)$ is at most 13. This completes the proof. □

We can now summarise the preceding lemmata.

Theorem 8.7. *Let K be a finite simple group of exceptional type. If K is a quotient of $\mathrm{SAut}(F_n)$, then $|K| > |\mathrm{L}_n(2)|$.*

In fact, using the A-rank we can say more: when $n > 16$ then the exceptional groups of Lie type are never quotients of $\mathrm{SAut}(F_n)$, see [**LS**].

9. Small values of n and the conclusion

We can now conclude the paper.

Theorem 9.1. *Let $n \geqslant 3$. Every non-trivial finite quotient of $\mathrm{SAut}(F_n)$ is either greater in cardinality than $\mathrm{L}_n(2)$, or isomorphic to $\mathrm{L}_n(2)$. Moreover, if the quotient is $\mathrm{L}_n(2)$, then the quotient map is the natural map postcomposed with an automorphism of $\mathrm{L}_n(2)$.*

Proof. Suppose that $n \geqslant 8$, and let K be a smallest non-abelian quotient of $\mathrm{SAut}(F_n)$. Since $\mathrm{SAut}(F_n)$ is perfect, K is simple. By Corollary 3.18, K is not an alternating group; by Proposition 4.3, K is not a sporadic group; by Theorem 7.13, K is not a classical group of Lie type in odd characteristic; by Theorem 8.7, K is not an exceptional group of Lie type. Finally, by Theorem 6.10, K is isomorphic to $\mathrm{L}_n(2)$, and the quotient map is obtained by postcomposing the natural map $\mathrm{SAut}(F_n) \to \mathrm{L}_n(2)$ by an automorphism of $\mathrm{L}_n(2)$.

For $3 \leqslant n < 8$, the result follows from Lemmata 9.3 to 9.7 below. □

As indicated above, we now verify Theorem 9.1 for $n \in \{3, \ldots, 7\}$. Note that in view of Proposition 6.6, it is enough to show that a smallest quotient of $\mathrm{SAut}(F_n)$ is isomorphic to $\mathrm{L}_n(2)$.

By Corollary 3.18, Proposition 4.3, and Theorem 8.7 we can assume that K is of classical type. We make use of the list of simple groups in order of size appearing in [**CCN**$^+$, pgs. 239–242]. Note that this list does not contain all members of the families of types $\mathtt{A}_1(q)$, $\mathtt{A}_2(q)$, ${}^2\mathtt{A}_2(q)$, $\mathtt{A}_3(q)$, $\mathtt{C}_2(q)$ or $\mathtt{G}_2(q)$; we can exclude $\mathtt{G}_2(q)$ by Lemma 8.2.

Lemma 9.2 ([**GLS**, Theorem 4.10.5])**.** *Let $K \leqslant \mathrm{L}_n(q)$ where q is odd. If $n \leqslant 4$, then the 2-rank of K is bounded above by n.*

The general strategy is exactly as described in the previous section. As before, we use types to denote the adjoint versions.

We now look at each value of n separately.

Lemma 9.3 ($n = 3$)**.** *Let K be a non-abelian finite simple group with $|K| \leqslant |\mathrm{L}_3(2)|$. If K is a quotient of $\mathrm{SAut}(F_3)$, then $K \cong \mathrm{L}_3(2)$.*

Proof. If K is a non-abelian simple group not isomorphic to $\mathrm{L}_3(2)$ and order at most $|\mathrm{L}_3(2)|$, then $K \cong A_5$. But A_5 is not a quotient of $\mathrm{SAut}(F_3)$ by Lemma 3.1. □

Lemma 9.4 ($n = 4$)**.** *Let K be a non-abelian finite simple group with $|K| \leqslant |\mathrm{L}_4(2)|$. If K is a quotient of* $\mathrm{SAut}(F_4)$*, then $K \cong \mathrm{L}_4(2)$.*

Proof. Let K be a simple group of order at most $|\mathrm{L}_4(2)|$ and a quotient of $\mathrm{SAut}(F_4)$. Assume that K is not isomorphic to $\mathrm{L}_4(2)$. By Lemma 9.2, K is not a subgroup of $\mathtt{A}_2(q)$ where q is odd. Hence, K is isomorphic to one of the following.

$$\mathtt{A}_1(8),\ \mathtt{A}_1(16),\ \mathtt{A}_2(4)$$

With the exception of $\mathrm{L}_3(4)$ (which has the same order as $\mathrm{L}_4(2)$), it is clear from the inspection of their maximal subgroups [**CCN**$^+$] that they do not contain subgroups isomorphic to D_4'. In the case of $\mathrm{L}_3(4)$, there is a subgroup isomorphic to $2^4 \rtimes A_5$, however this is not isomorphic to the group D_5'. It can be computed in GAP that $\mathrm{L}_3(4)$ does not contain subgroups isomorphic to D_4'. This completes the proof. □

Lemma 9.5 ($n = 5$)**.** *Let K be a non-abelian finite simple group with $|K| \leqslant |\mathrm{L}_5(2)|$. If K is a quotient of* $\mathrm{SAut}(F_5)$*, then $K \cong \mathrm{L}_5(2)$.*

Proof. Assume that K is not isomorphic to $\mathrm{L}_5(2)$. By Lemmata 6.1 and 9.2 we can exclude all but the following groups.

$$\mathtt{A}_2(4),\ {}^2\mathtt{A}_2(4),\ {}^2\mathtt{A}_2(8),\ \mathtt{A}_3(3),\ {}^2\mathtt{A}_3(2) \cong \mathtt{C}_2(3),\ \mathtt{C}_2(4),\ \mathtt{C}_2(5),\ {}^2\mathtt{A}_3(3),\ \mathtt{C}_3(2)$$

The groups $\mathtt{A}_3(3)$ and ${}^2\mathtt{A}_3(3)$ are dealt with in Lemma 7.4. Excluding those which also do not contain D_5' as subgroups we are left with the possibilities $\mathtt{C}_3(2)$ and $\mathtt{C}_2(5)$. If $K \cong \mathtt{C}_2(5)$ or $\mathtt{C}_3(2)$, then any non-abelian composition factor of an involution centraliser is isomorphic to A_5 or A_6, neither of which is a quotient of $\mathrm{SAut}(F_3)$ by Lemma 3.1, a contradiction. □

Lemma 9.6 ($n = 6$)**.** *Let K be a non-abelian finite simple group with $|K| \leqslant |\mathrm{L}_6(2)|$. If K is a quotient of* $\mathrm{SAut}(F_6)$*, then $K \cong \mathrm{L}_6(2)$.*

Proof. Assume that K is not isomorphic to $\mathrm{L}_6(2)$. By Lemmata 6.1 and 9.2 we can assume that K has dimension at least 4 in even characteristic, in order to contain a subgroup isomorphic to A_7, and dimension at least 5 in odd characteristic in order for the 2-rank to be at least 5. Hence K is isomorphic to one of the following:

$$\mathtt{C}_2(8),\ \mathtt{A}_3(4),\ {}^2\mathtt{A}_3(4),\ \mathtt{C}_3(3),\ \mathtt{C}_3(3),\ {}^2\mathtt{A}_4(2),\ \mathtt{D}_4(2),\ {}^2\mathtt{D}_4(2),\ {}^2\mathtt{A}_5(2)$$

Those groups which contain subgroups isomorphic to D_6' are isomorphic to ${}^2\mathtt{A}_5(2)$, $\mathtt{B}_3(3)$, $\mathtt{D}_4(2)$ and ${}^2\mathtt{D}_4(2)$. The simple factors of the centralisers of elements of order 3 in ${}^2\mathtt{A}_5(2)$ and $\mathtt{B}_3(3)$ can be computed in GAP and are isomorphic to $\mathtt{A}_1(9)$ or $\mathtt{C}_2(3)$, neither of which is a quotient of $\mathrm{SAut}(F_4)$ – this follows from Lemma 7.3 for $\mathtt{C}_2(3)$ and from Lemma 9.4 for $\mathtt{A}_1(9)$, since they are smaller in cardinality than $\mathrm{L}_4(2)$.

The simple factors of the involution centralisers of $\mathtt{D}_4(2)$ and ${}^2\mathtt{D}_4(2)$ are isomorphic to $\mathtt{A}_1(4)$ or $\mathtt{A}_1(9)$, neither of which is a quotient of $\mathrm{SAut}(F_4)$, by Lemma 9.4, since they are both smaller in cardinality than $\mathrm{L}_4(2)$. □

Lemma 9.7 ($n = 7$)**.** *Let K be a non-abelian finite simple group with $|K| \leqslant |\mathrm{L}_7(2)|$. If K is a quotient of* $\mathrm{SAut}(F_7)$*, then $K \cong \mathrm{L}_7(2)$.*

Proof. Assume that K is not isomorphic to $\mathrm{L}_7(2)$. Again we make use of the values listed in Lemma 6.1 for $R_p(A_8)$, hence we need to consider groups of dimension at least 7 in odd characteristic and dimension at least 4 in even characteristic. In even characteristic we are left with the following

$$\mathtt{A}_3(8),\ {}^2\mathtt{A}_3(8),\ {}^2\mathtt{A}_4(4),\ \mathtt{B}_2(16),\ \mathtt{C}_3(4),\ \mathtt{C}_4(2),\ \mathtt{D}_5(2),\ {}^2\mathtt{D}_5(2)$$

none of which is a quotient of $\mathrm{SAut}(F_7)$ by Theorem 6.9, with the exception of ${}^2\mathtt{D}_5(2)$.

Now let $K \cong \mathtt{D}_5(2)$. The non-abelian simple quotients of the involution centralisers in K are isomorphic to A_6, A_8 or $\mathtt{C}_3(2)$. Since all of these are smaller than $\mathrm{L}_5(2)$ we apply Lemma 9.5 which completes the proof.

In odd characteristic we have the groups

$$\mathtt{D}_4(3),\ {}^2\mathtt{D}_4(3)$$

Any non-abelian simple factor of a centraliser of an element of order 3 in either of these groups is isomorphic to $\mathtt{A}_1(9)$ or to $\mathtt{C}_2(3)$. By Lemma 9.5, neither of these groups is a quotient of $\mathrm{SAut}(F_5)$ since they are smaller in cardinality than $\mathrm{L}_5(2)$. □

10. Computations

This appendix contains all the necessary computations. Note that we use type symbols to denote the adjoint versions of the groups of Lie type.

Lemma 10.1. *For $n \geqslant 8$ we have*

$$2^{n-3} > \binom{n}{2}$$

Proof. It is enough to observe that the result is true for $n = 8$, and

$$\frac{\binom{n+1}{2}}{\binom{n}{2}} = \frac{n+1}{n-1} \leqslant 2$$

for all $n \geqslant 3$. □

Lemma 10.2. *For an even $n \geqslant 12$ we have*

$$\frac{1}{2}\binom{n}{\frac{n}{2}} \geqslant \min\left\{\binom{n}{\lfloor\frac{n}{4}\rfloor}, 2^{n-\lfloor\frac{n}{4}\rfloor-1}\right\}$$

Proof. Let $n = 2m$. We have

$$\begin{aligned}
\frac{\frac{1}{2}\binom{n}{\frac{n}{2}}}{\binom{n}{\lfloor\frac{n}{4}\rfloor}} &= \frac{(\lfloor\frac{m}{2}\rfloor)!(2m-\lfloor\frac{m}{2}\rfloor)!}{2\cdot m!m!} \\
&= \frac{1}{2}\cdot\prod_{i=1}^{\lceil\frac{m}{2}\rceil}\frac{m+i}{\lfloor\frac{m}{2}\rfloor+i} \\
&= \frac{(m+1)(m+2)}{2(\lfloor\frac{m}{2}\rfloor+1)(\lfloor\frac{m}{2}\rfloor+2)}\cdot\prod_{i=3}^{\lceil\frac{m}{2}\rceil}\frac{m+i}{\lfloor\frac{m}{2}\rfloor+i} \\
&\geqslant \frac{(m+1)(m+2)}{(m+2)(\frac{m}{2}+2)} \\
&\geqslant \frac{2m+2}{m+4} \\
&\geqslant 1
\end{aligned}$$

for any $m \geqslant 2$.

We also have

$$2^{n-\lfloor\frac{n}{4}\rfloor-1} \leqslant 2^{n-\frac{n}{4}-\frac{1}{2}} = 2^{\frac{3m-1}{2}}$$

and

$$\frac{2^{\frac{3(m+1)-1}{2}}}{2^{\frac{3m-1}{2}}} = 2^{\frac{3}{2}} < 3$$

Now

$$\frac{\frac{1}{2}\binom{n+2}{\frac{n+2}{2}}}{\frac{1}{2}\binom{n}{\frac{n}{2}}} = \frac{(2m+1)(2m+2)}{(m+1)^2} \geqslant 3$$

We conclude by remarking that $\frac{1}{2}\binom{n}{\frac{n}{2}} \geqslant 2^{\frac{3n-2}{4}}$ for $n = 12$. □

Lemma 10.3. *For $n \geqslant 7$ we have $\binom{n}{2}! \cdot \frac{1}{2} > |\operatorname{L}_n(2)|$.*

Proof. We have

$$2^{n^2} > |\operatorname{L}_n(2)|$$

since the left-hand side is the number of $n \times n$ matrices over the field of 2 elements. We also have

$$m! \geqslant (2\pi m)^{\frac{1}{2}} (\frac{m}{e})^m$$

by Stirling's approximation. Putting $m = \binom{n}{2} = \frac{n(n-1)}{2}$ we obtain

$$\begin{aligned}\binom{n}{2}! \cdot \frac{1}{2} &\geqslant \frac{1}{2}(\pi n(n-1))^{\frac{1}{2}} (\frac{n(n-1)}{2e})^{\frac{n(n-1)}{2}} \\ &= \frac{\pi^{\frac{1}{2}}}{2} \cdot \frac{(n(n-1))^{\frac{n^2-n+1}{2}}}{(2e)^{\frac{n(n-1)}{2}}} \\ &\geqslant \frac{1}{2} \cdot \frac{2^{\frac{5(n^2-n+1)}{2}}}{2^{\frac{5n(n-1)}{4}}} \\ &= 2^{\frac{10(n^2-n+1)-5n(n-1)-4}{4}} \\ &= 2^{\frac{5n^2-5n+6}{4}}\end{aligned}$$

where we have used the fact that $2^{\frac{5}{2}} > 2e$ and that $n(n-1) \geqslant 2^5$, as $n \geqslant 7$.

Now

$$5n^2 - 5n + 6 > 4n^2$$

holds for every $n \geqslant 4$ and we are done. □

We will now proceed to compute certain inequalities between orders of adjoint versions of finite groups of Lie type – these orders be found in [**CCN**$^+$, pg. xvi]. Let us start by some general remarks.

Firstly, if we fix the type, rank and characteristic, then enlarging the field always results in enlarging the group: this is obvious for the universal versions, and for adjoint versions requires comparing the sizes of centres of the universal versions; such a comparison can easily be performed. Since we will be looking at the smallest groups of a given type, rank and characteristic, we may therefore assume that the field is of prime cardinality.

In fact, arguing as above, we see that for odd characteristics we may assume that the field is of size 3, and for odd characteristics greater than 3 we may assume the field to be of size 5.

Secondly, if we fix the type and field, then increasing the rank always results in enlarging the group. The argument is precisely as above. The same holds for twisted rank, since to increase the twisted rank we have to increase the rank.

Lemma 10.4. *Let $n \geqslant 8$. Then every finite group of Lie type in characteristic 2 of twisted rank at least $n-2$ is larger than $\operatorname{L}_n(2)$, with the exception of $\mathtt{A}_{n-2}(2)$ and $\mathtt{A}_{n-1}(2)$.*

Proof. By the discussion above, it is enough to prove the result for

$$\mathtt{A}_{2n-2}(4), {}^2\mathtt{A}_{2n-3}(2), \mathtt{B}_{n-2}(2), \mathtt{C}_{n-2}(2), \mathtt{D}_{n-2}(2), {}^2\mathtt{D}_{n-1}(2)$$

and $\mathtt{E}_6(2), \mathtt{E}_7(2)$ and $\mathtt{E}_8(2)$ for small values of n.

For the groups of type $\mathtt{E}$ we confirm the result by a direct computation.

The orders of $\mathtt{B}_{n-2}(2)$ and $\mathtt{C}_{n-2}(2)$ are equal, and for all $n \geqslant 1$ we have the following identities

$$\frac{|\,\mathtt{B}_n(2)|}{2^n(2^n+1)} = |\,\mathtt{D}_n(2)| = \frac{|\,{}^2\mathtt{D}_{n+1}(2)|}{2^{2n}(2^{n+1}+1)(2^n+1)}$$

Furthermore, $\mathtt{D}_{n-2}(2)$ is a subgroup of $\mathtt{A}_{2n-5}(2)$, and $|\,\mathtt{A}_{2n-5}(2)| < |\,{}^2\mathtt{A}_{2n-4}(2)|$ when $n \geqslant 4$. We also have $|\,\mathtt{A}_{2n-5}(2)| < |\,{}^2\mathtt{A}_{2n-2}(4)|$ Therefore, $\mathtt{D}_{n-2}(2)$ is the smallest group we are considering, and so it remains to prove that $|\,\mathtt{D}_{n-2}(2)| > |\,\mathtt{A}_{n-1}(2)|$.

$$\begin{aligned}
\frac{|\,\mathtt{D}_{n-2}(2)|}{|\,\mathtt{A}_{n-1}(2)|} &= \frac{2^{(n-2)(n-3)}(2^{n-2}-1)\prod_{i=1}^{n-3}(2^{2i}-1)}{2^{n(n-1)/2}\prod_{i=1}^{n-1}(2^{i+1}-1)} \\
&= \frac{2^{(n-2)(n-3)}(2^{n-2}-1)\prod_{i=1}^{n-3}(2^{i}-1)\prod_{i=1}^{n-3}(2^{i}+1)}{2^{n(n-1)/2}\prod_{i=1}^{n-1}(2^{i+1}-1)} \\
&= 2^{\frac{n^2-9n+12}{2}}\frac{\prod_{i=1}^{n-3}(2^i+1)}{(2^n-1)(2^{n-1}-1)} \\
&> 2^{\frac{n^2-9n+12}{2}}\frac{2^{(n-2)(n-3)/2}}{2^{2n-1}} \\
&= 2^{\frac{n^2-9n+12}{2}}2^{\frac{(n^2-9n+8)}{2}} \\
&= 2^{n^2-9n+10}
\end{aligned}$$

which is at least 1 for all $n \geqslant 8$. □

Lemma 10.5. *Let K be any version of a finite classical group of type $\mathtt{A}_k$ or ${}^2\mathtt{A}_k$ in odd characteristic. For every $n \geqslant 6$, if $k \geqslant 2n-7$ then $|K| > |\,\mathrm{L}_n(2)|$.*

Proof. By the previous discussion, it is clear that it is enough to consider the smallest rank, that is $k = 2n-7$, and the simple group K. Also, it is enough to consider $q = 3$, as the orders increase with the field – this is obvious for the universal versions, and for the simple groups follows from inspecting the sizes of the centres of the universal versions.

We have $|{}^2\mathtt{A}_k(3)| \geqslant \frac{1}{2}|\mathtt{A}_k(3)|$, and

$$\begin{aligned}
\frac{1}{2}|\mathtt{A}_{2n-7}(3)| &\geqslant \frac{1}{4}\cdot 3^{\binom{2n-6}{2}} \cdot \prod_{i=1}^{2n-7}(3^{i+1}-1)\\
&\geqslant 2^{-2}\cdot 2^{\frac{3(2n-6)(2n-7)}{4}} \cdot \prod_{i=1}^{2n-7} 2^{\frac{3i}{2}}\\
&= 2^{\frac{-8+3(2n-6)(2n-7)+3(2n-7)(2n-6)}{4}}\\
&= 2^{6n^2-39n+61}\\
&= 2^{\binom{n}{2}}\cdot 2^{\frac{11n^2-77n+122}{2}}\\
&= 2^{\binom{n}{2}}\cdot \prod_{i=1}^{n} 2^{i+1}\cdot 2^{5n^2-40n+60}\\
&> 2^{\binom{n}{2}}\cdot \prod_{i=1}^{n}(2^{i+1}-1)\\
&= |\mathtt{A}_{n-1}(2)|
\end{aligned}$$

where the last inequality holds for $n \geqslant 6$. □

Lemma 10.6. *Let $n \geqslant 8$, and let K be any version of a finite classical group of of type $\mathtt{B}_k$, $\mathtt{C}_k$, $\mathtt{D}_k$ or ${}^2\mathtt{D}_k$ in odd characteristic. IF $k \geqslant n-3$ then $|K| > |\mathrm{L}_n(2)|$.*

Proof. By the previous discussion we take $k = n-4$, $q = 3$, and the adjoint version K.

Note that $|\mathtt{B}_k(3)| = |\mathtt{C}_k(3)|$; also $|\mathtt{B}_k(3)| > |\mathtt{D}_k(3)|$. We also have $|{}^2\mathtt{D}_k(3)| \geqslant \frac{1}{2}\cdot|\mathtt{D}_k(3)|$. We then have

$$\begin{aligned}
\frac{1}{2}\cdot|\mathtt{D}_{n-3}(3)| &\geqslant \frac{1}{8}\cdot 3^{(n-3)(n-4)}\cdot(3^{n-3}-1)\cdot\prod_{i=1}^{n-4}(3^{2i}-1)\\
&> 2^{-3}\cdot 2^{\frac{3(n-3)(n-4)}{2}}\cdot 2^{\frac{3(n-3)}{2}}\cdot\prod_{i=1}^{n-4}2^{3i}\\
&= 2^{\frac{3}{2}(-2+(n-3)(n-4)+n-3+(n-3)(n-4))}\\
&= 2^{\frac{3}{2}(2n^2-13n+19)}\\
&= 2^{\binom{n}{2}}2^{\frac{1}{2}(5n^2-38n+19)}\\
&= 2^{\binom{n}{2}}\cdot\prod_{i=1}^{n}2^{i+1}\cdot 2^{\frac{1}{2}(4n^2-41n+17)}\\
&> 2^{\binom{n}{2}}\cdot\prod_{i=1}^{n}(2^{i+1}-1)\\
&= |\mathtt{A}_{n-1}(2)|
\end{aligned}$$

where the last inequality holds for $n \geqslant 10$. In the cases $n = 8$ or 9, it can be verified directly that our claim holds. □

Lemma 10.7. *For $n \geqslant 4$ and $q > 3$ odd, the simple groups $\mathtt{B}_{\frac{n}{2}}(q)$ (when n is even), $\mathtt{C}_{\frac{n+1}{2}}(q)$, $\mathtt{D}_{\frac{n+1}{2}}(q)$ and ${}^2\mathtt{D}_{\frac{n+1}{2}}(q)$ (when n is odd) are larger in cardinality than $\mathrm{L}_n(2)$.*

Proof. When n is odd, all of the orders are bounded below by the order of $\mathtt{D}_{\frac{n+1}{2}}(5)$, which is

$$\begin{aligned} 5^{\frac{n^2-1}{4}} \cdot \prod_{i=1}^{\frac{n-1}{2}} (5^{2i}-1) \cdot (5^{\frac{n+1}{2}}-1) \cdot \frac{1}{4} &\geqslant 2^{\frac{n^2-1}{2}} \cdot \prod_{i=1}^{\frac{n-1}{2}} 2^{4i} \cdot 2^{n+1} \cdot \frac{1}{4} \\ &= 2^{\frac{n^2-1}{2}+\frac{n^2-n}{2}+n-1} \\ &= 2^{n^2+\frac{n}{2}-\frac{1}{2}} \\ &> 2^{n^2} \\ &= 2^{\binom{n}{2}} \cdot 2^{\binom{n+1}{2}} \\ &> 2^{\binom{n}{2}} \cdot \prod_{i=1}^{n} (2^i-1) \\ &= |\,\mathrm{L}_n(2)| \end{aligned}$$

When n is even, we have

$$\begin{aligned} |\,\mathtt{B}_{\frac{n}{2}}(5)| &= 5^{\frac{n^2}{4}} \prod_{i=1}^{\frac{n}{2}} (5^{2i}-1) \\ &> 2^{\frac{n^2}{2}} \prod_{i=1}^{\frac{n}{2}} 2^{4i} \\ &= 2^{\frac{1}{2}(n^2+n(n+2))} \\ &= 2^{n^2+1} \\ &> |L_n(2)| \end{aligned}$$

□

CHAPTER VI

The 6-strand braid group is CAT(0)

This is joint work with Thomas Haettel and Petra Schwer.

Abstract. We show that braid groups with at most 6 strands are CAT(0) using the close connection between these groups, the associated non-crossing partition complexes, and the embeddability of their diagonal links into spherical buildings of type A. Furthermore, we prove that the orthoscheme complex of any bounded graded modular complemented lattice is CAT(0), giving a partial answer to a conjecture of Brady and McCammond.

1. Introduction

A discrete group is called CAT(0) if it acts properly discontinuously and cocompactly by isometries on a CAT(0) space. The property of being CAT(0) has far reaching consequences for a group. Algorithmically, such groups have quadratic Dehn functions and hence soluble world problem; geometrically, all free-abelian subgroups are undistorted; algebraically, the centralisers of infinite cyclic subgroups split.

In [**Cha**] Charney asked whether all braid groups are CAT(0). In this paper we give a positive answer for braid groups with at most 6 strands.

Brady and McCammond showed in [**BM2**] that the n-strand braid groups are CAT(0) if $n = 4$ or 5. However, their proof for $n = 5$ relies heavily on a computer program. They also conjectured that the same statement should hold for arbitrary n [**BM2**, Conjecture 8.4].

This paper exploits the close relationship between braid groups, non-crossing partitions of a regular n-gon, and the geometry of spherical buildings; the latter relationship was discovered by Brady and McCammond [**BM2**].

More specifically, we look at the orthoscheme complex (a certain metric polyhedral complex) associated to NCP_n, the lattice of non-crossing partitions, whose geometry was studied in [**BM2**]. Brady and McCammond showed that the CAT(0) property for braid groups can be deduced from the fact that the orthoscheme complex of the non-crossing partition lattice NCP_n is a CAT(0) space. This can be done by inspecting the diagonal link of the orthoscheme complex of NCP_n and proving that this diagonal link is CAT(1).

The diagonal link of the orthoscheme complex of the lattice NCP_n can be embedded into a spherical building of type A_{n-2}. Our approach is based on investigating the relationship between the geometry of the diagonal link of NCP_n and the ambient building. Following the criterion of Gromov (see [**Gro1**]), made precise by Bowditch (see [**Bow**]) and Charney–Davis (see [**CD1**]), the result will be implied by two facts: it is enough to show that the diagonal link is locally CAT(1) and that it does not contain any unshrinkable locally geodesic loop of length smaller than 2π. We follow this strategy in the proof of the following theorem.

Theorem 4.17. *For every $n \leqslant 6$ the diagonal link in the orthoscheme complex of non-crossing partitions NCP_n is CAT(1).*

As a consequence we obtain

Corollary 5.6. *For every $n \leqslant 6$, the n-strand braid group is CAT(0).*

We are thus giving a new proof of the theorem in case $n = 4$ or 5, and provide more evidence (with the newly covered case $n = 6$) towards [**BM2**, Conjecture 8.4]. Note that our proof at no point relies on computer-assisted calculations; it is geometric in flavour.

Brady and McCammond conjectured further that the orthoscheme complex of any bounded graded modular lattice is CAT(0) [**BM2**, Conjecture 6.10]. We are able to give a partial result towards the solution of this problem.

Theorem 4.18. *The orthoscheme complex of any bounded graded modular complemented lattice is CAT(0).*

Outline of proof. To prove our main result Theorem 4.17 we first embed the diagonal link of the orthoscheme complex of non-crossing partitions NCP_n into a spherical building. Then we assume (for a contradiction) that the diagonal link contains an unshrinkable (and hence locally geodesic) short loop. The image of such a loop l contains a positive finite number of points of special interest (called turning points) which characterise the positions at which the loop fails to be locally geodesic in the ambient space.

By inspection we show that there is a short path p between any turning point of l and the point opposite to it in l, such that p passes through a special type of a vertex, called universal. We show that any short loop passing through a universal vertex is shrinkable, and thus the two new loops obtained by following half of l and then p are short and shrinkable. Then a result of Bowditch [**Bow**] concludes the argument.

Other Artin groups. In fact Charney [**Cha**] stated a more general question about the curvature of Artin groups, and suggested that all of them should be CAT(0). Several partial answers to this question are known. Brady and McCammond studied new presentations for certain three-generator Artin groups [**BM1**] and showed that the associated presentation 2-complex admits a metric of non-positive curvature.

Charney and Davis [**CD2**] introduced the Salvetti complex, a piecewise Euclidean cube complex, associated to an Artin group. They showed that its universal cover, on which the Artin group acts geometrically, is CAT(0) if and only if the Artin group is right-angled (i.e. the exponents appearing in the presentations are either equal to 2 or ∞).

Brady [**Bra**] studied a class of Artin groups with three generators and constructed certain complexes using the associated Coxeter groups. He showed that these complexes carry a piecewise Euclidean metric of non-positive curvature and have as fundamental group the Artin groups under consideration. A generalisation of this was proved by Bell [**Bel**].

Explicit examples of Artin groups with two-dimensional Eilenberg-McLane spaces which act geometrically on 3-dimensional CAT(0) complexes (but not so on 2-dimensional ones) were constructed by Brady and Crisp [**BC**].

Acknowledgements. We would like to thank Piotr Przytycki for introducing us to the problem, and all the help provided during our work. We would also like to thank Brian Bowditch who helped us access his 'Notes on locally CAT(1) spaces' [**Bow**]. Finally, we would like to thank Ursula Hamenstädt for useful conversations, and the referee, for pointing out numerous ways of improving the presentation of this paper.

2. Definitions and preliminaries

We use this section to collect definitions of the main objects of the paper as well as their most important properties. References are given for further reading as well as for all properties listed.

2.1. Posets and lattices. This first subsection is used to introduce partially ordered sets (posets) and the geometric realization considered in this paper.

Definition 2.1 (Intervals of integers)**.** We will use $[\![n, m]\!]$ to denote the interval in $\mathbb{Z}$ between n and m (with $n \leqslant m$), that is

$$[\![n, m]\!] = [n, m] \cap \mathbb{Z}.$$

Definition 2.2 (Posets)**.** A *poset* P is a set with a partial order. P is *bounded* if it has a (unique) minimum, denoted by 0, and a (unique) maximum, denoted by 1. That is for every $x \in P$ we have

$$0 \leqslant x \leqslant 1$$

Definition 2.3 (Subposets)**.** Let $P \subseteq Q$ be two posets. We say that P is a *subposet* of Q if and only if the order on P is induced by the order on Q.

Definition 2.4 (Rank)**.** A bounded poset P has *rank* n if and only if every chain is contained in a maximal chain with $n + 1$ elements. For $x \leqslant y$ in P, the *interval* between x and y is the subposet $P_{xy} = \{z \in P : x \leqslant z \leqslant y\}$. If every interval in P has a rank, P is *graded*. If x is an element in a bounded graded poset P, then the *rank* of x is the rank of the interval P_{0x}.

Definition 2.5 (Joins, meets, lattices)**.** A poset P is called a *lattice* if and only if for every $x, y \in P$ the following two conditions are satisfied:

- there exists a unique minimal element $x \vee y$ of the set
$$\{z \in P \mid x \leqslant z \text{ and } y \leqslant z\},$$
called the *join* of x and y, and
- there exists a unique maximal element $x \wedge y$ of the set
$$\{z \in P \mid x \geqslant z \text{ and } y \geqslant z\},$$
called the *meet* of x and y.

Definition 2.6 (Linear lattice)**.** If V is an n-dimensional vector space over a division algebra, we will denote by $S(V)$ the rank n lattice consisting of all vector subspaces of V, with the order given by inclusion. We call $S(V)$ the *linear lattice* of V.

It is easy to see that $S(V)$ is indeed a lattice, where the meet of two linear subspaces can be taken to be their intersection and the join is given by their common span.

Definition 2.7 (Failing modularity)**.** Let P be a subposet of a linear lattice $S(V)$. We say that two elements $x, y \in P$ *fail modularity* (with respect to P) if and only if their join or their meet in $S(V)$ is not contained in P.

When it is clear from the context in which pair $P \subseteq S(V)$ we are working we sometimes just say x and y fail modularity.

Definition 2.8 (Realisations)**.** Let P be a graded poset. The *simplicial realisation* $\|P\|$ of P is the simplicial complex whose vertex set is P, and whose k-simplices correspond to chains $x_0 < x_1 < \ldots < x_k$ of length k.

A *geometric realisation* of P is a metric space X together with a homeomorphism $X \to \|P\|$.

Note that in particular $\|P\|$ endowed with the standard piecewise-Euclidean metric is a geometric realisation of P. This is however *not* the metric we will study in this paper; the metric of our interest will be defined in Definition 2.17.

Observe that for a bounded poset P the edge connecting 0 to 1 is contained in every maximal simplex.

Notation 2.9 (Geometric realisation)**.** Given a (fixed) geometric realisation of P, we can (and will) treat X as a simplicial complex via the given homeomorphism $X \to \|P\|$. Thus we will continue to use the simplicial complex vocabulary when talking about X and we will write $|P|$ to denote X together with its simplicial structure inherited from $\|P\|$.

We will also use some standard buildings terminology in the setting of simplicial complexes:

Notation 2.10. In a simplicial complex a *vertex* is a 0-simplex, *faces* are simplices and *chambers* are maximal simplices.

As we will never look at more than one geometric realisation of any poset, we sometimes abuse notation by using P to denote both a poset and some fixed geometric realisation thereof.

Definition 2.11 (Adjacency)**.** Given a poset P, an element $y \in P$ is said to be *adjacent* to a chain $x_0 < \cdots < x_k$ in P if and only if y does not belong to the chain, and there exists a chain containing both y and all the elements x_i.

In the setting of the simplicial realisation $\|P\|$, a vertex y is adjacent to a face $F = \{x_0, \ldots, x_k\}$ if and only if y does not belong to F, and $F \cup \{y\}$ is itself a face. Equivalently, y is adjacent to F if and only if for each vertex $x_i \in F$ there is an edge connecting x_i to y.

Definition 2.12 (Diagonal link)**.** Given a geometric realisation $|P|$ of a bounded lattice P, we define the *diagonal link* of $|P|$ to be the link

$$LK(e_{01}, |P|)$$

of the *diagonal edge* e_{01}, that is the edge connecting the minimum 0 to the maximum 1.

Note that if $|P|$ has a piecewise Euclidean or spherical metric, then $LK(e_{01}, |P|)$ carries a natural angular metric and is hence itself a geometric realisation of the poset $P \setminus \{0, 1\}$.

Remark 2.13. A fact we will frequently use is that the vertices of $LK(e_{01}, |P|)$ are in a natural way in one to one correspondence with elements of $P \setminus \{0, 1\}$ as follows. A vertex, i.e. 0-simplex in $LK(e_{01}, |P|)$ corresponds to a 2-simplex in $|P|$ whose vertices are $0, 1$ and one additional vertex $p \in P$. We may thus label the corresponding 0-simplex in the link by p. So when we refer to a vertex p in $LK(e_{01}, |P|)$ what we mean is the 0-simplex in the link which corresponds to the 2-simplex in $|P|$ spanned by p and $0, 1$.

2.2. Spherical buildings and orthoscheme complexes. First we will very quickly introduce spherical buildings and some of their basic properties. In the rest of this section we focus on the spherical buildings of type A_n, which arise from the lattice of linear subspaces of a vector space. In the rest of the paper we will use the standard terminology of spherical buildings freely and refer the reader to the book by Abramenko and Brown [**AB**] for further details.

Definition 2.14. A (spherical) *building* is a simplicial complex B which is the union of a collection of subcomplexes A, called *apartments*, satisfying the following axioms:

(B0) Each apartment is isomorphic to a (finite) Coxeter complex.
(B1) For any two simplices c, d in B there exists an apartment containing both.
(B2) If A_1 and A_2 are two apartments containing simplices c and d, then there exists an isomorphism $A_1 \to A_2$ fixing c and d pointwise.

The maximal simplices in B are called *chambers*.

Note that c, d are allowed to be empty in axiom (B2) and hence any two apartments are isomorphic. We call the type of the Coxeter group the *type* of the building. Note further that B is a chamber complex, that is any two maximal simplices have the same dimension.

For any spherical building B there is a standard geometric realisation of B which induces on each apartment the round metric of a sphere. Throughout the paper we will consider a spherical building B simultaneously as a simplicial complex and a metric space using this standard geometric realisation.

Definition 2.15. Two points x, y in a spherical building X are called *opposite* if for some (any) apartment A containing x and y, x and y are opposite in the apartment A, seen as a round sphere. Equivalently, the distance between x and y in X is π. Two faces F,F' in X are called *opposite* if for some (any) apartment A containing F and F', F and F' are opposite in the apartment A.

Proposition 2.16. *Let B be a spherical building. Then*

(1) the link $\operatorname{lk}_B(c)$ of any simplex c in B is a spherical building.
(2) apartments are metrically convex, in other words for any apartment A containing a pair of points $x, y \in B$ one has $d_B(x, y) = d_A(x, y)$.
(3) B is a CAT(1) space when equipped with the standard metric.

Proof. For a proof of item 1 see Proposition 3 on page 79 in the book of Brown [**Bro1**]. Proofs of the other items are contained in the book of Bridson–Haefliger [**BH**]. Item 2 follows from Lemma II.10A.5 and item 3 is Theorem II.10A.4 therein. □

We are interested in one particular type of buildings namely the spherical buildings of type A_n. To see what they are recall that if V is an n-dimensional vector space over a division algebra, the linear lattice $S(V)$ of V is the rank n lattice consisting of all vector subspaces of V, with the order given by inclusion. One can equip the simplicial realization of a building with a so called orthoscheme metric that will allow us to explicitly describe the standard CAT(1) metric on buildings of type A.

Definition 2.17 (Orthoscheme complex)**.** Let P be a bounded graded poset. A maximal chain $x_0 < \cdots < x_n$ in P corresponds to an n-simplex F in the simplicial realisation $\|P\|$ of P. We endow this simplex with a metric in the following way: we identify each x_i with the vertex $(1, \ldots, 1, 0, \ldots, 0)$ (i times "1") in $\mathbb{R}^n$. We give F the metric of the Euclidean convex hull of the vertices in the cube. Equivalently, it is the metric on simplices of the barycentric subdivision of the Euclidean n-cube with side length 2.

Note that the distance between two vertices lying in a common simplex depends only on the difference in their rank. We can endow each maximal simplex (i.e. chamber) in $\|P\|$ with this metric in a coherent way. The induced length metric on the whole complex is the *orthoscheme metric*. This way $\|P\|$ becomes a geometric realisation of P, which is called the *orthoscheme complex* of P.

For more information about the orthoscheme complexes we refer the reader to Brady and McCammonds exposition in [**BM2**]. Brady–McCammond [**BM2**] show the following.

Proposition 2.18. *When V is an n-dimensional vector space over a division algebra, then the diagonal link $LK(e_{01}, |S(V)|)$ of the orthoscheme complex $|S(V)|$ is equal to the (standard CAT(1) realisation of the) spherical building associated to* $\mathrm{PGL}(V)$*, which is a spherical building of type A_{n-1}. The dimension of this building is $n-2$.*

The above proposition is crucial for the proof of our main result. It is precisely the geometry of this building and (a specific class of) its subcomplexes that we will focus on.

One can show that apartments of B, the building associated to $\mathrm{PGL}(V)$, are in one-to one correspondence with bases of V.

Note that, according to Remark 2.13, the vertex set of $LK(e_{01}, |S(V)|)$ has the structure of a bounded graded lattice if only we add to it the minimum 0 (corresponding to the trivial subspace) and the maximum 1 (corresponding to the improper subspace). Because of this deficiency let us use the following convention.

Definition 2.19. We say that a subset M of the vertex set of the diagonal link $LK(e_{01}, |S(V)|)$ is *stable under joins and meets* if and only if for every $x, y \in M$ we have $x \vee y \in M \cup \{1\}$ and $x \wedge y \in M \cup \{0\}$.

Definition 2.20. Let $M \subseteq B$ be a subset of a building. The *simplicial convex hull* of M is defined to be the intersection of all apartments of B containing M.

Lemma 2.21. *Let V be an n-dimensional vector space over a division algebra, and write $B = LK(e_{01}, |S(V)|)$ for the diagonal link of the orthoscheme complex $|S(V)|$ of $S(V)$ (and hence a spherical building of type A_{n-1}). Let F and F' be two faces in B. Consider the minimal set M of vertices of B containing the vertices of F and F' which is stable under joins and meets. Then the full subcomplex spanned by M is equal to the simplicial convex hull of $F \cup F'$.*

Proof. Let H denote the simplicial convex hull of $F \cup F'$ in B. Let A be an apartment in B containing $F \cup F'$, and let $\{e_1, \dots, e_n\}$ be a basis of V corresponding to A. As every element of M arises as a joins or meet of vertices of F and F' every element of M is spanned by some elements of $\{e_1, \dots, e_n\}$, so it belongs to the apartment A. Since this is true for every apartment containing $F \cup F'$ it holds for their intersections and hence we have proved that $M \subseteq H$.

To show the converse let F and F' be two simplices and suppose there exists some $v \in H \setminus M$. We will show that there exists an apartment containing M but not v.

Let A be an apartment containing M and let $\{e_1, \dots, e_n\}$ be a basis of V corresponding to A. Suppose that for each i with $\mathrm{e_i} \leqslant v$ there exists $m_i \in M$ with $e_i \leqslant m_i \leqslant v$. Then we have

$$v = \bigvee_I e_i \leqslant \bigvee_I m_i \leqslant v$$

where I denotes the set of i such that $e_i \leqslant v$. But then v is a join of elements in M, and thus is itself in M as M is closed under taking joins. We conclude that there exists $i \in [\![1, n]\!]$ such that $e_i \leqslant v$, and

$$\forall m \in M, e_i \leqslant m \Longrightarrow m \nleqslant v.$$

Consider $m_0 = \bigwedge\{m \in M \mid e_i \leqslant m\} \in M$, where $m_0 = V$ if there is no $m \in M$ such that $e_i \leqslant m$. Since $e_i \leqslant m_0$, we know that $m_0 \nleqslant v$, so there exists $j \in [\![1, n]\!] \setminus \{i\}$ such that $e_j \leqslant m_0$ and $e_j \nleqslant v$. Then the apartment A' corresponding to the basis $\{e_1, \dots, e_{i-1}, e_i + e_j, e_{i+1}, \dots, e_n\}$ contains M but not v. But then $v \notin H$, which is a contradiction. So we have proved that $H \subseteq M$. ☐

2.3. Non-crossing partitions. Let us now introduce the pivotal objects in this article, non-crossing partitions. We will see that they form a sublattice of the linear subspace lattice of a vector space.

Definition 2.22 (Partition lattice)**.** Let U_n be the set of n^{th} roots of unity inside the plane $\mathbb{C}$. A *partition* of U_n is a decomposition of the set U_n into disjoint subsets, called *blocks*, such that their union is U_n. Let P_n denote the set of all partitions of the set U_n. The set P_n forms a bounded graded lattice of rank $n-1$, where the order is given by: $p \leqslant p'$ if and only if every block of p is contained in a block of p'.

Definition 2.23 (Non-crossing partition lattice)**.** A partition of U_n is called *non-crossing* if for every distinct blocks x, y of the partition, the convex hulls $Hull(x)$ and $Hull(y)$ in $\mathbb{C}$ do not intersect. We define NCP_n to be the subposet of P_n consisting of non-crossing partitions of U_n. Then NCP_n is a bounded graded lattice of rank $n-1$.

Lemma 2.24 (NCP is a subposet of $S(V)$)**.** *For every $n \geqslant 2$, P_n and NCP_n are isomorphic to subposets of $S(V)$, where V is an $(n-1)$-dimensional vector space.*

Proof. Fix a field $\mathbb{F}$, and let $V = \{(y_i) \in \mathbb{F}^n \mid \sum_{i=1}^n y_i = 0\}$. Then V is an $(n-1)$-dimensional $\mathbb{F}$-vector space. Identify U_n with $[\![1, n]\!]$. If $x \in P_n$ let

$$f(x) = \big\{(y_i) \in V \mid \forall \text{ block } Q \in x : \sum_{i \in Q} y_i = 0\big\}.$$

Then f is an injective rank-preserving poset map from P_n to $S(V)$. It clearly restricts to $NCP_n \subseteq P_n$. □

Definition 2.25 (Non-crossing partition complex)**.** We will refer to the orthoscheme complexes of the non-crossing partition lattices NCP_n as the *non-crossing partition complexes*. It is the simplicial realization of NCP_n equipped with the orthoscheme metric, as defined in Definition 2.17.

Lemma 2.26 (Duality)**.** *For $n \geqslant 2$, there is a duality on NCP_n, i.e. an order-reversing bijection $x \mapsto x^*$ from NCP_n to itself.*

Proof. Denote by $\{\omega_k\}_{k \in \mathbb{Z}/n\mathbb{Z}} = U_n$ the n^{th} roots of unity. If x is a non-crossing partition of U_n, then its dual x^* is the partition of the shifted set

$$U_n^* = \{m_k = e^{\frac{\pi i}{n}} \omega_k\}_{k \in \mathbb{Z}/n\mathbb{Z}}$$

defined by: m_k and m_j belong to the same block of x^* if and only if the geodesic segment $[m_k, m_j]$ in $\mathbb{C}$ does not intersect the convex hull of any block of x. Then x^* is a non-crossing partition of U_n^*, with $\mathrm{rk}\,(x^*) = n - 1 - \mathrm{rk}\,(x)$, and $(x^*)^* = x$. Now choose some identification between U_n^* and U_n (like multiplying by $e^{\frac{-\pi i}{n}}$) to get a map from NCP_n to itself. □

Note that we will only use duality to reduce the number of cases that will need checking in the later stage of our proof.

2.4. CAT(0) and CAT(1) spaces. In this section we will state the definitions of CAT(0) and CAT(1) spaces and recall some of Bowditch's results about locally CAT(1) spaces (see [**Bow**]). Moreover we recall how Brady and McCammond use Bowditch's criteria to give a sufficient condition for braid groups to be CAT(0) (see [**BM2**]). For a general discussion of CAT(κ) spaces we refer the reader to the book by Bridson and Haefliger [**BH**].

Definition 2.27. Let X be a geodesic metric space. A *geodesic triangle* Δ is formed by three geodesic segments, $\gamma_i\colon [0, l_i] \to X$ with $i \in \mathbb{Z}/3\mathbb{Z}$, such that $\gamma_i(l_i) = \gamma_{i+1}(0)$.

Given such a triangle, we form the *Euclidean comparison triangle* $\Delta' \subset \mathbb{R}^2$ by taking any triangle whose vertices x_1, x_2, x_3 satisfy

$$d_{\mathbb{R}^2}(x_i, x_{i+1}) = l_i.$$

There is an obvious map $c\colon \Delta \to \Delta'$, isometric on edges, sending $\gamma_i(0) \mapsto x_i$; we will refer to it as a *comparison map.*

We say that X is *CAT(0)* if and only if for any two points x, y on any geodesic triangle Δ, we have

$$d_X(x, y) \leqslant d_{\mathbb{R}^2}(c(x), c(y)).$$

Definition 2.28. Given a geodesic triangle Δ with notation as above, with the additional condition that $l_1 + l_2 + l_3 \leqslant 2\pi$, we form the *spherical comparison triangle* $\Delta'' \subset S^2$ by taking any triangle whose vertices x_1, x_2, x_3 satisfy

$$d_{S^2}(x_i, x_{i+1}) = l_i.$$

Again there is an obvious map $c\colon \Delta \to \Delta''$, isometric on edges, sending $\gamma_i(0) \mapsto x_i$; we will refer to it as a *comparison map.*

We say that X is *CAT(1)* if and only if for any two points x, y on any geodesic triangle Δ with perimeter at most 2π, we have

$$d_X(x, y) \leqslant d_{S^2}(c(x), c(y))$$

Definition 2.29. A group G has the *CAT(0) property*, or *is CAT(0)*, if and only if it acts properly discontinuously and cocompactly by isometries on a CAT(0) space.

Definition 2.30 (Locally CAT(1))**.** A complete, locally compact, path-metric space X is said to be *locally* CAT(1) if each point of X has a CAT(1) neighbourhood.

Definition 2.31 (Shrinking and shrinkable loops)**.** Let X be a complete, locally compact path-metric space. A rectifiable loop l in X is said to be *shrinkable to l'* if and only if l' is another rectifiable loop in X, and there exists a homotopy between l and l' going through rectifiable loops of non-increasing lengths.

A rectifiable loop l is *shrinkable* if and only if it is shrinkable to a constant loop.

The loop l is said to be *short* if its length is smaller than 2π.

Theorem 2.32 (Bowditch [**Bow**, Theorem 3.1.2])**.** *Let X be a locally CAT(1) space. Then X is CAT(1) if and only if every short loop is shrinkable.*

The following theorem will be an important tool in our argument.

Theorem 2.33 (Bowditch [**Bow**, Theorem 3.1.1])**.** *Let X be a locally CAT(1) space. Let $x, y \in X$, and consider three paths $\alpha_1, \alpha_2, \alpha_3\colon [0,1] \to X$ joining x to y. For all $i \in \{1,2,3\}$, consider the loop $\gamma_i = \alpha_{i+1}^{-1} \circ \alpha_i$ based at x (with indices modulo 3). Assume that for all $i \in \{1,2,3\}$ the loop γ_i is short. Assume further that γ_1 and γ_2 are shrinkable. Then γ_3 is shrinkable.*

Theorem 2.34 (Brady, McCammond [**BM2**, Theorem 5.10 and Lemma 5.8])**.** *Assume that for all $3 \leqslant k \leqslant n$, the diagonal link of $|NCP_k|$ does not contain any unshrinkable short loop. Then $|NCP_n|$ is CAT(0).*

Proposition 2.35 (Brady, McCammond [**BM2**, Proposition 8.3])**.** *If $|NCP_m|$ is CAT(0) for all $m \leqslant n$, then the n-strand braid group is CAT(0), that is it acts geometrically on a CAT(0) space.*

3. Turning points and turning faces

In order to get ready for the proof of our main result Theorem 4.17, we introduce further tools: turning points and turning faces. Some of their properties hold more generally in arbitrary path-connected subset of metric spaces equipped with the induced length metric. We collect definitions and properties in the following two subsections.

3.1. Turning points. Turning points are points on locally geodesic loops in a subspace of a metric space where said loop fails to be a local geodesic in the ambient space. The precise definition is a follows.

Definition 3.1 (Turning points)**.** Let X be a path-connected subspace of a geodesic metric space B, and endow X with the induced length metric. Suppose that $l\colon D \to X$ is a local isometry, where D is a metric space. We say that a point $t \in D$ is a *turning point* of l in B if and only if $i \circ l$ fails to be a local isometry at t, where $i\colon X \to B$ is the inclusion map.

Definition 3.2 (Locally geodesic loops)**.** Let X be a metric space. We say that $l\colon S^1 \to X$ is a *locally geodesic loop* in X if and only if l is a local isometry, where S^1 is given the length metric of the quotient of some closed interval I of $\mathbb{R}$ by its endpoints. The *length* of l is defined to be the length of I.

We say that $l\colon I \to X$ is a *locally geodesic path* in X if and only if l is a local isometry, where I is a closed interval of $\mathbb{R}$. The *length* of l is defined to be the length of I.

Lemma 3.3. *Let X be a path-connected subset of a geodesic metric space B, and endow X with the induced length metric. Suppose that $l\colon D \to X$ is a locally geodesic path or loop, with D being respectively I or S^1. Let $t \in D$. Suppose that there exists a subset $N \subseteq X$, such that N contains the convex hull in B of the image under l of some neighbourhood of t in D. Then t is not a turning point.*

Proof. Suppose (for a contradiction) that t is a turning point. As $i \circ l$ fails to be a local geodesic at t, there exist $t_1, t_2 \in D$ in a neighbourhood of t such that

$$d_B\big(l(t_1), l(t_2)\big) < d(t_1, t_2)$$

and such that $l(t_j) \in N$ for $j \in \{1, 2\}$.

Since B is a geodesic metric space, we can realise the distance between $l(t_1)$ and $l(t_2)$ with a geodesic segment g in B. Since N contains the endpoints of g, it contains the whole of g. Hence in particular g lies in X, which (as was claimed) contradicts the fact that l was a local geodesic. □

Remark 3.4. We will often identify (isometrically) a neighbourhood of a point $t \in S^1$ with an interval in $\mathbb{R}$ containing t in its interior. We will therefore feel free to write $[t - \varepsilon, t + \varepsilon]$ etc. (for a small ε) to denote a subset of S^1.

Definition 3.5 (Consecutive turning points)**.** Suppose that we have a subset $T \subset S^1$. We will say that $t, t' \in T$ are *consecutive* if and only if there is a path in S^1 with endpoints t and t' not containing any other point in T. A shortest such path will be denoted by $[t, t']$.

Remark 3.6. Note that $[t, t']$ defined above is unique provided that the cardinality of T is at least 3.

Lemma 3.7. *Let X be a path-connected subset of a CAT(1) space B, and endow X with the induced length metric. If $l\colon S^1 \to X$ is a locally geodesic loop in X of length $0 < L < 2\pi$, then the cardinality of the set of turning points T of l is greater than 2. Moreover, $l|_{[t,t']}$ is a geodesic in B for any pair of consecutive turning points t, t'.*

Proof. Suppose that we can find three distinct points $t_1, t_2, t_3 \in S^1$ such that each pairwise distance is strictly bounded above by $\frac{1}{2}L < \pi$, and such that T is contained in $[t_1, t_2] \cup \{t_3\}$, where $[t_i, t_j]$ denotes the shortest segment of S^1 with endpoints t_i and t_j not containing t_k for

$$\{i, j, k\} = [\![1, 3]\!].$$

Note that $l|_{[t_1,t_3]}$ and $l|_{[t_2,t_3]}$ are geodesics in B – this follows from the fact that local geodesics of length at most π are geodesics in CAT(1) spaces (essentially because the statement is true for the 2-sphere S^2).

Consider a geodesic triangle $\Delta = l(t_1)l(t_2)l(t_3)$ in B, and let Δ' be the comparison triangle in S^2. Since $L < 2\pi$, the perimeter of Δ (and hence also of Δ') is smaller than 2π. Therefore Δ' cannot be a great circle in S^2.

Suppose that $t_3 \notin T$. Then the angle of Δ at $l(t_3)$ is equal to π, and the same is true in the comparison triangle Δ' (by the CAT(1) inequality). But then the triangle is degenerate, and hence so is Δ. In particular the geodesic from $l(t_1)$ to $l(t_2)$ goes via $l(t_3)$. But this contradicts the assumption that the distance (in B) between $l(t_1)$ and $l(t_2)$ is smaller than $\frac{1}{2}L$. So $t_3 \in T$. We will now use this trick to prove our claims.

If $|T| \leqslant 1$ then we immediately get a contradiction by taking either any three points in S^1 satisfying the conditions above, or the turning point and two other points so that the triple satisfies the condition.

If $|T| = 2$ and the two points are not antipodal in S^1, then we can always (very easily indeed) find a third point so that the triple satisfies our condition. If the turning points are antipodal, then the two local geodesics given by l, which connect the images of the turning points, coincide. This is because local geodesics of length smaller than π are geodesics in B, and such geodesics in CAT(1) spaces are unique. But then l cannot be a local geodesic in X. We have thus shown that $|T| \geqslant 3$.

Now suppose we have two consecutive turning points, t_1 and t_2. If $[t_1, t_2]$ is of length at most π, then $l|_{[t_1,t_2]}$ is a geodesic as before. If the length is larger than π, then in particular it is larger than $\frac{1}{2}L$, and so we can take the midpoint $t_3 \in [t_1, t_2]$ and (applying the argument above) conclude that $t_3 \in T$, which in turn contradicts the definition of $[t_1, t_2]$. □

3.2. Turning faces. We are mainly interested in locating turning points on loops in linearly embedded subcomplexes of the orthoscheme complex of linear subspaces of a vector space. To understand their behaviour we use properties of the supporting faces which will be called turning faces and are introduced in this section.

Definition 3.8. Let V be an n-dimensional vector space over a division algebra and denote by B be the diagonal link $LK(e_{01}, |S(V)|)$ of the orthoscheme complex $|S(V)|$ of $S(V)$. Hence B is a spherical building of type A_{n-1} equipped with the standard CAT(1) metric. We say that a geometric realization X of a simplicial complex is *linearly embedded in* B if and only if there exists a bounded graded lattice P of rank n, such that

(1) P is a subposet of $S(V)$ with a geometric realisation $|P|$ isometric to the full subcomplex of $|S(V)|$ spanned by P, and
(2) X is isometric (and isomorphic as simplicial complexes) to the diagonal link $LK(e_{01}, |P|) \subseteq B$ equipped with the length metric induced from B.

We will call the metric on X the *spherical orthoscheme* metric. Note that X has dimension $n - 2$.

Since a linearly embedded X is a subcomplex of a building, we will use the standard buildings vocabulary when talking about X. Hence *a chamber* or *an* *apartment* in X is a chamber or an apartment of B, respectively, which is contained in X. Note also that since the ranks of P and $S(V)$ agree, the complex X is a union of chambers.

For the remainder of this subsection let X be linearly embedded in B.

Definition 3.9 (Rank and corank)**.** Let F be a face of codimension m in X. Then F is the span of vertices $x_1, \dots, x_{n-1-m}$ of ranks $r_1, \dots, r_{n-1-m}$. We define the set $\mathrm{rk}\,(F) = \{r_1, \dots, r_{n-1-m}\}$ to be the *rank* of F, and the set $\mathrm{crk}(F) = [\![1, n-1]\!] \smallsetminus \mathrm{rk}\,(F)$ to be the *corank* of F.

Definition 3.10 (Turning face)**.** Let t be a turning point of a locally geodesic loop $l\colon S^1 \to X$. The smallest (with respect to inclusion) intersection F of a chamber in X containing $l([t, t-\varepsilon))$ and a chamber in X containing $l([t, t+\varepsilon))$, for sufficiently small $\varepsilon > 0$, will be called the *turning face* of t.

Lemma 3.11. *Let l be a locally geodesic loop or path in X. Then the set T of turning points of l is finite.*

Proof. Let $l\colon D \to X$, where $D = S^1$ or $D = I = [0, L]$ and $l(0) \neq l(L)$. If $D = I$, notice that for $\varepsilon > 0$ small enough $l([0, \varepsilon])$ is contained in a chamber in X, hence by Lemma 3.3 the point $0 \in I$ is not a turning point of l, and similarly neither is L.

We claim that T is a discrete subset of D. Suppose that $t \in T$ is a turning point. Let F be the turning face of t. By definition, there exists $\varepsilon > 0$ such that F is the intersection of a chamber C^- in X containing $l([t, t-\varepsilon))$ and a chamber C^+ in X containing $l([t, t+\varepsilon))$.

Then $l|_{[t-\varepsilon, t]}$ is the geodesic segment from $l(t-\varepsilon)$ to $l(t)$ in $C^- \subset X$. In particular it is also locally geodesic in B, so there is no turning point in $(t-\varepsilon, t)$, and similarly in $(t, t+\varepsilon)$ and C^+. Therefore T is discrete.

Note that T is closed – this follows directly from the fact that if $t \in D$ is not a turning point, then l is a geodesic (in B) at some open neighbourhood of t in D, and so in particular none of the points in this open neighbourhood are turning points themselves. Hence T is closed, and therefore compact since D is compact. We have thus shown that T is compact and discrete, and so it is finite. □

The following lemma is the first result which gives us some combinatorial control over the turning points. It is precisely this type of control which will allow us to perform the inspection in the proof of the main theorem.

Lemma 3.12. *Let l be a locally geodesic loop in X. Then for every turning point of l in B, its turning face has a corank which contains two consecutive integers.*

Proof. Suppose that we have $t \in S^1$, a turning point of l, whose image x under l is contained in the turning face F in X. By definition, there exists $\varepsilon > 0$ such that $l((t-\varepsilon, t]) \subseteq C^-$ and $l([t, t+\varepsilon)) \subseteq C^+$, where C^- and C^+ are chambers of X such that $F = C^+ \cap C^-$.

Assume that the corank of F does not contain two consecutive integers. Then the sets of vertices of C^- and C^+ differ at vertices of ranks

$$1 \leqslant r_1 < \dots < r_k \leqslant n-2$$

with $\forall\ i \in [\![1, k-1]\!]\ :\ r_{i+1} - r_i \geqslant 2$. Then for every $\varepsilon = (\varepsilon_1, \dots, \varepsilon_k) \in \{\pm\}^k$, consider the chamber C^ε spanned by $C^+ \cap C^-$ and, for every $1 \leqslant i \leqslant k$, by the vertex of rank r_i in C^{ε_i}. Since all vertices of C^ε belong to X, we know that C^ε belongs to X.

By item 1 of Proposition 2.16 the link of $F = C^+ \cap C^-$ in B is itself a spherical building. This is easily seen by taking successive links of single vertices of F.

Taking the first such link, we either get a building of type A_{n-2} (when the vertex we removed was of rank 1 or $n-1$), or a building of type $A_l \times A_{n-2-l}$ (when the vertex was neither of minimal nor maximal rank). Repeating this process we obtain a spherical building whose type is determined precisely by the structure of the corank of F; more specifically, when the corank has no consecutive integers, the building is of type A_1^k. Thus the apartments are spherical joins of k copies of the 0-sphere.

Let $N = \bigcup_{\varepsilon \in \{\pm\}^k} C^\varepsilon$. Observe that the image of N in the link of F in B is precisely one of the apartments, and therefore it is convex. Since the link $\mathrm{lk}(x, N)$ of x in N is isometric to the spherical join $\mathrm{lk}(x, F) * \mathrm{lk}(F, N)$, it is convex in the link $\mathrm{lk}(x, B) \simeq \mathrm{lk}(x, F) * \mathrm{lk}(F, B)$ of x in B. For $\delta > 0$ small, the δ-ball around x is isometric to the δ-ball around the cone point in the cone over the link of x, according to [**BH**, Theorem 7.16]. Since $\mathrm{lk}(x, N)$ is convex in $\mathrm{lk}(x, B)$, we conclude that the δ-ball around x in N is convex in B. Since N contains the image under l of some neighbourhood of t in S^1, using Lemma 3.3 we show that t is not a turning point. □

Lemma 3.13. *Let $l\colon I \to X$ be a locally geodesic segment in X with a turning point t in B. Let E^+ (respectively E^-) be minimal faces in X containing the image under l of a right (respectively left) ε-neighbourhood of t for some $\varepsilon > 0$. Then the simplicial convex hull of $E^+ \cup E^-$ is not contained in X.*

Proof. Let N denote the simplicial convex hull of $E^- \cup E^+$ in B. Suppose (for a contradiction) that N is contained in X. The subcomplex N is metrically convex (as it is an intersection of apartments, which are metrically convex in B), and contains the image under l of some neighbourhood of t in S^1. Therefore using Lemma 3.3, we show that t is not a turning point, and this concludes the proof. □

4. Proof of the main theorem

In this section we will prove our main result. Let us first fix some notation and recall a few facts.

Denote by $|NCP_n|$ the orthoscheme complex of the non-crossing partition lattice for some $n \geqslant 3$, and let X denote the diagonal link $LK(e_{01}, |NCP_n|)$ equipped with the spherical orthoscheme metric. Thus X is the geometric realization of an $n-3$-dimensional simplicial complex whose vertices are in one-to-one correspondence with the partitions in $NCP_n \setminus \{0, 1\}$, see Remark 2.13. By the rank of X we mean the rank of the poset $NCP_n \setminus \{1\}$, which is $n-2$.

Recall that X is linearly embedded in a spherical building B, which is the diagonal link of the linear subspace lattice of an $(n-1)$-dimensional vector space V by Lemmata 2.24 and 2.18. Note that this building B has type A_{n-2} and is a simplicial complex of dimension $n-3$.

Moreover, B is a CAT(1) metric space, which is why (as Brady and McCammond remarked, see [**BM2**, Remark 8.5]) the spherical orthoscheme metric on X is a good candidate to be CAT(1) for all $n \geqslant 3$.

Recall from Remark 2.13 that the vertices of X are naturally labeled by elements of NCP_n. We may thus talk about partitions in X.

Remark 4.1. Note that for $n = 3$, the diagonal link X in NCP_3 is the disjoint union of 3 points, so it is CAT(1). For $n = 4$, the diagonal link X in NCP_4 is a subgraph of the incidence graph of the Fano plane, so it has combinatorial girth 6. Since each edge has length $\frac{\pi}{3}$, its girth is $6\frac{\pi}{3} = 2\pi$, so X is CAT(1). A picture of the diagonal link of NCP_4 can be found in Figure 4.1.

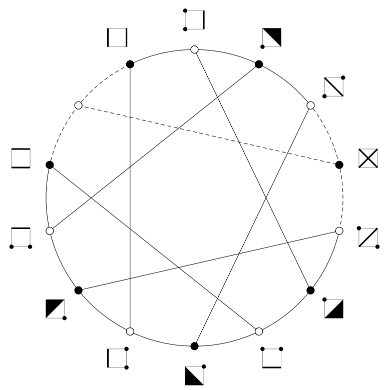

FIGURE 4.1. The diagonal link of NCP_4 is shown with solid lines. Dotted lines represent the missing two vertices of the Fano plane, one of which is a vertex of the diagonal link of the geometric realisation of the partition lattice.

Definition 4.2 (Non-crossing trees). A *non-crossing forest* of U_n is a metric forest embedded in $\mathbb{C}$ with vertex set U_n, whose edges are geodesic segments in $\mathbb{C}$. When such a forest has only one connected component, we call it a *non-crossing tree*

Remark 4.3. Note that every non-crossing forest corresponds to an element in NCP_n. The correspondence is obtained by saying that two points in U_n lie in the same block if and only if they lie in a single connected component of the forest. In particular this gives a one-to-one correspondence between vertices of rank 1 in X (that is the corresponding partition is of rank 1) and non-crossing forests with only one edge.

This way we can also associate a subset of X to a non-crossing tree by taking the span of all vertices associated to proper subforests of our non-crossing tree.

Proposition 4.4. *Let A be an apartment in B. Then A is included in X if and only if its $(n-1)$ rank 1 vertices lie in X and correspond to the edges of a non-crossing tree.*

Proof. Suppose A is an apartment lying in X. Each rank 1 vertex v_i of A corresponds to a basis vector ϵ_i of V, the vector space that is used to define B as the diagonal link in $S(V)$. Each such vertex also corresponds to an edge $e_i \subset \mathbb{C}$ connecting two points in U_n, as explained in the remark above. We claim that T, the union of edges e_i, is an embedded tree.

Let v_i and v_j be two distinct vertices of A of rank 1. Then their join in B has rank 2 (it is the plane $\langle \epsilon_i, \epsilon_j \rangle$), and lies in A. But $A \subseteq X$, and so the partition $v_i \vee v_j$ has rank 2. Observe that if e_i intersects e_j away from U_n, then the join $v_i \vee v_j$ has to contain the convex hull of $e_i \cup e_j$ as a block (since it is non-crossing), and therefore its rank is at least 3. Hence e_i can intersect e_j only at U_n, and therefore T is embedded.

Now suppose that T contains a cycle. Without loss of generality let us suppose that the shortest cycle is given by the concatenation of edges $e_1, \ldots, e_k$ for some k. Then note that the joins in X satisfy

$$\bigvee_{i=1}^{k} v_i = \bigvee_{i=1}^{k-1} v_i.$$

But, as before, they are equal to the joins in B (since $A \subseteq X$). This yields the equality

$$\langle \epsilon_1, \ldots, \epsilon_k \rangle = \langle \epsilon_1, \ldots, \epsilon_{k-1} \rangle,$$

which contradicts the fact that the vectors ϵ_i are linearly independent.

We have thus shown that T is an embedded forest. But T consists of $n-1$ edges, and so an Euler characteristic count yields that it has exactly one connected component. Therefore T is a tree as required.

Now suppose the vertices of rank 1 of an apartment A lie in X and form a non-crossing tree T. The apartment is the span of the closure of the set of its rank 1 vertices under taking joins in B. The fact that T is non-crossing tells us that the joins of these vertices taken in B or X coincide, and hence all vertices of A lie in X. Therefore $A \subseteq X$. □

Definition 4.5 (Universal points)**.** A point $x \in X$ is said to be *universal* if it belongs to a face in X, all of whose vertices are partitions with exactly one block containing more than one element, and such that this block only contains consecutive elements of U_n. Such a face is also called *universal.*

Note that every universal face is contained in a universal chamber.

Example 4.6. Consider for $n = 6$ the edge between the two partitions shown in Figure 4.2, that is

$$\{\{1,2,3\},\{4\},\{5\},\{6\}\} < \{\{1,2,3,4,5\},\{6\}\}.$$

Then any point on the edge they form is universal in the sense just defined.

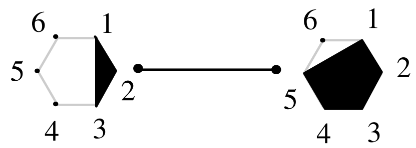

FIGURE 4.2. All points on this edge in X are universal.

Lemma 4.7. *Let $x \in X$ be a universal point, and let $y \in X$ be non opposite to x in B, that is $d_B(x, y) < \pi$. Then the geodesic in B between x and y lies in X. In other words, X is π-star-shaped at x.*

Proof. Let C be a universal chamber with $x \in C$, and let C' be a chamber containing y. We will construct an apartment $A \subseteq X$ containing both.

Let $\{x_1, \ldots, x_{n-2}\}$ denote the vertices of C, with indices corresponding to ranks. We are going to construct a total order on U_n. Note that, seen as partitions, x_{i+1} is obtained from x_i by expanding the unique block with multiple elements (which we will refer to a as the big block of x_i) by an element adjacent to the block. We will call this element the new element of x_{i+1}. We take our order to be one in which an element $v \in U_n$ is larger than u whenever there exists i such that v is new for x_{i+1}, and u belongs to the big block of x_i (we allow $i = n-1$ and set $x_{n-1} = 1$, the partition with one element). Note that there are precisely two such total orders, depending on how we order vertices in the big block of x_1. Note also that given any element $v \in U_n$, we get a non-crossing partition o_v with blocks

$$\{w \in U_n \mid w \leqslant v\} \text{ and } \{w \in U_n \mid w > v\}.$$

Now let $\{y_1, \ldots, y_{n-2}\}$ denote the vertices of C', with indices corresponding to ranks. Note that, seen as partitions, y_{i+1} is obtained from y_i by combining two blocks into one. We are now going to inductively construct embedded forests with vertex set U_n, and edges given by geodesic segments.

We set T_1 to be the forest with vertex set U_n, and a single edge connecting the two elements of the unique non-trivial block of y_1. Suppose we have already defined T_i. Then T_{i+1} is obtained from T_i by adding an edge connecting elements v and w such that

- v and w do not lie in a common block in y_i;
- v and w do lie in a common block in y_{i+1};
- $v > w$;
- v is the minimal element in its block in y_i; and
- the new forest T_{i+1} is embedded.

To show that such a pair v, w exists let us look at minimal vertices in the two blocks of y_i that become one in y_{i+1}. We let v be the larger of the two. Then we know that the block not containing v contains at least one smaller vertex. Together with the fact that o_v defined above is non-crossing, the existence of a suitable w is guaranteed.

Now it is clear that $T = T_{n-1}$ (with $y_{n-1} = 1$, the full partition) is an embedded tree with vertex set U_n. It is also clear that the apartment A defined by T (using Proposition 4.4) contains C'.

Observe that every element (except the minimal one) is connected with an edge to a smaller element. This is due to the fact that every element except the minimal one stops being the smallest element in its block for some i (when we add $y_{n-1} = 1$ to our considerations). When it stops being minimal, it plays the role of v above, and so is connected to a smaller element. From this we easily deduce that $x_i \in A$ for every i, and so that $C \subseteq A$.

Now both points x and y lie in a common apartment A, and the distance between them is smaller than π. Hence there exists a unique geodesic in B between them, and it lies in A. But $A \subseteq X$, so this concludes the proof. □

Lemma 4.8. *Let $x \in X$ be a universal point, and let l be a short loop in X through x. Then l is shrinkable in X.*

Proof. The Arzelà–Ascoli theorem tells us that we can assume without loss of generality that l cannot be shrunk to a shorter loop going through x. Then, since every point in X has a neighbourhood isometric to a metric cone over the point, the loop l is a locally geodesic path in X except possibly at x. We claim that l is constant.

By contradiction, assume that l is not constant, and has length $0 < L < 2\pi$. View l as a path $l : [0, L] \to X$ from x to x. As observed above, l is a local geodesic. If l is not locally geodesic in B, then it has a turning point in $(0, L)$. According to Lemma 3.11, the set of turning points of l is finite. Consider a turning point closest to 0 or L; without loss of generality assume that $0 < t \leqslant \frac{L}{2} < \pi$ is a turning point such that there is no turning point in $(0, t)$. Then $l|_{[0,t]}$ is a locally geodesic segment in B of length smaller than π, hence it is a geodesic segment in B.

Then for $\varepsilon > 0$ small, the geodesic segments $[x, l(t + \alpha\varepsilon)] \subset B$ for $\alpha \in (0, 1]$ lie in X by Lemma 4.7, and are shorter than $l|_{[0,t+\alpha\varepsilon]}$. They also vary continuously with α, since B is CAT(1) (compare Figure 4.3). Therefore l can be shrunk by replacing $l|_{[0,t+\alpha\varepsilon]}$ by $[x, l(t + \alpha\varepsilon)]$, which contradicts the assumption on l.

So l is locally geodesic in B, and therefore $l|_{[0,\frac{L}{2}]}$ and $l^{-1}|_{[\frac{L}{2},L]}$ are two locally geodesic paths in B from x to $l(\frac{L}{2})$ of length smaller than π. They must be equal, since B is CAT(1), and hence l is constant. □

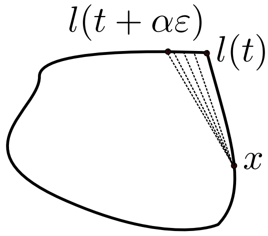

FIGURE 4.3. Illustration of curve shortening in the proof of Lemma 4.8.

Lemma 4.9 (Failing modularity with respect to NCP_n)**.** *Let E, F be two faces of X. The simplicial convex hull of $E \cup F$ in B is contained in X if and only if no two vertices of $E \cup F$ fail modularity with respect to NCP_n.*

Proof. In Lemma 2.21 we have already identified the simplicial convex hull N of $E \cup F$ with the (full subcomplex spanned by) the smallest subset of the vertex set of B stable under taking joins and meets, and containing the vertices of $E \cup F$. If a pair of vertices of $E \cup F$ fails modularity, then we immediately see that N does not lie in X.

Let us now assume that no two vertices of $E \cup F$ fail modularity. This means that for vertices x and y in $E \cup F$ their joins and meets taken in NCP_n agree with the joins and meets taken in $S(V)$.

Let us focus on joins for the moment; the situation for meets is analogous, and the results for joins are easily transferred to results for meets using the duality of NCP_n. The join $x \vee y$ taken in NCP_n equals the one taken in $S(V)$ if and only if they are of the same rank. The rank can be easily read off the block structure of the partition, and thus we immediately see that the two joins agree if and only if whenever the convex hull of a block of x intersects the convex hull of a block of y, then this intersection contains some point of U_n, i.e. no two blocks are crossing.

We claim that $N \subseteq X$, that is that we can perform sequences of meets and joins (in $S(V)$ on vertices of $E \cup F$ and never leave NCP_n. Let us suppose that this is not the case. Without loss of generality we may assume that there are vertices z and w in X, each obtainable from the vertices of $E \cup F$ by a sequence of meets and joins, and such that $z \vee w$ (taken in $S(V)$) does not lie in X. The discussion above tells us that z and w contain crossing blocks. In particular, there exist points $\alpha_1, \alpha_2, \beta_1, \beta_2 \in U_n$ such that $\alpha_1 \neq \alpha_2$ lie in a block of z, $\beta_1 \neq \beta_2$ lie in a block of w, and these two blocks cross.

Since z is obtained from vertices of $E \cup F$ by taking joins and meets, there exists a vertex therein in which α_1 and α_2 are contained in a single block, and this block does not contain both β_1 and β_2 (it might however contain one of them). Without loss of generality let us assume that there exists a vertex of E satisfying this property; let e denote the minimal such vertex in E.

Let us suppose that there exists a vertex in E satisfying the analogous property for β_1 and β_2; let e' denote the minimal such vertex. Now if $e' \leqslant e$ then the block of e containing α_1 and α_2 must also contain β_1 and β_2, since otherwise e is not a non-crossing partition. This is a contradiction. Similarly, when $e < e'$, then e' cannot be as defined. We conclude that there exists a vertex in F such that β_1 and β_2 lie in a common block thereof, and this block does not contain both α_1 and α_2. Let f be the minimal such vertex. Note that f is in fact the minimal vertex of $E \cup F$ satisfying the above property; using an analogous argument we show that e is also the minimal vertex of $E \cup F$ in which α_1 and α_2 lie in a common block, which does not contain both β_1 and β_2.

We now claim that $e \vee f$ (taken in $S(V)$) does not lie in NCP_n. It is enough to find a block in e which crosses a block in f. We already have two candidate blocks, the one containing α_1 and α_2 in e and the one containing β_1 and β_2 in f. It could happen however that these blocks have an element in U_n, say γ, in common. But then, using the minimality of e and f, we conclude that the crossing blocks of z and w also contain γ, and thus are not crossing. This is a contradiction.

We have thus found two vertices in $E \cup F$ which fail modularity with respect to NCP_n. □

Lemma 4.10. *When $n = 5$, turning faces in X are universal vertices.*

Proof. When $n = 5$, by Lemma 3.12 we know that the corank of a turning face contains at least 2 consecutive integers. Since the rank of X is equal to 3, we conclude that F is a vertex of rank either 1 or 3. By Lemmata 3.13 and 4.9, we know that F has two neighbours which fail modularity with respect to NCP_5, hence F is necessarily (by inspection) a universal vertex. □

Corollary 4.11. *The non-crossing partition complex NCP_5 is CAT(0).*

Proof. Assume there is an unshrinkable short loop in X, the diagonal link of NCP_5. Then it has a turning point by Lemma 3.7, which is a universal vertex by Lemma 4.10, so the loop can be shrunk by Lemma 4.8. Hence by Theorem 2.34 (and Remark 4.1), NCP_5 is CAT(0). □

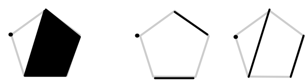

FIGURE 4.4. On the left the turning vertex $\{\{1,2,3,4\},\{5\}\}$ in NCP_5, with two neighbours that fail modularity with respect to NCP_5. Compare Example 4.12

Example 4.12 (Vertices failing modularity with respect to NCP_5)**.** Figure 4.4 pictures the partition $\{\{1,2,3,4\},\{5\}\}$ in NCP_5, with a pair of neighbours given by the partitions $\{\{1,2\},\{3,4\},\{5\}\}$ and $\{\{2,3\},\{1,4\},\{5\}\}$. These neighbours fail modularity with respect to NCP_5 as defined in Definition 2.7. More explicitly this means the following: recall that the lattice of non-crossing partitions can be linearly embedded into a linear lattice $S(V)$. The partitions $\{\{1,2\},\{3,4\},\{5\}\}$ and $\{\{2,3\},\{1,4\},\{5\}\}$ then represent linear subspaces of the underlying vector space V. Their common span in the linear lattice is a linear subspace which cannot be represented (under our fixed embedding $NCP_n \to S(V)$) by a non-crossing partition.

The same vertex with crossing neighbours that fail modularity is illustrated in Figure 4.5 where we also show how these vertices fit into the diagonal link of the partition complex. Pictured are three apartments A_1, A_2 and A_3 of B. The apartments A_1 and A_2 are contained in X and intersect in the gray-shaded region. The geodesic in X connecting u and w runs via v. The apartment A_3 (which does not lie in X) contains u, v and w and of course also the (now shorter) geodesic in the diagonal link of the partition complex connecting u and w. The intersection $A_1 \cap A_3$ is shown in blue while the yellow area highlights $A_2 \cap A_3$.

Definition 4.13 (Dominant vertex)**.** A vertex v of a face F of X is called *dominant* if and only if every apartment A in X, with $v \in A$, contains F.

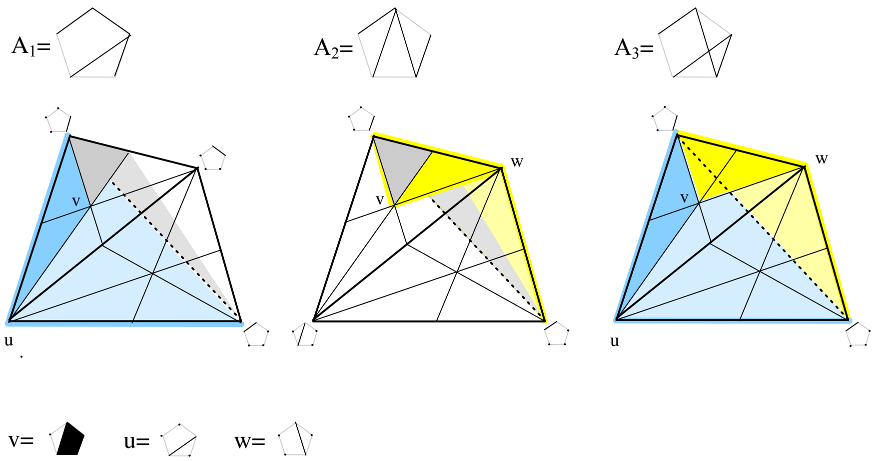

FIGURE 4.5. This figure shows the vertex $v = \{\{1,2,3,4\},\{5\}\}$ in NCP_5 with two neighbours u, w that fail modularity with respect to NCP_5. For more details see Example 4.12.

The following lemma will be accompanied by figures illustrating the cases and subcases. In each case we look at a turning face which can either be one of the two edges illustrated, or a single vertex. Next to each turning edge (subcases (b) and (c)) we depict a pair of adjacent vertices that fail modularity. For the turning vertices (subcases (a)) at least one of the examples next to subcases (b) and (c) gives a pair of adjacent vertices which fail modularity.

The pictured partitions should be read as follows: the vertex in the upper left hand corner will be labeled by 1 and all the other vertices will be labeled clockwise from 2,...,6. Hence the partitions shown in case (1b) are (from left to right): {{1}, {2,3,4}, {5}, {6}}, connected to {{1},{2,4}, {3}, {5}, {6}}, then {{1,5}, {2,3,4}, {6}} and {{1}, {2,3,4,6}, {5}}.

Lemma 4.14. *When $n = 6$, a non-universal turning face F in X, up to the symmetries of U_6, falls into one of the following cases.*

(1a) F is the vertex $v = \{\{2,4\},\{1\},\{3\},\{5\},\{6\}\}$ of rank 1

(1b) F is the following edge of rank $(1,2)$ with dominant vertex v as in (1a).

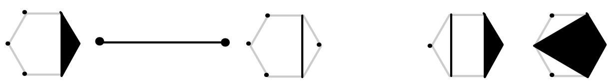

(1c) F is the following edge of rank $(1,4)$ with dominant vertex v as in (1a).

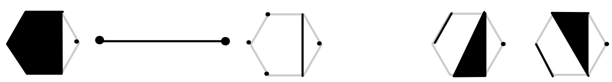

(2a) F is a single vertex of rank 2, namely either $v = \{\{1,2\},\{3,4\},\{5\},\{6\}\}$ or $v' = \{\{1,2\},\{4,5\},\{3\},\{6\}\}$

(2b) F is the following edge of rank $(1,2)$ with dominant vertex v' as in (2a)

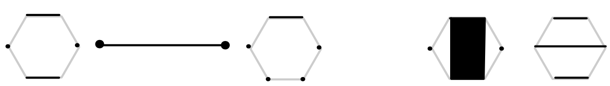

(2c) *F is the following edge of rank* $(1, 2)$ *with dominant vertex* v *as in (2a)*

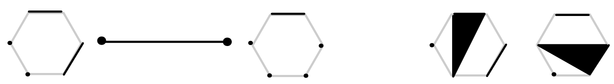

(3a) *F is a single vertex of rank 3, namely either* $v = \{\{1, 2, 3, 5\}, \{4\}, \{6\}\}$ *or* $v' = \{\{1, 2, 4, 5\}, \{3\}, \{6\}\}$

(3b) *F is the following edge of rank* $(3, 4)$ *with dominant vertex* v *as in (3a)*

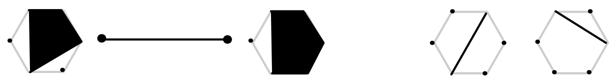

(3c) *F is the following edge of rank* $(3, 4)$ *with dominant vertex* v' *as in (3a)*

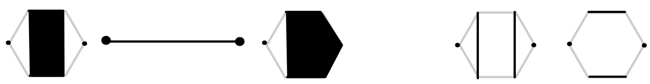

(4a) *F is a single vertex* $v = \{\{1, 2, 3, 4\}, \{5, 6\}\}$ *of rank* 4

(4b) *F is the following edge of rank* $(3, 4)$ *with dominant vertex* v *as in (4a)*

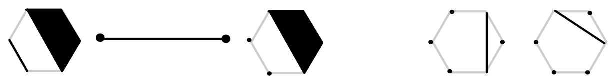

(4c) *F is the following edge of rank* $(1, 4)$ *with dominant vertex* v *as in (4a)*

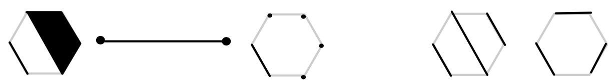

Proof. The proof of this lemma is essentially an inspection.

When $n = 6$, by Lemma 3.12 we know that the corank of a turning face contains at least 2 consecutive integers. Since the rank of X is equal to 4, we conclude that F is either a vertex or an edge. Using duality, we can restrict to the cases where F is a vertex of rank 1 or 2, or an edge of rank $\{1, 2\}$ or $\{1, 4\}$. In these cases it is not hard to see that the cases listed are the only ones which allow for a pair of adjacent vertices failing modularity, in view of Lemmata 3.13 and 4.9. Once we have listed all possible cases, we observe (again by inspection) that each turning face contains a dominant vertex.

Let us look at one example more closely; it is typical in the sense that in all the cases the argument is essentially the same.

Suppose that the turning face is a vertex of rank 1. Then, up to symmetry, it is either the vertex of subcase (1a), that is the partition $\{\{2, 4\}, \{1\}, \{3\}, \{5\}, \{6\}\}$, or the vertex $\{\{2, 5\}, \{1\}, \{3\}, \{4\}, \{6\}\}$. In the former case we can indeed find two vertices adjacent to our vertex which fail modularity. In the latter it is impossible: the lattice interval between that vertex v and the maximal element 1 is isomorphic to the product of two copies of the lattice NCP_3. As a consequence, no two vertices adjacent to v can fail modularity, since the way NCP_3 embeds into $S(V)$ preserves meets and joins. $\square$

When F is a face in X and $i \in U_n$, let $F_i \subseteq U_n$ denote the smallest subset of U_n that appears as a block in a vertex of F and that contains i properly. If the set of such subsets is empty we set $F_i = U_n$.

Lemma 4.15. *Let C be a chamber in X, and let i, j be consecutive elements of U_n. If C_i contains j, then there exists an apartment in X containing C, v and w, where v is the universal vertex having the single nontrivial block $\{i, j\}$ and w is the universal vertex opposite to v in B given by the partition $w = \{\{i\}, [\![1, n]\!] \setminus \{i\}\}$.*

Proof. Write $v_1, \ldots, v_{n-2}$ for the vertices of C, with the indices corresponding to the ranks, and let $v_{n-1} = 1$ be the maximal element in NCP_n. Denote the edge $\{i, j\}$ by e.

Let k be minimal such that C_i is a block of v_k.

Any apartment in X containing both v and w is represented by a non-crossing tree T which contains the edge e, and such that the subforest obtained by removing e from the tree corresponds to w (since w is opposite v), using the correspondence from Remark 4.3. We will now construct such a tree T by inductively picking edges $e_l =$ for $l = 1, \ldots, n-1$ in U_n.

Take e_1 to be the edge corresponding to v_1. For each $2 \leqslant l \leqslant k-1$ choose e_l to be an edge such that the edges $e_1, \ldots, e_l$ form a non-crossing forest corresponding to v_l (again using Remark 4.3). This is possible since vertices of C are non-crossing partitions.

Choose $e_k = e$. Observe that the edges $e_1, \ldots, e_k$ still form a non-crossing forest, since the edge e cannot cross any other edge, and the vertex i was isolated in the forest formed by $e_1, \ldots, e_{k-1}$, and so no cycles appear.

Now we continue choosing edges e_l for $k+1, \ldots, n-1$ as before, with the additional requirement that none of the edges e_l with $l \geqslant k+1$ connects to i. Choosing the remaining edges like this is possible since in each step two blocks of v_l are joined to form a block of v_{l+1}, and the block containing i always contains at least also the vertex j, hence if a block is joined to the one containing i then we may do this using an edge emanating from j (or some other vertex in this block different from i). The resulting tree is by construction non-crossing. The apartment A spanned by T contains v and C. Further, since e_k is the only edge connected to i, the apartment A does also contain the vertex w. □

Lemma 4.16. *When $n = 6$, let F be a non-universal turning face in X, and let C be any chamber in X. Then there exists a pair v, w of universal vertices in X, which are opposite in B, such that F, v, w are contained in an apartment in X, and C, v, w are contained in a (possibly different) apartment in X.*

Proof. According to Lemma 4.14, every non-universal turning face contains a dominant vertex; let u denote the dominant vertex of F. Our strategy here is to find consecutive $i, j \in U_n$ such that $j \in C_i \cap u_i$. Then Lemma 4.15 will give a pair v, w of universal vertices in X, which are opposite in B, such that there exists an apartment in X containing C, v, w, and another apartment in X containing u, v, w. Since u is dominant for F, this last apartment contains F, v, w.

The following table lists all possibilities for u_i (up to duality), depending on i and the dominant vertex u of F (listed as in Lemma 4.14).

Dominant vertex u	$i = 1$	2	3	4	5	6
$\{\{2,4\},\{1\},\{3\},\{5\},\{6\}\}$	U_6	$\{2,4\}$	U_6	$\{2,4\}$	U_6	U_6
$\{\{1,2\},\{3,4\},\{5\},\{6\}\}$	$\{1,2\}$	$\{1,2\}$	$\{3,4\}$	$\{3,4\}$	U_6	U_6
$\{\{1,2\},\{4,5\},\{3\},\{6\}\}$	$\{1,2\}$	$\{1,2\}$	U_6	$\{4,5\}$	$\{4,5\}$	U_6

Now let us consider C_i. If $5 \in C_6$ or $1 \in C_6$ then our table tells us that we are done. Suppose that neither of these two occurs. If $4 \in C_6$ then $4 \in C_5$ and again we are done. Similarly if $2 \in C_6$ then $2 \in C_1$.

We are left with the case $C_6 = \{3, 6\}$. Here $6 \in C_5$ or $4 \in C_5$, which deals with the first two possibilities for u_5. The third one requires the observation that if $4 \notin C_5$ then $5 \in C_4$. □

Theorem 4.17. *The non-crossing partition complex NCP_6 is CAT(0).*

Proof. Assume there is an unshrinkable short loop $l : S^1 \to X$ of length $L < 2\pi$, where X is the diagonal link of $|NCP_6|$. Then by Lemma 3.7 this loop has a turning point with image x in X. Let us reparametrise l so that the domain of l is $[0, L]$ and $x = l(0) = l(L)$. Consider $y = l(L/2)$. By Lemma 4.16, there exists a pair v, w of universal vertices in X, which are opposite in B, such that both $\{x, v, w\}$ and $\{y, v, w\}$ lie in apartments in X. Hence we know that

$$d(x, v) + d(x, w) = d(y, v) + d(y, w) = \pi.$$

So at least one element of $\{v, w\}$, say v, satisfies

$$d(x, v) + d(v, y) \leqslant \pi.$$

Let $\alpha_1 = l|_{[0,L/2]}$ and $\alpha_2 = l|^{-1}_{[L/2,L]}$ be the two subpaths of l from x to y. Let $\alpha_3 \colon [0, d(x, v) + d(v, y)] \to X$ denote the concatenation of the geodesic segments $[x, v]$ and $[v, y]$.

Consider the loop $\alpha_3^{-1} \circ \alpha_1$. Since it is short and passes through the universal vertex v, by Lemma 4.8 it can be shrunk. Similarly, the loop $\alpha_3^{-1} \circ \alpha_2$ can be shrunk. Now we can apply Theorem 2.33, which tells us that the loop $l = \alpha_2^{-1} \circ \alpha_1$ can be shrunk. Hence by Theorem 2.34, the diagonal link in NCP_6 is CAT(1), and the result follows. □

Now we apply Proposition 2.35 to conclude the following.

Corollary 4.18. *For every $n \leqslant 6$, the n-strand braid group is CAT(0).*

5. The orthoscheme complex of a modular complemented lattice is CAT(0)

We now prove that the orthoscheme complex of a bounded graded modular complemented lattice is CAT(0), thus giving a partial confirmation of [**BM2**, Conjecture 6.10]. The conjecture states that the result should be true without assuming that the lattice is complemented, however we need this extra assumption to embed the diagonal link of the orthoscheme complex into a spherical building (or a CAT(1) graph in a pathological case).

Definition 5.1 (Modular lattice)**.** A lattice P is said to be *modular* if

$$\forall x, y, z \in P,\ x \geqslant z \Longrightarrow x \wedge (y \vee z) = (x \wedge y) \vee z.$$

Suppose P is a modular lattice which is linearly embedded in some $S(V)$. Then it is easy to check that joins and meets in P need to coincide with joins, respectively meets in $S(V)$. Hence there is no pair $x, y \in P$ which fails modularity in the sense of Definition 2.7.

Definition 5.2 (Complemented lattice)**.** A bounded lattice P is said to be *complemented* if

$$\forall x \in P, \exists y \in P,\ x \wedge y = 0 \text{ and } x \vee y = 1.$$

Definition 5.3 (Plane lattice)**.** A lattice P is said to be a *plane lattice* if it is bounded, and graded of rank 3.

Theorem 5.4 (Frink's embedding Theorem)**.** *Let P be a bounded graded modular complemented lattice. Then P is isomorphic to a direct product*

$$P = \prod_{i=1}^{r} P_i$$

of bounded graded modular complemented lattices, such that for all $i \in [\![1, r]\!]$, the lattice P_i can be embedded as a subposet of a linear lattice (over a division algebra) or of a plane lattice, where the embedding preserves the meets and the joins.

Note that it is absolutely crucial for us that the embeddings preserve joins and meets. It immediately implies that if the diagonal link X in $|P|$ is linearly embedded, then no two vertices in X fail modularity with respect to P.

Proof. According to [**Grä**, Theorem 279] and [**Grä**, Lemma 99], P is isomorphic to a direct product $P = \prod_{i=1}^{r} P_i$ of simple bounded graded modular complemented lattices, where the fact that each P_i is simple implies that P_i cannot be embedded non-trivially as a subposet of a non-trivial product of lattices. According to [**Grä**, Corollary 439], each P_i is then embedded as a subposet of a product of linear lattices (over a division algebra) or of plane lattices, such that the joins and the meets are preserved. Since P_i is simple, the product consists of only one non-trivial factor. □

For a more precise version of Frink's embedding Theorem, we refer the reader to [**Grä**].

Corollary 5.5. *The diagonal link in the orthoscheme complex of a plane lattice is CAT(1).*

Proof. The diagonal link of the orthoscheme complex of a plane lattice is a graph, since the orthoscheme complex of a rank 3 poset has dimension 3, and so the diagonal link has dimension 1. Moreover, any cycle in the graph is of even length, since the lattice is graded, there are no 2-cycles, since it is a simplicial complex, and no 4-cycles, since the poset is a lattice. Thus the girth of the diagonal link is at least 6. Also, each edge has the same length, namely $\frac{\pi}{3}$. Thus all loops shorter than 2π are shrinkable. The graph is also clearly locally CAT(1). □

Theorem 5.6. *The orthoscheme complex of a bounded graded modular complemented lattice is CAT(0).*

Proof. Let P be a bounded graded modular complemented lattice, and let $|P|$ be its orthoscheme complex. By Theorem 5.4, write $P = \prod_{i=1}^{r} P_i$. Since the orthoscheme complex of $|P|$ is the Euclidean product of the orthoscheme complexes of the posets P_i (thanks to [**BM2**, Remark 5.3]), we only need to show that each $|P_i|$ is CAT(0).

According to [**BM2**, Theorem 5.10], it is enough to check that the diagonal links of the full subcomplexes of $|P_i|$ spanned by intervals in P_i are CAT(1). Since every such interval is itself a bounded graded modular complemented lattice by [**Grä**, Lemma 98], and since this subcomplex is isometric to the orthoscheme complex of the interval, we only need to check this property for the diagonal link of $|P_i|$ itself (formally, we proceed by induction on the rank of the lattice).

Fix $i \in [\![1, r]\!]$. The lattice P_i is embedded as a subposet of a linear lattice $S(V)$ (over a division algebra) or of a plane lattice L, where the embedding preserves the meets and the joins.

In the first case, P_i is linearly embedded in $S(V)$ in the sense of Definition 3.8. Since joins and meets coincide in P_i and in $S(V)$, we deduce that P_i has no pair of elements failing modularity in the sense of Definition 2.7. Let

$$X = LK(e_{01}, |P_i|) \subseteq LK(e_{01}, |S(V)|) = B$$

be their diagonal links. Since B is a spherical building, it is CAT(1). Assume that there is a short locally geodesic loop l in X. By Lemma 3.7, the loop l has a turning point in B. By Lemma 3.13, the image of that turning point has neighbours which fail modularity, which is a contradiction. Hence l cannot exist, and so according to [**BM2**, Theorem 5.10] this implies that $|P_i|$ is CAT(0).

In the second case, the diagonal link of $|P_i|$ is a subgraph of the diagonal link of a plane lattice, which is a CAT(1) graph. Thus so is the diagonal link of $|P_i|$, and thus $|P_i|$ is CAT(0) as before. □

Automorphisms of RAAGs

CHAPTER VII

Outer actions of $\mathrm{Out}(F_n)$ on small RAAGs

Abstract. We determine the precise conditions under which $\mathrm{SOut}(F_n)$, the unique index two subgroup of $\mathrm{Out}(F_n)$, can act non-trivially via outer automorphisms on a RAAG whose defining graph has fewer than $\frac{1}{2}\binom{n}{2}$ vertices.
We also show that the outer automorphism group of a RAAG cannot act faithfully via outer automorphisms on a RAAG with a strictly smaller (in number of vertices) defining graph.
Along the way we determine the minimal dimensions of non-trivial linear representations of congruence quotients of the integral special linear groups over algebraically closed fields of characteristic zero, and provide a new lower bound on the cardinality of a set on which $\mathrm{SOut}(F_n)$ can act non-trivially.

1. Introduction

The main purpose of this article is to study the ways in which $\mathrm{Out}(F_n)$ can act via outer automorphisms on a right-angled Artin group A_Γ with defining graph Γ. Such actions have previously been studied for the extremal cases: when the graph Γ is discrete, we have $\mathrm{Out}(A_\Gamma) = \mathrm{Out}(F_m)$ for some m, and homomorphisms

$$\mathrm{Out}(F_n) \to \mathrm{Out}(F_m)$$

have been investigated by Bogopolski–Puga [**BP**], Khramtsov [**Khr2**], Bridson–Vogtmann [**BV3**], and the author [**Kie1, Kie3**]. When the graph Γ is complete, we have $\mathrm{Out}(A_\Gamma) = \mathrm{GL}_m(\mathbb{Z})$, and homomorphisms

$$\mathrm{Out}(F_n) \to \mathrm{GL}_m(\mathbb{Z})$$

or more general representation theory of $\mathrm{Out}(F_n)$ have been studied by Grunewald–Lubotzky [**GL**], Potapchik–Rapinchuk [**PR**], Turchin–Wilwacher [**TW**], and the author [**Kie1, Kie3**].

There are two natural ways of constructing non-trivial homomorphisms

$$\varphi\colon \mathrm{Out}(F_n) \to \mathrm{Out}(A_\Gamma)$$

When Γ is a join of two graphs, Δ and Σ say, then $\mathrm{Out}(A_\Gamma)$ contains

$$\mathrm{Out}(A_\Delta) \times \mathrm{Out}(A_\Sigma)$$

as a finite index subgroup. When additionally Δ is isomorphic to the discrete graph with n vertices, then $\mathrm{Out}(A_\Delta) = \mathrm{Out}(F_n)$, and so we have an obvious embedding φ.

In fact this method works also for a discrete Δ with a very large number of vertices, since there are injective maps $\mathrm{Out}(F_n) \to \mathrm{Out}(F_m)$ constructed by Bridson–Vogtmann [**BV3**] for specific values of m growing exponentially with n.

The other way of constructing non-trivial homomorphisms φ becomes possible when Γ contains n vertices with identical stars. In this case it is immediate that these vertices form a clique Θ, and we have a map

$$\mathrm{GL}_n(\mathbb{Z}) = \mathrm{Aut}(A_\Theta) \to \mathrm{Aut}(A_\Gamma) \to \mathrm{Out}(A_\Gamma)$$

We also have the projection

$$\mathrm{Out}(F_n) \to \mathrm{Out}(H_1(F_n)) = \mathrm{GL}_n(\mathbb{Z})$$

and combining these two maps gives us a non-trivial (though also non-injective) φ.

This second method does not work in other situations, due to the following result of Wade.

Theorem 1.1 ([**Wad**]). *Let $n \geqslant 3$. Every homomorphism*

$$\mathrm{SL}_n(\mathbb{Z}) \to \mathrm{Out}(A_\Gamma)$$

has finite image if and only if Γ does not contain n distinct vertices with equal stars.

In fact Wade proved a much more general result, in which the domain of the homomorphism is allowed to be any irreducible lattice in a real semisimple Lie group with finite centre and without compact factors, and with real rank $n-1$.

The aim of this paper is to prove

Theorem 3.7. *Let $n \geqslant 6$. Suppose that Γ is a simplicial graph with fewer than $\frac{1}{2}\binom{n}{2}$ vertices, which does not contain n distinct vertices with equal stars, and is not a join of the discrete graph with n vertices and another (possibly empty) graph. Then every homomorphism $\mathrm{SOut}(F_n) \to \mathrm{Out}(A_\Gamma)$ is trivial.*

Here $\mathrm{SOut}(F_n)$ denotes the unique index two subgroup of $\mathrm{Out}(F_n)$.

The proof is an induction, based on an observation present in a paper of Charney–Crisp–Vogtmann [**CCV**], elaborated further in Chapter VIII, which states that, typically, the graph Γ contains many induced subgraphs Σ which are *invariant up to symmetry*, in the sense that the subgroup of A_Γ the vertices of Σ generate is invariant under any outer action up to an automorphism induced by a symmetry of Γ (and up to conjugacy).

To use the induction we need to show that such subgraphs are really invariant, that is that we do not need to worry about the symmetries of Γ. To achieve this we prove

Theorem 2.28. *Every action of $\mathrm{Out}(F_n)$ (with $n \geqslant 6$) on a set of cardinality $m \leqslant \binom{n+1}{2}$ factors through $\mathbb{Z}/2\mathbb{Z}$.*

Since $\mathrm{SOut}(F_n)$ is the unique index two subgroup of $\mathrm{Out}(F_n)$, the conclusion of this theorem is equivalent to saying that $\mathrm{SOut}(F_n)$ lies in the kernel of the action.

A crucial ingredient in the proof of this theorem is the following.

Theorem 2.27. *Let V be a non-trivial, irreducible $\mathbb{K}$-linear representation of*

$$\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$$

where $n \geqslant 3$, q is a power of a prime p, and where $\mathbb{K}$ is an algebraically closed field of characteristic 0. Then

$$\dim V \geqslant \begin{cases} 2 & \text{if} \quad (n,p) = (3,2) \\ p^{n-1} - 1 & \text{otherwise} \end{cases}$$

This result seems not to be present in the literature; it extends a theorem of Landazuri–Seitz [**LS**] yielding a very similar statement for $q = p$ (see Theorem 2.26).

At the end of the paper we also offer

Theorem 4.1. *There are no injective homomorphisms $\mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_{\Gamma'})$ when Γ' has fewer vertices than Γ.*

This theorem follows from looking at the $\mathbb{Z}/2\mathbb{Z}$-rank, i.e. the largest subgroup isomorphic to $(\mathbb{Z}/2\mathbb{Z})^k$.

2. The tools

2.1. Automorphisms of free groups.

Definition 2.1 ($\mathrm{SOut}(F_n)$)**.** Consider the composition

$$\mathrm{Aut}(F_n) \to \mathrm{GL}_n(\mathbb{Z}) \to \mathbb{Z}/2\mathbb{Z}$$

where the first map is obtained by abelianising F_n, and the second map is the determinant. We define $\mathrm{SAut}(F_n)$ to be the kernel of this map; we define $\mathrm{SOut}(F_n)$ to be the image of $\mathrm{SAut}(F_n)$ in $\mathrm{Out}(F_n)$.

It is easy to see that both $\mathrm{SAut}(F_n)$ and $\mathrm{SOut}(F_n)$ are index two subgroups of, respectively, $\mathrm{Aut}(F_n)$ and $\mathrm{Out}(F_n)$.

The group $\mathrm{SAut}(F_n)$ has a finite presentation given by Gersten [**Ger**], and from this presentation one can immediately obtain the following result.

Proposition 2.2 (Gersten [**Ger**])**.** *The abelianisation of* $\mathrm{SAut}(F_n)$*, and hence of* $\mathrm{SOut}(F_n)$*, is trivial for all* $n \geqslant 3$*.*

It follows that $\mathrm{SOut}(F_n)$ is the unique subgroup of $\mathrm{Out}(F_n)$ of index two.

We will now look at symmetric and alternating subgroups of $\mathrm{Out}(F_n)$, and list some corollaries of their existence.

Proposition 2.3 ([**BV1**, Proposition 1])**.** *Let* $n \geqslant 3$*. There exists a symmetric subgroup of rank* n

$$\mathrm{Sym}_n < \mathrm{Out}(F_n)$$

such that any homomorphism $\varphi\colon \mathrm{Out}(F_n) \to G$ *that is not injective on* Sym_n *has image of cardinality at most 2.*

The symmetric group is precisely the symmetric group operating on some fixed basis of F_n. It is easy to see that it intersects $\mathrm{SOut}(F_n)$ in an alternating group Alt_n. Whenever we talk about the alternating subgroup Alt_n of $\mathrm{SOut}(F_n)$, we mean this subgroup. Note that $\mathrm{SOut}(F_n)$ actually contains an alternating subgroup of rank $n+1$, which is a supergroup of our Alt_n; we will denote it by Alt_{n+1}. There is also a symmetric supergroup Sym_{n+1} of Alt_{n+1} contained in $\mathrm{Out}(F_n)$.

The proof of [**BV1**, Proposition 1] actually allows one to prove the following proposition.

Proposition 2.4. *Let* $n \geqslant 3$*. Then* $\mathrm{SOut}(F_n)$ *is the normal closure of any non-trivial element of* Alt_n*.*

Following the proof of [**BV1**, Theorem A], we can now conclude

Corollary 2.5. *Let*

$$\varphi\colon \mathrm{SOut}(F_n) \to \mathrm{GL}_k(\mathbb{Z})$$

be a homomorphism, with $n \geqslant 6$ *and* $k < n$*. Then* φ *is trivial.*

Proof. For $n \geqslant 6$, the alternating group Alt_{n+1} does not have non-trivial complex representations below dimension n. Thus $\varphi|_{\mathrm{Alt}_{n+1}}$ is not injective, and therefore trivial, as Alt_{n+1} is simple. Now we apply Proposition 2.4. □

More can be said about linear representations of $\mathrm{Out}(F_n)$ in somewhat larger dimensions – see [**Kie1, Kie3, TW**].

Another related result that we will use is the following.

Theorem 2.6 ([**Kie1**])**.** *Let* $n \geqslant 6$ *and* $m < \binom{n}{2}$*. Then every homomorphism* $\mathrm{Out}(F_n) \to \mathrm{Out}(F_m)$ *has image of cardinality at most 2, provided that* $m \neq n$*.*

In fact, we will need to go back to the proof of the above theorem and show:

Theorem 2.7. *Let $n \geqslant 6$ and $m < \frac{1}{2}\binom{n}{2}$. Then every homomorphism*

$$\mathrm{SOut}(F_n) \to \mathrm{Out}(F_m)$$

is trivial, provided that $m \neq n$.

The proof of this result forms the content of the next section.

2.2. Homomorphisms $\mathrm{SOut}(F_n) \to \mathrm{Out}(F_m)$. To study such homomorphisms we need to introduce finite subgroups B_n and B of $\mathrm{SOut}(F_n)$ that will be of particular use. Let F_n be freely generated by $\{a_1, \dots, a_n\}$.

Definition 2.8. Let us define $\delta \in \mathrm{Out}(F_n)$ by $\delta(a_i) = {a_i}^{-1}$ for each i. (Formally speaking, this defines an element in $\mathrm{Aut}(F_n)$; we take δ to be the image of this element in $\mathrm{Out}(F_n)$.) Define $\sigma_{12} \in \mathrm{Sym}_n < \mathrm{Out}(F_n)$ to be the transposition swapping a_1 with a_2. Define $\xi \in \mathrm{SOut}(F_n)$ by

$$\xi(a_i) = \begin{cases} \delta & \text{if } n \text{ is even} \\ \delta\sigma_{12} & \text{if } n \text{ is odd} \end{cases}$$

and set $B_n = \langle \mathrm{Alt}_{n+1}, \xi \rangle \leqslant \mathrm{SOut}(F_n)$.

We also set A to be either Alt_{n-1}, the pointwise stabiliser of $\{1, 2\}$ when Alt_{n+1} acts on $\{1, 2, \dots, n+1\}$ in the natural way (in the case of odd n), or Alt_{n+1} (in the case of even n). Furthermore, we set $B = \langle A, \xi \rangle$.

It is easy to see that B_n is a finite group – it is a subgroup of the automorphism group of the (suitably marked) $(n+1)$-*cage graph*, that is a graph with 2 vertices and $n+1$ edges connecting one to another.

To prove Theorem 2.7 we need to introduce some more notation from [**Kie1**]. Throughout, when we talk about modules or representations, we work over the complex numbers.

Definition 2.9. A B-module V admits a *convenient split* if and only if V splits as a B-module into

$$V = U \oplus U'$$

where U is a sum of trivial A-modules and ξ acts as minus the identity on U'.

Definition 2.10. A graph X with a G-action is called G-*admissible* if and only if it is connected, has no vertices of valence 2, and any G-invariant forest in X contains no edges. Here by 'invariant' we mean setwise invariant.

Proposition 2.11 ([**Kie1**])**.** *Let $n \geqslant 6$. Suppose that X is a B_n-admissible graph of rank smaller than $\binom{n+1}{2}$ such that*

(1) the B-module $H_1(X; \mathbb{C})$ admits a convenient split; and
(2) any vector in $H_1(X; \mathbb{C})$ which is fixed by Alt_{n+1} is also fixed by ξ; and
(3) the action of B_n on X restricted to A is non-trivial.

Then X is the $(n+1)$-cage.

The above proposition does not (unfortunately) feature in this form in [**Kie1**] – it does however follow from the proof of [**Kie1**, Proposition 6.7].

Proof of Theorem 2.7. Let $\varphi \colon \mathrm{SOut}(F_n) \to \mathrm{Out}(F_m)$ be a homomorphism. Using Nielsen realisation for free groups (due to, independently, Culler [**Cul**], Khramtsov [**Khr1**] and Zimmermann [**Zim1**]) we construct a finite connected graph X with fundamental group F_m, on which B_n acts in a way realising the outer action $\varphi|_{B_n}$. We easily arrange for X to be B_n-admissible by collapsing invariant forests. Note that $V = H_1(F_m; \mathbb{C})$ is naturally isomorphic to $H_1(X; \mathbb{C})$ as a B_n-module.

We have a linear representation

$$\mathrm{SOut}(F_n) \to \mathrm{Out}(F_m) \to \mathrm{GL}(V)$$

where the first map is φ. We can induce it to a linear representation

$$\mathrm{Out}(F_n) \to \mathrm{GL}(W)$$

of dimension $\dim W = 2 \dim V = 2m$. Since we are assuming that

$$m < \frac{1}{2}\binom{n}{2}$$

the combination of [**Kie1**, Lemma 3.8 and Proposition 3.11] tells us that W splits as an $\mathrm{Out}(F_n)$-module as

$$W = W_0 \oplus W_1 \oplus W_{n-1} \oplus W_n$$

where the action of $\mathrm{Out}(F_n)$ is trivial on W_0 but not on W_n, and the action of the subgroup $\mathrm{SOut}(F_n)$ is trivial on both. Moreover, as Sym_{n+1} modules, W_1 is the sum of standard and W_{n-1} of signed standard representations. We also know that δ acts on W_i as multiplication by $(-1)^i$.

When n is even this immediately tells us that, as a $B = B_n$-module, we have

$$W = U \oplus U'$$

where $U = W_0 \oplus W_n$ is sum of trivial $A = \mathrm{Alt}_{n+1}$-modules, and $\xi = \delta$ acts on

$$U' = W_1 \oplus W_{n-1}$$

as minus the identity.

When n is odd we can still write

$$W = U \oplus U'$$

as a B-module, with A acting trivially on U and ξ acting as minus the identity on U'. Here we have $W_0 \oplus W_n < U$, but U also contains the trivial A-modules contained in $W_1 \oplus W_{n-1}$. The module U' is the sum of the standard A-modules. Thus W admits a convenient split.

Now we claim that V also admits a convenient split as a B-module. To define the induced $\mathrm{Out}(F_n)$-module W we need to pick en element $\mathrm{Out}(F_n) \smallsetminus \mathrm{SOut}(F_n)$; we have already defined such an element, namely σ_{12}. The involution σ_{12} commutes with ξ and conjugates A to itself. Thus, as an A module, V could only consist of the trivial and standard representations, since these are the only A-modules present in W. Moreover, any trivial A-module in V is still a trivial A-module in W, and so ξ acts as minus the identity on it. Therefore V also admits a convenient split as a B-module. This way we have verified assumption (1) of Proposition 2.11.

Observe that the $\mathrm{SOut}(F_n)$-module V embeds into W. In W every Alt_{n+1}-fixed vector lies in $W_0 \oplus W_n$, and here ξ acts as the identity. Thus assumption (2) of Proposition 2.11 is satisfied in W, and therefore also in V.

We have verified the assumptions (1) and (2) of Proposition 2.11; we also know that the conclusion of Proposition 2.11 fails, since the $n+1$-cage has rank n, which would force $m = n$, contradicting the hypothesis of the theorem. Hence we know that assumption (3) of Proposition 2.11 fails, and so A acts trivially on X. But this implies that $A \leqslant \ker \varphi$.

Note that A is a subgroup of the simple group Alt_{n+1}, and so we have

$$\mathrm{Alt}_{n+1} \leqslant \ker \varphi$$

But then Proposition 2.4 tells us that φ is trivial. $\square$

2.3. Automorphisms of RAAGs. Throughout the paper, Γ will be a simplicial graph, and A_Γ will be the associated RAAG, that is the group generated by the vertices of Γ, with a relation of two vertices commuting if and only if they are joined by an edge in Γ.

We will often look at subgraphs of Γ, and we always take them to be induced subgraphs. Thus we will make no distinction between a subgraph of Γ and a subset of the vertex set of Γ.

Given an induced subgraph $\Sigma \subseteq \Gamma$ we define A_Σ to be the subgroup of A_Γ generated by (the vertices of) Σ. Abstractly, A_Σ is isomorphic to the RAAG associated to Σ (since Σ is an induced subgraph).

Definition 2.12 (Links, stars, and extended stars)**.** Given a subgraph $\Sigma \subseteq \Gamma$ we define

- $\mathrm{lk}(\Sigma) = \{w \in \Gamma \mid w \text{ is adjacent to } v \text{ for all } v \in \Sigma\}$;
- $\mathrm{st}(\Sigma) = \Sigma \cup \mathrm{lk}(\Sigma)$;
- $\widehat{\mathrm{st}}(\Sigma) = \mathrm{lk}(\Sigma) \cup \mathrm{lk}(\mathrm{lk}(\Sigma))$.

Definition 2.13 (Joins and cones)**.** We say that two subgraphs $\Sigma, \Delta \subseteq \Gamma$ form a *join* $\Sigma * \Delta \subseteq \Gamma$ if and only if $\Sigma \subseteq \mathrm{lk}(\Delta)$ and $\Delta \subseteq \mathrm{lk}(\Sigma)$.

A subgraph $\Sigma \subseteq \Gamma$ is a *cone* if and only if there exists a vertex $v \in \Sigma$ such that $\Sigma = v * (\Sigma \smallsetminus \{v\})$. In particular, a singleton is a cone.

Definition 2.14 (Join decomposition)**.** Given a graph Σ we say that

$$\Sigma = \Sigma_1 * \cdots * \Sigma_k$$

is the *join decomposition* of Σ when each Σ_i is non-empty, and is not a join of two non-empty subgraphs.

Each of the graphs Σ_i is called a *factor*, and the join of all the factors which are singletons is called the *clique factor*.

We will often focus on a specific finite index subgroup $\mathrm{Out}^0(A_\Gamma)$ of $\mathrm{Out}(A_\Gamma)$, called the *group of pure outer automorphisms of* A_Γ. To define it we need to discuss a generating set of $\mathrm{Out}(A_\Gamma)$ due to Laurence [**Lau**] (it was earlier conjectured to be a generating set by Servatius [**Ser2**]).

$\mathrm{Aut}(A_\Gamma)$ is generated by the following classes of automorphisms:

(1) Inversions
(2) Partial conjugations
(3) Transvections
(4) Graph symmetries

Here, an *inversion* maps one generator of A_Γ to its inverse, fixing all other generators.

A *partial conjugation* needs a vertex v; it conjugates all generators in one connected component of $\Gamma \smallsetminus \mathrm{st}(v)$ by v, and fixes all other generators.

A *transvection* requires vertices v, w with $\mathrm{st}(v) \supseteq \mathrm{lk}(w)$. For such v and w, a transvection is the automorphism which maps w to wv, and fixes all other generators.

A *graph symmetry* is an automorphism of A_Γ which permutes the generators according to a combinatorial automorphism of Γ.

The group $\mathrm{Aut}^0(A_\Gamma)$ of pure automorphisms is defined to be the subgroup generated by generators of the first three types, i.e. without graph symmetries. The group $\mathrm{Out}^0(A_\Gamma)$ of pure outer automorphisms is the quotient of $\mathrm{Aut}^0(A_\Gamma)$ by the inner automorphisms.

Let us quote the following result of Charney–Crisp–Vogtmann:

Proposition 2.15 ([**CCV**, Corollary 3.3])**.** *There exists a finite subgroup*

$$Q < \operatorname{Out}(A_\Gamma)$$

consisting solely of graph symmetries, such that

$$\operatorname{Out}(A_\Gamma) = \operatorname{Out}^0(A_\Gamma) \rtimes Q$$

Corollary 2.16. *Suppose that any action of G on a set of cardinality at most k is trivial, and assume that Γ has k vertices. Then any homomorphism*

$$\varphi \colon G \to \operatorname{Out}(A_\Gamma)$$

has image contained in $\operatorname{Out}^0(A_\Gamma)$*.*

Proof. Proposition 2.15 tells us that

$$\operatorname{Out}(A_\Gamma) = \operatorname{Out}^0(A_\Gamma) \rtimes Q$$

for some group Q acting faithfully on Γ. Hence we can postcompose φ with the quotient map

$$\operatorname{Out}^0(A_\Gamma) \rtimes Q \to Q$$

and obtain an action of G on the set of vertices of Γ. By assumption this action has to be trivial, and thus $\varphi(G)$ lies in the kernel of this quotient map, which is $\operatorname{Out}^0(A_\Gamma)$. □

Definition 2.17 (G-invariant subgraphs)**.** Given a homomorphism $G \to \operatorname{Out}(A_\Gamma)$ we say that a subgraph $\Sigma \subseteq \Gamma$ is *G-invariant* if and only if the conjugacy class of A_Σ is preserved (setwise) by G.

Definition 2.18. Having an invariant subgraph $\Sigma \subseteq \Gamma$ allows us to discuss two additional actions:

- Since, for any subgraph Σ, the normaliser of A_Σ in A_Γ is equal to $A_\Sigma C(A_\Sigma)$, where $C(A_\Sigma)$ is the centraliser of A_Σ (see e.g. [**CSV**, Proposition 2.2]), any invariant subgraph Σ gives us an *induced (outer) action* $G \to \operatorname{Out}(A_\Sigma)$.
- When Σ is invariant, we also have the *induced quotient action*

$$G \to \operatorname{Out}(A_\Gamma / \langle\langle A_\Sigma \rangle\rangle) \simeq \operatorname{Out}(A_{\Gamma \smallsetminus \Sigma})$$

Let us quote the following.

Lemma 2.19 ([**HK1**, Lemmata 4.2 and 4.3])**.** *For any homomorphism $G \to \operatorname{Out}^0(A_\Gamma)$ we have:*

1. *for every subgraph $\Sigma \subseteq \Gamma$ which is not a cone, $\operatorname{lk}(\Sigma)$ is G-invariant;*
2. *connected components of Γ which are not singletons are G-invariant;*
3. *$\widehat{\operatorname{st}}(\Sigma)$ is G-invariant for every subgraph Σ;*
4. *if Σ and Δ are G-invariant, then so is $\Sigma \cap \Delta$;*
5. *if Σ is G-invariant, then so is $\operatorname{st}(\Sigma)$.*

Definition 2.20 (Trivialised subgraphs)**.** Let $\varphi \colon G \to \operatorname{Out}(A_\Gamma)$ be given. We say that a subgraph $\Sigma \subseteq \Gamma$ is *trivialised* if and only if Σ is G-invariant, and the induced action is trivial.

Lemma 2.21. *Let $\varphi \colon G \to \operatorname{Out}(A_\Gamma)$ be a homomorphism. Suppose that Σ is a connected component of Γ which is trivialised by G. Consider the graph*

$$\Gamma' = (\Gamma \smallsetminus \Sigma) \sqcup \{s\}$$

were s denotes a new vertex not present in Γ. There exists an action

$$\psi \colon G \to \operatorname{Out}(A_{\Gamma'})$$

for which $\{s\}$ is invariant, and such that the quotient actions

$$G \to \operatorname{Out}(A_{\Gamma \smallsetminus \Sigma})$$

induced by φ and ψ by removing, respectively, Σ and s, coincide.

Proof. Consider an epimorphism $f\colon A_\Gamma \to A_{\Gamma'}$ defined on vertices of Γ by

$$f(v) = \begin{cases} v & \text{if} \quad v \notin \Sigma \\ s & \text{if} \quad v \in \Sigma \end{cases}$$

The kernel of f is normally generated by elements vu^{-1}, where $v, u \in \Sigma$ are vertices. Since the induced action of G on A_Σ is trivialised, the action preserves each element vu^{-1} up to conjugacy. But this in particular means that G preserves the (conjugacy class of) the kernel of f, and hence φ induces an action

$$G \to \mathrm{Out}(A_{\Gamma'})$$

which we call ψ. It is now immediate that ψ is as required. □

2.4. Finite groups acting on RAAGs.

Definition 2.22. Suppose that Γ has k vertices. Then the abelianisation of A_Γ is isomorphic to $\mathbb{Z}^k$, and we have the natural map

$$\mathrm{Out}(A_\Gamma) \to \mathrm{Out}(H_1(A_\Gamma)) = \mathrm{GL}_k(\mathbb{Z})$$

We will refer to the kernel of this map as the *Torelli subgroup.*

We will need the following consequence of independent (and more general) results of Toinet [**Toi**] and Wade [**Wad**].

Theorem 2.23 (Toinet [**Toi**]; Wade [**Wad**])**.** *The Torelli group is torsion free.*

Lemma 2.24. *Let $\varphi\colon H \to \mathrm{Out}(A_\Gamma)$ be a homomorphism with a finite domain. Suppose that $\Gamma = \Sigma_1 \cup \cdots \cup \Sigma_m$, and each Σ_i is trivialised by H. Then so is Γ.*

Proof. Consider the action

$$\psi\colon H \to \mathrm{Out}(H_1(A_\Gamma)) = \mathrm{GL}_k(\mathbb{Z})$$

obtained by abelianising A_Γ, where k is the number of vertices of Γ. This $\mathbb{Z}$-linear representation ψ preserves the images of the subgroups A_{Σ_i}, and is trivial on each of them. Thus the representation is trivial, and so $\varphi(H)$ lies in the Torelli group. But the Torelli subgroup is torsion free. Hence φ is trivial. □

Lemma 2.25. *Let $\varphi\colon G \to \mathrm{Out}(A_\Gamma)$ be a homomorphism. Let*

$$\Gamma = (\Gamma_1 \cup \cdots \cup \Gamma_n) \sqcup \Theta$$

where $n \geqslant 1$, each Γ_i is trivialised by G, and where Θ is a discrete graph with m vertices. Suppose that for some $l \in \{m, m+1\}$ any homomorphism

$$G \to \mathrm{Out}(F_l)$$

is trivial. Then Γ is trivialised, provided that G is the normal closure of a finite subgroup H, and that G contains a perfect subgroup P, which in turn contains H.

Proof. We can quotient out all of the groups A_{Γ_i}, and obtain an induced quotient action

$$G \to \mathrm{Out}(A_\Theta) \tag{$*$}$$

We claim that this map is trivial. To prove the claim we have to consider two cases: the first case occurs when $l = m$ in the hypothesis of our lemma, that is every homomorphism

$$G \to \mathrm{Out}(F_m)$$

is trivial. Since Θ is a discrete graph with m vertices, we have $\mathrm{Out}(A_\Theta) = \mathrm{Out}(F_m)$ and so the homomorphism $(*)$ is trivial.

The second case occurs when $l = m+1$ in the hypothesis of our lemma. In this situation we quotient A_Γ by each subgroup A_{Γ_i} for $i > 1$, but instead of quotienting out A_{Γ_1}, we use Lemma 2.21. This way we obtain an outer action on a free group with $m+1$ generators, and such an action has to be trivial by assumption. Thus we can take a further quotient and conclude again that the induced quotient action $(*)$ on A_Θ is trivial. This proves the claim.

Now consider the action of G on the abelianisation of A_Γ. We obtain a map

$$\psi\colon G \to \mathrm{GL}_k(\mathbb{Z})$$

where k is the number of vertices of Γ. Since each Γ_i is trivialised, and the induced quotient action on A_Θ is trivial, we see that $\psi(G)$ lies in the abelian subgroup of $\mathrm{GL}_n(\mathbb{Z})$ formed by block-upper triangular matrices with identity blocks on the diagonal, and a single non-trivial block of fixed size above the diagonal. But P is perfect, and so $\psi(P)$ must lie in the Torelli subgroup of $\mathrm{Out}(A_\Gamma)$. This is however torsion free by Theorem 2.23, and so H must in fact lie in the kernel of φ. We conclude that the action of G on Γ is also trivial, since G is the normal closure of H. □

2.5. Some representation theory. Let us mention a result about representations of $\mathrm{PSL}_n(\mathbb{Z}/p\mathbb{Z})$, for prime p, due to Landazuri and Seitz:

Theorem 2.26 ([**LS**])**.** *Suppose that we have a non-trivial, irreducible projective representation* $\mathrm{PSL}_n(\mathbb{Z}/p\mathbb{Z}) \to \mathrm{PGL}(V)$*, where* $n \geqslant 3$*,* p *is prime, and* V *is a vector space over a field* $\mathbb{K}$ *of characteristic other than* p*. Then*

$$\dim V \geqslant \begin{cases} 2 & \text{if} \quad (n,p) = (3,2) \\ p^{n-1} - 1 & \text{otherwise} \end{cases}$$

We offer an extension of their theorem for algebraically closed fields of characteristic 0, which we will need to discuss actions of $\mathrm{Out}(F_n)$ and $\mathrm{SOut}(F_n)$ on finite sets.

Theorem 2.27. *Let* V *be a non-trivial, irreducible* $\mathbb{K}$*-linear representation of* $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$*, where* $n \geqslant 3$*,* q *is a power of a prime* p*, and where* $\mathbb{K}$ *is an algebraically closed field of characteristic 0. Then*

$$\dim V \geqslant \begin{cases} 2 & \text{if} \quad (n,p) = (3,2) \\ p^{n-1} - 1 & \text{otherwise} \end{cases}$$

Proof. Let $\varphi\colon \mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z}) \to \mathrm{GL}(V)$ denote our representation. Consider Z, the subgroup of $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$ generated by diagonal matrices with all non-zero entries equal. Note that Z is the centre of $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$. Hence V splits as an $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$-module into intersections of eigenspaces of all elements of Z. Since V is irreducible, we conclude that $\varphi(Z)$ lies in the centre of $\mathrm{GL}(V)$.

First suppose that $q = p$. Consider the composition

$$\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z}) \to \mathrm{GL}(V) \to \mathrm{PGL}(V)$$

We have just showed that Z lies in the kernel of this composition, and so our representation descends to a representation of $\mathrm{PSL}_n(\mathbb{Z}/p\mathbb{Z}) \cong \mathrm{SL}_n(\mathbb{Z}/p\mathbb{Z})/Z$. This new, projective representation is still irreducible. It is also non-trivial, as otherwise V would have to be a 1-dimensional non-trivial $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$-representation. There are no such representations since $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$ is perfect when $p = q$. Now Theorem 2.26 yields the result.

Suppose now that $q = p^\alpha$, where $\alpha > 1$. Let $N \trianglelefteq \mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$ be the kernel of the natural map $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z}) \to \mathrm{SL}_n(\mathbb{Z}/p\mathbb{Z})$. As an N-module, by Maschke's Theorem,

V splits as

$$V = \bigoplus_{i=1}^{k} U_i$$

where each $U_i \neq \{0\}$ is a direct sum of irreducible N-modules, and irreducible submodules $W \leqslant U_i, W' \leqslant U_j$ are isomorphic if and only if $i = j$.

Observe that we get an induced action of $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})/N \cong \mathrm{SL}_n(\mathbb{Z}/p\mathbb{Z})$ on the set $\{U_i, U_2, \ldots, U_k\}$. As V is an irreducible $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$-module, the action is transitive.

Note that an action of a group on a finite set S induces a representation on the vector space with basis S. If $k > 1$ then this representation is not the sum of trivial ones, because of the transitivity just described, and so

$$k \geqslant \begin{cases} 2 & \text{if} \quad (n,p) = (3,2) \\ p^{n-1} - 1 & \text{otherwise} \end{cases}$$

since our theorem holds for $\mathrm{SL}_n(\mathbb{Z}/p\mathbb{Z})$. Since $\dim U_i \geqslant 1$ for all i, we get $\dim V \geqslant k$ and our result follows.

Let us henceforth assume that $k = 1$. We have

$$V = U_1 = \bigoplus_{j=1}^{l} W$$

where W is an irreducible N-module.

Note that we have an alternating group $\mathrm{Alt}_n < \mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$ satisfying

$$\mathrm{Alt}_n \cap N = \{1\}$$

Let $\sigma \in \mathrm{Alt}_n$ be an element of order $o(\sigma)$ equal to 2 or 3.

Consider the group $M = \langle N, \sigma \rangle < \mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$. Note that $M \cong N \rtimes \mathbb{Z}_{o(\sigma)}$. The module V splits as a direct sum of irreducible M-modules by Maschke's theorem. Let X be such an irreducible M-module.

Note that X as an N-module is a direct sum of, say, m copies of the N-module W (with $m \geqslant 1$). Frobenius Reciprocity (see e.g. [**Wei**, Corollary 4.1.17]) tells us that the multiplicity m of W (as an N-module) in X is equal to the multiplicity of the M-module X in the M-module induced from the N-module W. Hence the multiplicity of W in the M-module induced from the N-module W is at least m^2. But it is bounded above by $o(\sigma)$ and $o(\sigma) \leqslant 3$, which forces $m = 1$, as $m \geqslant 1$.

This shows in particular that X as an N-module is isomorphic to W. It also shows that the M-module induced from W contains a submodule isomorphic to X. Since

$$M \cong N \rtimes \mathbb{Z}_{o(\sigma)}$$

an easy calculation shows that σ acts on this copy of X as a scalar multiple of the identity matrix, i.e. via a central matrix. This is true for every irreducible M-submodule X of V, and hence σ commutes with N when acting on V. Since the above statement is true for each $\sigma \in \mathrm{Alt}_n$ of order 2 or 3, we conclude that φ factors through $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})/[N, \mathrm{Alt}_n]$. Note that we need to consider elements σ of order 3 when we are dealing with the case $n = 4$.

Mennicke's proof of the Congruence Subgroup Property [**Men**] tells us that N is normally generated (as a subgroup of $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$) by the p^{th} powers of the elementary matrices. Now $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$ itself is generated by elementary matrices; let us denote such a matrix by E_{ij} with the usual convention. Observe that for all $\sigma \in \mathrm{Alt}_n$ we have

$$\varphi(E_{\alpha\beta}^{-1} E_{ij}^{p} E_{\alpha\beta}) = \varphi(\sigma^{-1} E_{\alpha\beta}^{-1} E_{ij}^{p} E_{\alpha\beta} \sigma) = \varphi(E_{\sigma(\alpha)\sigma(\beta)}^{-1} E_{ij}^{p} E_{\sigma(\alpha)\sigma(\beta)})$$

Choose $\sigma \in A_n$ such that $\sigma(\alpha) = i$ and $\sigma(\beta) = j$. We conclude that $\varphi(N)$ lies in the centre of $\varphi\big(\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})\big)$. In particular, $\varphi(N)$ is abelian, and hence (as $\mathbb{K}$ is algebraically closed) $\dim W = 1$, as W is an irreducible N-module. Since V is a direct sum of N-modules isomorphic to W, the group N acts via matrices in the centre of $\mathrm{GL}(V)$. Hence N lies in the kernel of the composition

$$\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z}) \xrightarrow{\varphi} \mathrm{GL}(V) \longrightarrow \mathrm{PGL}(V)$$

We have already shown that Z lies in this kernel, and so our representation descends to a projective representation of $\mathrm{PSL}_n(\mathbb{Z}/p\mathbb{Z})$. If we can show that this representation is non-trivial, we can then apply Theorem 2.26 and our proof will be finished.

Suppose that this projective representation is trivial. This means that V is a 1-dimensional, non-trivial $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$-representation. This is however impossible, since the abelianisation of $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$ is trivial when $n \geqslant 3$. □

2.6. Actions of $\mathrm{Out}(F_n)$ on finite sets.

Theorem 2.28. *Every action of* $\mathrm{Out}(F_n)$ *(with* $n \geqslant 6$*) on a set of cardinality* $m \leqslant \binom{n+1}{2}$ *factors through* $\mathbb{Z}/2\mathbb{Z}$.

Proof. Suppose that we are given such an action. It gives us

$$\mathrm{Out}(F_n) \to \mathrm{Sym}_m \hookrightarrow \mathrm{GL}_{m-1}(\mathbb{C})$$

where Sym_m denotes the symmetric group of rank m, and the second map is the standard irreducible representation of Sym_m. Since

$$m - 1 < \binom{n+1}{2}$$

the composition factors through the natural map $\mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z})$ induced by abelianising F_n, by [**Kie1**, Theorem 3.13]. Thus we have

$$\mathrm{Out}(F_n) \to \mathrm{GL}_n(\mathbb{Z}) \to \mathrm{GL}_{m-1}(\mathbb{C})$$

with finite image. The Congruence Subgroup Property [**Men**] tells us that the map $\mathrm{GL}_n(\mathbb{Z}) \to \mathrm{GL}_{m-1}(\mathbb{C})$ factors through a congruence map

$$\mathrm{GL}_n(\mathbb{Z}) \to \mathrm{GL}_n(\mathbb{Z}/p^\alpha\mathbb{Z})$$

for some positive integer α and some prime p. Now

$$m - 1 < 2^{n-1} - 1 \leqslant p^{n-1} - 1$$

and so the restricted map $\mathrm{SL}_n(\mathbb{Z}/p^\alpha\mathbb{Z}) \to \mathrm{GL}_{m-1}(\mathbb{C})$ must be trivial by Theorem 2.27. Thus the given action factors through $\mathrm{GL}_n(\mathbb{Z}/p^\alpha\mathbb{Z})/\mathrm{SL}_n(\mathbb{Z}/p^\alpha\mathbb{Z})$, which is an abelian group. Therefore $\mathrm{SOut}(F_n)$ lies in the kernel of φ, since it is perfect (Proposition 2.2), and we are finished. □

Corollary 2.29. *Every action of* $\mathrm{SOut}(F_n)$ *(with* $n \geqslant 6$*) on a set of cardinality* $m \leqslant \frac{1}{2}\binom{n+1}{2}$ *is trivial.*

Proof. Every action of an index k subgroup of a group G on a set of cardinality m can be induced to an action of G on a set of cardinality km. □

3. The main result

Definition 3.1. Let D_n denote the discrete graph with n vertices.

Definition 3.2. Let $\varphi\colon G \to \mathrm{Out}(A_\Gamma)$ be a homomorphism, and let n be fixed. We define two properties of the action (with respect to n):

- $\mathfrak{C}$ For every G-invariant clique Σ in Γ with at least n vertices there exists a G-invariant subgraph Θ of Γ, such that $\Theta \cap \Sigma$ is a proper non-empty subgraph of Σ.

$\mathfrak{D}$ For every G-invariant subgraph Δ of Γ isomorphic to D_n, there exists a G-invariant subgraph Θ of Γ, such that $\Theta \cap \Delta$ is a proper non-empty subgraph of Δ.

The notation $\mathfrak{C}$ stands for 'clique', and $\mathfrak{D}$ for 'discrete'.

Lemma 3.3. *Let $\varphi\colon G \to \mathrm{Out}(A_\Gamma)$ be an action satisfying $\mathfrak{C}$ and $\mathfrak{D}$. Let Ω be a G-invariant subgraph of Γ. Then both the induced action and the induced quotient action satisfy $\mathfrak{C}$ and $\mathfrak{D}$.*

Proof. Starting with a subgraph Σ or Δ in either Ω or $\Gamma \smallsetminus \Omega$, we observe that the subgraph is a subgraph of Γ, and so using the relevant property we obtain a G-invariant subgraph Θ. We now only need to observe that $\Theta \cap \Omega$ is G-invariant by Lemma 2.19(4), and the image of Θ in $\Gamma \smallsetminus \Omega$ is invariant under the induced quotient action

$$G \to \mathrm{Out}(A_{\Gamma \smallsetminus \Omega})$$ □

Theorem 3.4. *Let us fix positive integers n and $m \geqslant n$. Suppose that a group G satisfies all of the following:*

(1) G is the normal closure of a finite subgroup H.

(2) All homomorphisms

$$G \to \mathrm{Out}(F_k)$$

are trivial when $k \neq n$ and $k < m$.

(3) All homomorphisms

$$G \to \mathrm{GL}_k(\mathbb{Z})$$

are trivial when $k < n$.

(4) Any action of G on a set of cardinality smaller than m is trivial.

Let

$$\varphi\colon G \to \mathrm{Out}(A_\Gamma)$$

be a homomorphism, where Γ has fewer than m vertices. Then φ is trivial, provided that the action satisfies properties $\mathfrak{C}$ and $\mathfrak{D}$ (with respect to n).

Proof. Formally, the proof is an induction on the number of vertices of Γ, and splits into two cases.

Before we proceed, let us observe that assumption (4) allows us to apply Corollary 2.16, and hence to use Lemma 2.19 whenever we need to.

Case 1: Suppose that Γ does not admit proper non-empty G-invariant subgraphs.

Note that this is in particular the case when Γ is a single vertex, which is the base case of our induction.

We claim that Γ is either discrete, or a clique. To prove the claim, let us suppose that Γ is not discrete.

Let v be a vertex of Γ with a non-empty link. Lemma 2.19(3) tells us that $\widehat{\mathrm{st}}(v)$ is G-invariant, and thus it must be equal to Γ. Hence Γ is a join, and therefore admits a join decomposition.

If each factor of the decomposition is a singleton, then Γ is a clique as claimed. Otherwise, the decomposition contains a factor Σ which is not a singleton and not a join, and so in particular not a cone. Thus Lemma 2.19(1) informs us that $\mathrm{lk}(\Sigma)$ is G-invariant. This is a contradiction, since this link is a proper non-empty subgraph. We have thus shown the claim.

Suppose that Γ is a clique, with, say, k vertices. Property $\mathfrak{C}$ immediately tells us that $k < n$, and so we are dealing with a homomorphism

$$\varphi\colon G \to \mathrm{Out}(A_\Gamma) = \mathrm{GL}_k(\mathbb{Z})$$

where $k < n$. Such a homomorphism is trivial by assumption (3).

Suppose that Γ is a discrete graph, with, say, k vertices. Property $\mathfrak{D}$ immediately tells us that $k \neq n$, and so we are dealing with a homomorphism

$$\varphi\colon G \to \mathrm{Out}(A_\Gamma) = \mathrm{Out}(F_k)$$

where $k \neq n$ and $k < m$. Such a homomorphism is trivial by assumption (2).

Case 2: Suppose that Γ admits a proper non-empty G-invariant subgraph Σ.

Lemma 3.3 guarantees that the induced action

$$G \to \mathrm{Out}(A_\Sigma)$$

satisfies the assumptions of our theorem, and thus, using the inductive hypothesis, we conclude that this induced action is trivial.

We argue in an identical manner for the induced quotient action

$$G \to \mathrm{Out}(A_{\Gamma \smallsetminus \Sigma})$$

and conclude that it is also trivial.

These two observations imply that in particular the restriction of these two actions to the finite group H from assumption (1) is trivial. Now Lemma 2.24 tells us that H lies in the kernel of φ, and hence so does G, as it is a normal closure of H by assumption (1). □

Lemma 3.5. *When Γ does not contain n distinct vertices with identical stars, then property $\mathfrak{C}$ is satisfied for any action $G \to \mathrm{Out}^0(A_\Gamma)$.*

Proof. Let Σ be a G-invariant clique in Γ with at least n vertices. Since we know that no n vertices of Γ have identical stars, we need to have distinct vertices of Σ, say v and w, with $\mathrm{st}(v) \neq \mathrm{st}(w)$. Without loss of generality we may assume that there exists $u \in \mathrm{st}(v) \smallsetminus \mathrm{st}(w)$. In particular this implies that u and w are not adjacent.

Consider $\Lambda = \mathrm{lk}(\{u, w\})$: it is invariant by Lemma 2.19(1), since $\{u, w\}$ is not a cone; it intersects Σ non-trivially, since the intersection contains v; the intersection is also proper, since $w \notin \Lambda$. Thus property $\mathfrak{C}$ is satisfied. □

Proposition 3.6. *In Theorem 3.4, we can replace the assumption on the action satisfying $\mathfrak{D}$ by the assumption that Γ is not a join of D_n and another (possibly empty) graph, provided that G satisfies additionally*

(5) G contains a perfect subgroup P, which in turn contains H.

Proof. We are going to proceed by induction on the number of vertices of Γ, as before. Assuming the inductive hypothesis, we will either show the conclusion of the theorem directly, or we will show that in fact property $\mathfrak{D}$ holds.

Note that the base case of induction (Σ being a singleton) always satisfies $\mathfrak{D}$.

Let Δ be as in property $\mathfrak{D}$, and suppose that the property fails for this subgraph.

Case 1: suppose that there exists a vertex u of Δ with a non-empty link.

Let v be a vertex of $\Gamma \smallsetminus \Delta$ joined to some vertex of Δ. Consider $\widehat{\mathrm{st}}(v)$. It is G-invariant by Lemma 2.19(3). If it intersects Δ in a proper subset thereof, then Δ does satisfy property $\mathfrak{D}$. We may thus assume that $\Delta \subseteq \widehat{\mathrm{st}}(v)$.

We would like to apply induction to $\widehat{\mathrm{st}}(v)$, and conclude that it, and hence Δ, are trivialised. This would force Δ to satisfy property $\mathfrak{D}$.

There are two cases in which we cannot apply the inductive hypothesis to $\widehat{\mathrm{st}}(v)$: this subgraph might be equal to Γ, or it might be a join of a subgraph isomorphic to D_n and another subgraph.

In the former case, Γ is a join of two non-empty graphs. If there exists a factor Θ of the join decomposition of Γ which is not a singleton, and which does not contain Δ, then let us look at $\mathrm{lk}(\Theta)$. This is a proper subgraph of Γ, it is

G-invariant by Lemma 2.19(1), and is not a join of D_n and another graph since Γ is not. Thus we may apply the inductive hypothesis to $\mathrm{lk}(\Theta)$ and conclude that it is trivialised. But $\Delta \subseteq \mathrm{lk}(\Theta)$, and so Δ is also trivialised, and thus satisfies $\mathfrak{D}$.

If Γ has no such factor Θ in its join decomposition, then $\Gamma = \mathrm{st}(\Sigma)$, where Σ is a non-empty clique. The clique Σ is a proper subgraph, since it does not contain Δ. It is G-invariant by Lemma 2.19(1) and so the inductive hypothesis tells us that it is trivialised.

The induced quotient action $G \to \mathrm{Out}(A_{\Gamma \smallsetminus \Sigma})$ is also trivialised by induction, as $\Gamma \smallsetminus \Sigma$ cannot be a join of D_n and another graph as before. We now apply Lemma 2.24 for the subgroup H, and conclude that H, and hence its normal closure G, act trivially.

Now we need to look at the situation in which $\widehat{\mathrm{st}}(v)$ is a proper subgraph of Γ, but it is a join of D_n and another graph.

Let us look at Λ, the intersection of $\widehat{\mathrm{st}}(v)$ with the link of all factors of the join decomposition of $\widehat{\mathrm{st}}(v)$ isomorphic to D_n. The subgraph Λ is G-invariant by Lemma 2.19(1) and (4). It is a proper subgraph of Γ, and so the inductive hypothesis tells us that Λ is trivialised. If Λ contains Δ then we are done.

The graph Λ does not contain Δ if and only if Δ is a factor of the join decomposition of $\widehat{\mathrm{st}}(v)$. Observe that we can actually use another vertex of $\Gamma \smallsetminus \Delta$ in place of v, provided that this other vertex is joined by an edge to some vertex of Δ. Thus we may assume that Δ is a factor of the join decomposition of every $\widehat{\mathrm{st}}(v)$ where v is as described. This is however only possible when $\mathrm{st}(\Delta)$ is a connected component of Γ. There must be at least one more component, since Γ is not a join of Δ and another graph.

Note that the component $\mathrm{st}(\Delta)$ is invariant by Lemma 2.19(5).

Suppose that the clique factor Σ of $\mathrm{lk}(\Delta)$ is non-trivial. As before, Σ is trivialised. Observing that $\Gamma \smallsetminus \Sigma$ is disconnected, and if it is discrete then it is has more than n vertices, allows us to apply the inductive hypothesis to the quotient action induced by Σ, and so, arguing as before, we see that Γ is trivialised.

Now suppose that $\mathrm{lk}(\Delta)$ has a trivial clique component. The join decomposition of the component $\mathrm{st}(\Delta)$ consists of at least two factors, each of which is invariant by Lemma 2.19(1). Let Θ be such a factor. Removing Θ leaves us with a disconnected graph smaller than Γ. Thus, we may apply the inductive hypothesis, provided that $\Gamma \smallsetminus \Theta$ is not D_n. This might however occur: in this situation $\mathrm{st}(\Delta) \smallsetminus \Theta$ fulfils the role of the graph Θ from the definition of $\mathfrak{D}$, and so we can use the inductive hypothesis nevertheless.

We now apply Lemma 2.24 to the subgroup H and the induced quotient actions determined by removing two distinct factors of $\mathrm{st}(\Delta)$, and conclude that H, and hence its normal closure G, act trivially on A_Γ.

Case 2: $\mathrm{lk}(u) = \emptyset$ for every vertex u of Δ.

We write $\Gamma = \Gamma_1 \sqcup \cdots \sqcup \Gamma_k \sqcup \Theta$ where the subgraphs Γ_i are non-discrete connected components of Γ, and Θ is discrete. By assumption $\Delta \subseteq \Theta$.

If $k \geqslant 2$, then removing any component Γ_i leaves us with a smaller graph, to which we can apply the inductive hypothesis. Then we use Lemma 2.25.

If $k = 0$ then Θ is not isomorphic to D_n by assumption. Then we know that the action φ is trivial by assumption (2).

If $k = 1$, then we need to look more closely at Γ_1. If Γ_1 does not have factors isomorphic to D_n in its join decomposition, then by induction we know that Γ_1 is trivialised. Now we use Lemma 2.25.

Suppose that Γ_1 contains a subgraph Ω isomorphic to D_n in its join decomposition. If Γ_1 has a non-trivial clique factor, then this factor is invariant, induction tells us that it is trivialised, and the induced quotient action is also trivial. Thus the entire action of H is trivial, thanks to Lemma 2.24, and thus the action of G is trivial, as G is the normal closure of H.

If the clique factor is trivial, then taking links of different factors of the join decomposition of Γ_1 allows us to repeat the argument we just used, and conclude that H, and thus G, act trivially. □

Theorem 3.7. *Let $n \geqslant 6$. Suppose that Γ is a simplicial graph with fewer than $\frac{1}{2}\binom{n}{2}$ vertices. Let $\varphi\colon \mathrm{SOut}(F_n) \to \mathrm{Out}(A_\Gamma)$ be a homomorphism. Then φ is trivial, provided that there are no n vertices in Γ with identical stars, and that Γ is not a join of the discrete graph with n vertices and another (possibly empty) graph.*

Proof. We start by showing that $G = \mathrm{SOut}(F_n)$ satisfies the assumptions (1)–(4) of Theorem 3.4 and (5) of Proposition 3.6, with $m = \frac{1}{2}\binom{n}{2}$.

(1) Let $H = \mathrm{Alt}_n$. The group G is the normal closure of H by Proposition 2.4.
(2) All homomorphisms
$$G \to \mathrm{Out}(F_k)$$
are trivial when $k \neq n$ and $k < m$ by Theorem 2.7.
(3) All homomorphisms
$$G \to \mathrm{GL}_k(\mathbb{Z})$$
are trivial when $k < n$ by Corollary 2.5.
(4) Any action of G on a set of cardinality smaller than m is trivial by Corollary 2.29.
(5) G is perfect by Proposition 2.2.

To verify property $\mathfrak{C}$ we use Lemma 3.5, and property $\mathfrak{D}$ we replace using Proposition 3.6. Now we apply Theorem 3.4. □

4. From larger to smaller RAAGs

In this section we will look at homomorphisms $\mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_{\Gamma'})$, where Γ' has fewer vertices than Γ.

Theorem 4.1. *There are no injective homomorphisms $\mathrm{Out}(A_\Gamma) \to \mathrm{Out}(A_{\Gamma'})$ when Γ' has fewer vertices than Γ.*

Proof. For a group G we define its $\mathbb{Z}_2$-*rank* to be the largest n such that $(\mathbb{Z}_2)^n$ embeds into G.

We claim that the $\mathbb{Z}_2$-rank of $\mathrm{Out}(A_\Gamma)$ is equal to $|\Gamma|$, the number of vertices of Γ.

Firstly, note that for every vertex of Γ we have the corresponding inversion in $\mathrm{Out}(A_\Gamma)$, and these inversions commute; hence the $\mathbb{Z}_2$-rank of $\mathrm{Out}(A_\Gamma)$ is at least $|\Gamma|$.

For the upper bound, observe that the $\mathbb{Z}_2$-rank of $\mathrm{GL}_n(\mathbb{R})$ is equal to n, since we can simultaneously diagonalise commuting involutions in $\mathrm{GL}_n(\mathbb{R})$. Thus, the $\mathbb{Z}_2$-rank of $\mathrm{GL}_n(\mathbb{Z})$ is equal to n as well (since it is easy to produce a subgroup of this rank).

Finally, note that the kernel of the natural map $\mathrm{Out}(A_\Gamma) \to \mathrm{GL}_n(\mathbb{Z})$ with $n = |\Gamma|$ is torsion free by Theorem 2.23, and so the $\mathbb{Z}_2$-rank of $\mathrm{GL}_n(\mathbb{Z})$ is bounded below by the $\mathbb{Z}_2$-rank of $\mathrm{Out}(A_\Gamma)$. □

Remark 4.2. The proof of the above theorem works for many subgroups of $\mathrm{Out}(A_\Gamma)$ as well; specifically it applies to $\mathrm{Out}^0(A_\Gamma)$, the group of *untwisted* outer automorphisms $\mathrm{U}(A_\Gamma)$, and the intersection $\mathrm{U}^0(A_\Gamma) = \mathrm{U}(A_\Gamma) \cap \mathrm{Out}^0(A_\Gamma)$.

It also works when the domain of the homomorphisms is $\mathrm{Aut}(A_\Gamma)$, or more generally any group with $\mathbb{Z}_2$-rank larger than the number of vertices of Γ'.

CHAPTER VIII

Nielsen realisation for untwisted automorphisms of right-angled Artin groups

This is joint work with Sebastian Hensel.

ABSTRACT. We prove Nielsen realisation for finite subgroups of the groups of untwisted outer automorphisms of RAAGs in the following sense: given any graph Γ, and any finite group $G \leqslant \mathcal{U}^0(A_\Gamma) \leqslant \mathrm{Out}^0(A_\Gamma)$, we find a non-positively curved cube complex with fundamental group A_Γ on which G acts by isometries, realising the action on A_Γ.

1. Introduction

A right-angled Artin group (RAAG) A_Γ is a group given by a very simple presentation, which is defined by a graph Γ: the group A_Γ has one generator for each vertex of Γ, and two generators commute if and only if the corresponding vertices are joined by an edge in Γ.

RAAGs have been an object of intense study over the last years, and indeed seem to be ubiquitous in geometry and topology. The most striking example is possibly the role they played in the recent solution of the virtual Haken conjecture by Agol [**Ago**]. They also possess a rich intrinsic structure, maybe most visibly so in the variety of surprising properties their subgroups can exhibit (see e.g. the work of Bestvina and Brady [**BB**]).

A general RAAG A_Γ can be seen as interpolating between a non-abelian free group F_n (corresponding to the graph with n vertices and no edges) and a free Abelian group $\mathbb{Z}^n$ (defined by the complete graph on n vertices). If a property holds for both F_n and $\mathbb{Z}^n$, it is then natural to look for an analogue that works for all RAAGs.

In this article we investigate *Nielsen realisation* from this point of view. For free groups this takes the following form: suppose one is given a finite subgroup $H < \mathrm{Out}(F_n)$. Is there a graph X with $\pi_1(X) = F_n$ on which the group H acts by isometries, inducing the given action on the fundamental group? The answer turns out to be yes (as shown independently by Culler [**Cul**], Khramtsov [**Khr1**], and Zimmermann [**Zim1**]; see also [**HOP**] for a more recent, topological proof).

Let us note here that Nielsen realisation for free groups is equivalent to the statement that in the action of $\mathrm{Out}(F_n)$ on the Culler–Vogtmann Outer Space every finite subgroup fixes a point. The result is also an essential tool in the work of Bridson–Vogtmann [**BV2**] and the second-named author [**Kie1, Kie3**], and is used to prove certain rigidity phenomena for $\mathrm{Out}(F_n)$.

The corresponding statement for free abelian groups follows from the (classical) fact that any finite (in fact compact) subgroup of $\mathrm{GL}_n(\mathbb{R})$ can be conjugated to be a subgroup of the orthogonal group. This implies that any finite $H < \mathrm{Out}(\mathbb{Z}^n) = \mathrm{GL}_n(\mathbb{Z})$ acts isometrically on an n-torus, and the induced action on the fundamental group is the given one.

For RAAGs the natural analogue is as follows: suppose one is given a finite subgroup $H < \mathrm{Out}(A_\Gamma)$ in the outer automorphism group of a RAAG. Is there a compact non-positively curved metric space on which H acts by isometries, realising the action on the fundamental group?

The close relationship between RAAGs and cube complexes tempts one to ask the above question with cube complexes in place of metric spaces. This is however bound to lead to a negative answer, since already for general finite subgroups of $\mathrm{GL}_n(\mathbb{Z})$ the action on the torus described above cannot be made cubical and cocompact simultaneously.

The main result of this article proves Nielsen Realisation for a large class of RAAGs. The restrictions are chosen in a way allowing us to use cube complexes, and we obtain

THEOREM. *Suppose Γ is a simplicial graph, and let $H < \mathcal{U}^0(A_\Gamma)$ be finite. Then there is a compact non-positively curved cube complex realising the action of H. Moreover, the dimension of the complex is the same as the dimension of the Salvetti complex of A_Γ.*

The group $\mathcal{U}^0(A_\Gamma) < \mathrm{Out}(A_\Gamma)$ is the intersection of the group $\mathcal{U}(A_\Gamma)$ of untwisted outer automorphisms (introduced by Charney–Stambaugh–Vogtmann [**CSV**]) with the finite index subgroup $\mathrm{Out}^0(A_\Gamma)$.

Observe that when Γ has no symmetries, and the link of any vertex in Γ is not a cone, then $\mathrm{Out}(A_\Gamma) = \mathcal{U}^0(A_\Gamma)$. In particular, our result holds for connected triangle- and symmetry-free defining graphs.

1.1. Outline of the proof. Since this article is rather substantial in length, let us offer here a somewhat informal outline of the proof of the main theorem.

The proof (Sections 8 through 10) is inductive on the dimension of Γ, that is the maximal size of a maximal clique in Γ.

We proceed by identifying maximal proper subgraphs $\Gamma_1, \dots, \Gamma_n$ of Γ, which are *invariant*, that is we have an induced action of H on (conjugacy classes in) A_{Γ_i}, the subgroup of A_Γ generated by vertices of Γ_i, for each i. We assume that the result holds for Γ_i.

Now one of the following situations occurs: the subgraphs Γ_i might be disjoint, or they might intersect in some non-trivial fashion. The first of these cases is essentially covered by Relative Nielsen Realisation for free products, the main theorem in our previous article [**HK2**, Theorem 5.4]; the precise statement we use here is Proposition 6.1.

The second situation requires a different type of argument. Here we take some maximal invariant proper subgraph Γ', and build the cube complex for Γ from the one for Γ' (given by induction), by gluing to it cube complexes for invariant subgraphs which intersect Γ' and its complement non-trivially. This process presents two types of difficulties: firstly, we need the complexes to agree on the overlap; this requirement forces us to study Relative Nielsen Realisation, which is a way of building cube complexes for our actions having some prescribed subcomplexes. To make sense of such statements we introduce *cubical systems* (Section 5), which are precisely cube complexes with subcomplexes realising the induced actions on relevant invariant subgraphs.

The second difficulty arises when we try to glue two cube complexes over a subcomplex they both posses; here care needs to be taken to make sure that the object we obtain from the gluing realises the given action, and not some other action related to the given one by a partial conjugation. This is the content of Section 7.

For the reasons outlined here, our paper is rich with technical details. The (positive) side effect of this is that the cube complexes realising the action of a finite group $H \to \mathcal{U}^0(A_\Gamma)$ we start with come equipped with a plethora of invariant subcomplexes, which gives our realisation an extra layer of potential applicability.

Acknowledgments. The authors would like to thank Piotr Przytycki and Ruth Charney for many helpful comments and discussions. The authors would furthermore like to thank Karen Vogtmann for discussions and suggesting the statement and use of adapted realisation.

2. Preliminaries

2.1. Graphs and RAAGs. Throughout the paper Γ will denote a fixed simplicial graph. We define the associated RAAG A_Γ to be the group generated by the vertices of Γ, and with a presentation in which the only relations occurring are commutators of vertices adjacent in Γ.

The only subgraphs of Γ we will encounter will be induced subgraphs; such a subgraph is uniquely determined by its vertex set (since Γ is fixed). Hence we will use $\cap, \cup, \smallsetminus$ etc. of two graphs to denote the induced subgraph spanned by the corresponding operation applied to the vertices of the two graphs.

Definition 2.1. Let Δ, Σ be two induced subgraphs of Γ. We say that they form a *join* if and only if each vertex in Δ is connected by an edge (in Γ) to each vertex of Σ. The induced subgraph spanned by all vertices of Δ and Σ will be denoted by $\Delta * \Sigma$.

An induced subgraph Θ is a *join* if and only if we have $\Theta = \Delta * \Sigma$ for some non-empty induced subgraphs Δ and Σ; furthermore Θ is a cone if Δ can be taken to be a singleton.

Note that, in accordance with our convention, we have $\Delta * \Sigma = \Delta \cup \Sigma$, with the join notation indicating the presence of the relevant edges.

Note that induced subgraphs (their vertices to be more specific) generate subgroups of A_Γ; given such a subgraph Δ we will call the corresponding subgroup A_Δ. This subgroup is abstractly isomorphic to the RAAG defined by Δ. We adopt the convention $A_\emptyset = \{1\}$.

Throughout the paper we use the (standard) convention of denoting the normaliser, centraliser, and centre of a subgroup $H \leqslant A_\Gamma$ by, respectively, $N(H), C(H)$ and $Z(H)$. We will also use $c(x) \in \mathrm{Aut}(A_\Gamma)$ to denote conjugation by $x \in A_\Gamma$.

We will need the following definitions throughout the paper. Some of them are new; others may be non-standard.

Definition 2.2. Suppose $\Delta \subseteq \Gamma$ is an induced subgraph.

i) The *link* of Δ is
$$\mathrm{lk}(\Delta) = \bigcap_{v \in \Delta} \mathrm{lk}(v)$$

ii) The *star* of Δ is
$$\mathrm{st}(\Delta) = \mathrm{lk}(\Delta) * \Delta$$

iii) The *extended star* of Δ is
$$\widehat{\mathrm{st}}(\Delta) = \mathrm{lk}(\Delta) * \mathrm{lk}(\mathrm{lk}(\Delta)) = \mathrm{st}(\mathrm{lk}(\Delta))$$

iv) Given a second full subgraph $\Theta \subseteq \Gamma$ with $\Delta \subseteq \Theta$ we define the *restricted link* and *restricted star* of Δ in Θ to be respectively
$$\mathrm{lk}_\Theta(\Delta) = \mathrm{lk}(\Delta) \cap \Theta \text{ and } \mathrm{st}_\Theta(\Delta) = \mathrm{st}(\Delta) \cap \Theta$$

Let us observe the following direct consequences of the definition.

Lemma 2.3. *Let Δ and Θ be two induced subgraphs of Γ, and let v be a vertex of Γ. Then*

(1) $\Delta \subseteq \mathrm{lk}(\Theta) \Leftrightarrow \Theta \subseteq \mathrm{lk}(\Delta)$
(2) $\Delta \subseteq \mathrm{st}(v) \Rightarrow v \in \mathrm{st}(\Delta)$
(3) $\mathrm{lk}(\Delta) \subseteq \mathrm{st}(v) \Rightarrow v \in \widehat{\mathrm{st}}(\Delta)$

Proof.

(1) Both statements are equivalent to saying that each vertex in Δ is connected to each vertex in Θ.
(2) If $v \in \Delta$ then the result follows trivially. If not, then $\Delta \subseteq \mathrm{lk}(v)$ and the result follows from the previous one.
(3) $\mathrm{lk}(\Delta) \subseteq \mathrm{st}(v) \Rightarrow v \in \mathrm{st}(\mathrm{lk}(\Delta)) = \widehat{\mathrm{st}}(\Delta)$. □

Definition 2.4 (Join decomposition)**.** Let $\Delta \subseteq \Gamma$ be an induced subgraph. We say that

$$\Delta = \Delta_1 * \cdots * \Delta_k$$

is a *join decomposition* of Δ if and only if each Δ_i is an induced subgraph of Γ which is not a join.

We define $Z(\Delta)$ to be the union all subgraphs Δ_i which are singletons.

Such a decomposition is unique up to reordering the factors.

Proposition 2.5 ([**CSV**, Proposition 2.2])**.** *Given $\Delta \subseteq \Gamma$ we have the following identifications*

- $N(A_\Delta) = A_{\mathrm{st}(\Delta)} = A_\Delta \times A_{\mathrm{lk}(\Delta)}$
- $Z(A_\Delta) = A_{Z(\Delta)}$
- $C(A_\Delta) = A_{Z(\Delta)} \times A_{\mathrm{lk}(\Delta)}$

Given another induced subgraph $\Sigma \subseteq \Gamma$ we also have

$$x^{-1} A_\Delta x \leqslant A_\Sigma \Longleftrightarrow x \in N(A_\Delta) N(A_\Sigma) \text{ and } \Delta \subseteq \Sigma$$

Definition 2.6. Given a simplicial graph Γ we define its *dimension* $\dim \Gamma$ to be the number of vertices in a largest clique in Γ.

The dimension of Γ coincides with the dimension of the Salvetti complex of A_Γ.

2.2. Words in RAAGs. Since A_Γ is given in terms of a presentation, its elements are equivalence classes of words in the alphabet formed by vertices of Γ (which we will refer to simply as the alphabet Γ). There is a robust notion of normal form based on reduced and cyclically reduced words. We will only mention the results necessary for our arguments; for further details see the work of Servatius [**Ser2**].

Definition 2.7. Given a word $w = v_1 \cdots v_n$, where each v_i is a letter, i.e. a vertex of Γ or its inverse, we define two *basic moves*:

- *reduction*, which consists of removing v_i and v_j from w (with $i < j$), provided that $v_i = v_j^{-1}$, and that v_k commutes with v_i in A_Γ for each $i < k < j$.
- *cyclic reduction*, which consists of removing v_i and v_j from w (with $i < j$), provided that $v_i = v_j^{-1}$, and that v_k commutes with v_i in A_Γ for each $k < i$ and $j < k$.

A word w which does not allow for any reduction is called *reduced*; if in addition it does not allow for any cyclic reduction, it is called *cyclically reduced*.

Servatius shows that, starting with a word w, there is a unique reduced word obtainable from w by reductions, and a unique cyclically reduced word obtainable from w by basic moves. It is clear that the former gives the same element of A_Γ as w did, and the latter gives the same conjugacy class.

He also shows that two reduced words give the same element in A_Γ if and only if they differ by a sequence of moves replacing a subword vv' by $v'v$ with v, v' being commuting letters; let us call those *swaps*.

Lemma 2.8. *Let $\Sigma \subseteq \Gamma$ be an induced subgraph, and suppose that an element $x \in A_\Gamma$ satisfies $x \in y^{-1}A_\Sigma y$ for some $y \in A_\Gamma$. Then the elements of A_Γ given by cyclically reduced words representing the conjugacy class of x lie in A_Σ.*

Proof. Take a reduced word w in Σ representing yxy^{-1}; let w_y be a word in Γ representing y. Then $w_y w w_y^{-1}$ represents x, and it is clear that there is a series of cyclic reductions taking this word to w. Further reductions and cyclic reduction will yield another word in Σ representing the conjugacy class of x. Hence there exists a cyclically reduced word representing the conjugacy class of x as required.

Now suppose that we have two cyclically reduced words, w and w', representing the same conjugacy class in A_Γ, and such that w is a word in Σ. In A_Γ we have the equation

$$w = z^{-1}w'z$$

where z is some reduced word in Γ. Let us take z of minimal word length.

Suppose that the word $z^{-1}w'z$ is not reduced. Then there is a reduction allowed, and it cannot happen within w', since w' is reduced. Thus there exists a letter v such that, without loss of generality, it occurs in z, its inverse occurs in w', and the two can be removed. Note that if the inverse of v occurred only in z^{-1} then performing the reduction would yield a word z', shorter than z, which satisfies the equation

$$w = (z')^{-1}w'z'$$

over A_Γ. This contradicts the minimality of z.

Since v can be removed, we can perform a number of swaps to w' and obtain a reduced word $w''v^{-1}$; we can do the same for z and obtain a reduced word vz'. We now have the following equality in A_Γ

$$w = z'^{-1}v^{-1}w''v^{-1}vz' = z'^{-1}v^{-1}w''z'$$

with $v^{-1}w''$ consisting of exactly the same letters as w' (it is a cyclic conjugate of $w''v^{-1}$), and z' shorter than z. We repeat this procedure until we obtain a reduced word. But then we know that it differs from w by a sequence of swaps, and hence is a word in Σ. Thus w' must have been a word in Σ as well. □

Now we can prove the following proposition.

Proposition 2.9. *Let $\varphi \in \mathrm{Aut}(A_\Gamma)$. Let $\Sigma \subseteq \Gamma$ be such that for all $x \in A_\Sigma$, the element $\varphi(x)$ is conjugate to some element of A_Σ. Then there exists $y \in A_\Gamma$ such that*

$$\varphi(A_\Sigma) \leqslant y^{-1}A_\Sigma y$$

Proof. Let the vertex set of Σ be $\{v_1, \ldots, v_m\}$; these letters are then generators of A_Σ. Let $Z_\Gamma = H_1(A_\Gamma; \mathbb{Z})$ denote the abelianisation of A_Γ, and let $Z_\Sigma \leqslant Z_\Gamma$ denote the image of A_Σ in the abelianisation. Note that Z_Σ is also generated by $\{v_1, \ldots, v_m\}$ in a natural way.

Let $\varphi_* \colon Z_\Gamma \to Z_\Gamma$, be the induced isomorphism on abelianisations. Now, by assumption, φ_* induces a surjection $Z_\Gamma/Z_\Sigma \to Z_\Gamma/Z_\Sigma$; observe that Z_Γ/Z_Σ is isomorphic to $\mathbb{Z}^n$ for some n, and such groups are Hopfian, so this induced morphism is an isomorphism. Hence $\varphi_*|_{Z_\Sigma}$ is an isomorphism (since φ_* is), and so there exists an element $w \in A_\Sigma$, such that its image in Z_Σ is mapped by φ_* to the element $v_1 \cdots v_m$. This implies that w is mapped by φ to a conjugate (by some element y^{-1}) of a cyclically reduced word x, which contains each letter v_i. Crucially, an element $z \in A_\Gamma$ commutes with x if and only if $z \in C(A_\Sigma)$ (as two reduced words define the same element if and only if one can be obtained from the other by a sequence of swaps described above).

Consider $\psi = c(y)\varphi$, so that $\psi(w) = x$. We now aim to show that $\psi(A_\Sigma) \leqslant A_\Sigma$.

Suppose for a contradiction that there exists $u \in A_\Sigma$ such that $\psi(u) \notin A_\Sigma$.

It could be possible that $\psi(u) \in A_{\mathrm{st}(\Sigma)} = A_\Sigma \times A_{\mathrm{lk}(\Sigma)}$. But the only elements in $A_\Sigma \times A_{\mathrm{lk}(\Sigma)}$ conjugate to elements in A_Σ are in fact the elements of A_Σ. Hence we can assume that $\psi(u) \notin A_{\mathrm{st}(\Sigma)}$.

Since $\psi(u)$ is conjugate to an element in A_Σ, yet is not in $A_{\mathrm{st}(\Sigma)}$, we can write

$$\psi(u) = a^{-1}b^{-1}v^{-1}x'vba$$

where the word is reduced, a is a subword containing only letters in

$$Z(\Sigma) * \mathrm{lk}(\Sigma)$$

the subword b contains only letters in $\Sigma \smallsetminus Z(\Sigma)$, the letter v does not lie in $\mathrm{st}(\Sigma)$, and x' is any subword. We will obtain a contradiction from this form of the word.

By assumption $\psi(wu) = xa^{-1}b^{-1}v^{-1}x'vba$ is conjugate to an element of A_Σ. In particular it lies in A_Σ after a sequence of reductions and cyclic reductions, by Lemma 2.8. We are going to visualise the situation as follows. We take a polygon with the number of vertices matching the length of the word describing $\psi(wu)$; now we label the vertices by letters so that going around the polygon clockwise and reading the labels gives us a cyclic conjugate of our word.

In this picture, a reduction or cyclic reduction consists of a deletion of two vertices V_1, V_2 labeled by the same letter but with opposite signs, and such that there is a path between V_1 and V_2 whose vertices are only labeled by letters commuting with the letter labeling V_1.

Recall that we have a finite sequence of basic moves which takes the word

$$xa^{-1}b^{-1}v^{-1}x'vba$$

to a word in A_Σ. Note that after every successive move we can still identify which part of our new word came from x, and which from $a^{-1}b^{-1}v^{-1}x'vba$.

We claim that we can never remove two occurrences of letters in

$$a^{-1}b^{-1}v^{-1}x'vba$$

along a path lying in this subword (and the same is true for the subword x)

Consider the first time we use a move which deletes the occurrences of a letter and its inverse in (what remains of) the subword $a^{-1}b^{-1}v^{-1}x'vba$ along a path lying in this subword. Let q denote the letter we are removing. Since the subword is reduced, such a move was not possible until a prior removal of a letter q' lying between the two letters we are removing, and such that q and q' do not commute. But now q' must have been removed by a path not contained in our subword, since the removal of q is the first move of this type. Therefore the path used to remove q' must contain q, and so q and q' must commute. This is a contradiction which shows the claim.

It is clear that we can remove all letters in a and a^{-1} with a path going over x; let us perform these moves first.

Since our finite sequence of moves takes us to a word in A_Σ, it must at some point remove the letters v and v^{-1}. Consider the first move removing occurrences of this letter. Note that both these occurrences must lie in the subword $v^{-1}x'v$. Hence our move removes this occurrences along a path containing all of x. We have however assumed that v does not commute with x, and so we must have first removed an occurrence of a letter q from x, where q does not commute with v.

In fact all occurrences of q must be removed from bxb^{-1} by reductions before we can remove v. This implies that x contains at least two occurrences of q and its inverse, which is a contradiction. □

2.3. Automorphisms of a RAAG. Let us here briefly discuss a generating set for the group $\mathrm{Aut}(A_\Gamma)$.

By work of Servatius [**Ser2**] and Laurence [**Lau**], $\mathrm{Aut}(A_\Gamma)$ is generated by the following classes of automorphisms:

i) Inversions
ii) Partial conjugations
iii) Transvections
iv) Graph symmetries

Here, an *inversion* maps one generator of A_Γ to its inverse, fixing all other generators.

A *partial conjugation* requires a vertex v in Γ whose star disconnects Γ. For such a v, a partial conjugation is an automorphism which conjugates all generators in one of the complementary components of $\mathrm{st}(v)$ by v and fixes all other generators.

A *transvection* requires vertices v, w with $\mathrm{st}(v) \supseteq \mathrm{lk}(w)$. For such v, w, a transvection is the automorphism which maps w to wv, and fixes all other generators. Transvections come in two types: *folds* (or *type I*) occur when

$$\mathrm{lk}(v) \supseteq \mathrm{lk}(w)$$

and *twists* (or *type II*) when $v \in \mathrm{lk}(w)$.

The group $\mathrm{UAut}(A_\Gamma)$ is defined to be the subgroup generated by all generators from our list except the twists. The group $\mathcal{U}(A_\Gamma)$ is its quotient by inner automorphisms.

A *graph symmetry* is an automorphism of A_Γ which permutes the generators according to a combinatorial automorphism of Γ.

The group $\mathrm{Aut}^0(A_\Gamma)$ is defined to be the subgroup generated by generators of the first four types, i.e. without graph symmetries. Again, $\mathrm{Out}^0(A_\Gamma)$ is its quotient by inner automorphisms.

The group $\mathcal{U}^0(A_\Gamma)$ is defined to be the quotient by inner automorphisms of the subgroup $\mathrm{UAut}^0(A_\Gamma)$ of $\mathrm{Aut}(A_\Gamma)$ generated by all generators from our list except the twists and graph symmetries.

Note that $\mathcal{U}^0(A_\Gamma) = \mathcal{U}(A_\Gamma) \cap \mathrm{Out}^0(A_\Gamma)$, since conjugating any of our generators by a graph symmetry gives a generator of the same kind.

Lemma 2.10. *Suppose that* $\mathrm{lk}(v)$ *is not a cone for all vertices* v *of* Γ*. Then* $\mathrm{UAut}(A_\Gamma) = \mathrm{Aut}(A_\Gamma)$ *and* $\mathrm{UAut}^0(A_\Gamma) = \mathrm{Aut}^0(A_\Gamma)$*.*

Proof. It is enough to show that the assumption prohibits the existence of twists. Let us suppose (for a contradiction) that such a transvection exists; this is equivalent to assuming that there exist vertices v and w such that $v \in \mathrm{lk}(w) \subseteq \mathrm{st}(v)$. But in this case we have

$$\mathrm{lk}(w) = (v \cap \mathrm{lk}(w)) * (\mathrm{lk}(v) \cap \mathrm{lk}(w)) = v * (\mathrm{lk}(w) \smallsetminus v)$$

which is a cone. □

2.4. Markings. We also need the notion of marked topological spaces in the sense introduced in Chapter IV.

Definition 2.11. We say that a path-connected topological space X with a universal covering $\widetilde{X}$ is *marked* by a group A if and only if it comes equipped with an isomorphism between A and the deck transformation group of $\widetilde{X}$.

Remark 2.12. Given a space X marked by a group A, we obtain an isomorphism $A \cong \pi_1(X, p)$ by choosing a basepoint $\widetilde{p} \in \widetilde{X}$ (where p denotes its projection in X). We adopt the notation that the image of a point or set under the universal covering map will be denoted as its *projection*. To keep the notation uniform, we will also call X the projection of $\widetilde{X}$.

Conversely, an isomorphism $A \cong \pi_1(X, p)$ together with a choice of a lift $\widetilde{p} \in \widetilde{X}$ of p determines the marking in the sense of the previous definition.

Definition 2.13. Suppose that we are given an embedding $\pi_1(X) \hookrightarrow \pi_1(Y)$ of fundamental groups of two path-connected spaces X and Y, both marked. A map $\iota \colon X \to Y$ is said to *respect the markings via the map* $\widetilde{\iota}$ if and only if $\widetilde{\iota} \colon \widetilde{X} \to \widetilde{Y}$ is $\pi_1(X)$-equivariant (with respect to the given embedding $\pi_1(X) \hookrightarrow \pi_1(Y)$), and satisfies the commutative diagram

$$\begin{array}{ccc} \widetilde{X} & \xrightarrow{\widetilde{\iota}} & \widetilde{Y} \\ \downarrow & & \downarrow \\ X & \xrightarrow{\iota} & Y \end{array}$$

We say that ι *respects the markings* if and only if such an $\widetilde{\iota}$ exists.

To keep then notation (slightly) more clean, given a space X_Δ we will denote its universal cover by $\widetilde{X}_\Delta$ (rather than $\widetilde{X_\Delta}$).

Next we describe a construction that allows us to glue two marked spaces.

Lemma 2.14. *Suppose that for each $i \in \{0, 1, 2\}$ we are given a group A_i and a space X_i marked by this group. Suppose further that for each $i \in \{1, 2\}$ we are given a monomorphism $\varphi_i \colon A_0 \hookrightarrow A_i$, and a continuous embedding $\iota_i \colon X_0 \hookrightarrow X_i$ which respect the markings via a map $\widetilde{\iota}_i$.*

Then there exists a group A, a space X marked by A, and maps $\varphi'_i \colon A_i \to A$ and $\iota'_i \colon X_i \to X$, the latter respecting the markings via maps $\widetilde{\iota}'_i$, such that the following diagrams commute

$$\begin{array}{ccc} A_0 & \xrightarrow{\varphi_1} & A_1 \\ \downarrow{\scriptstyle \varphi_2} & & \downarrow{\scriptstyle \varphi'_1} \\ A_2 & \xrightarrow{\varphi'_2} & A \end{array} \qquad \begin{array}{ccc} X_0 & \xrightarrow{\iota_1} & X_1 \\ \downarrow{\scriptstyle \iota_2} & & \downarrow{\scriptstyle \iota'_1} \\ X_2 & \xrightarrow{\iota'_2} & X \end{array} \qquad \begin{array}{ccc} \widetilde{X}_0 & \xrightarrow{\widetilde{\iota}_1} & \widetilde{X}_1 \\ \downarrow{\scriptstyle \widetilde{\iota}_2} & & \downarrow{\scriptstyle \widetilde{\iota}'_1} \\ \widetilde{X}_2 & \xrightarrow{\widetilde{\iota}'_2} & \widetilde{X} \end{array}$$

We will refer to the construction in this lemma as *obtaining $\widetilde{X}$ from $\widetilde{X}_1$ and $\widetilde{X}_2$ by gluing* $\mathrm{im}(\widetilde{\iota}_1)$ *to* $\mathrm{im}(\widetilde{\iota}_2)$. Note that the projection X of $\widetilde{X}$ is in fact obtained by an honest gluing of the projections X_1 and X_2 along the respective subspaces.

Proof. We define A and X to be the push-outs of the appropriate diagrams. Take a point $p \in X_0$, and its lift $\widetilde{p} \in \widetilde{X}_0$. The point p has its copies in X_1, X_2 and X; let $\widetilde{q}$ denote some lift of p in $\widetilde{X}$. We also have copies of the point $\widetilde{p}$ in $\widetilde{X}_1$ and $\widetilde{X}_2$. We define the maps $\widetilde{\iota}'_i$ to be the unique maps satisfying the following commutative diagrams of pointed spaces

$$\begin{array}{ccc} \widetilde{X}_i, \widetilde{p} & \xrightarrow{\widetilde{\iota}'_i} & \widetilde{X}, \widetilde{q} \\ \downarrow & & \downarrow \\ X_i, p & \xrightarrow{\iota_i} & X, p \end{array}$$

where the vertical maps are the coverings. The verification that these maps satisfy the third commutative diagram above is easy. □

3. Relative Nielsen Realisation

One of the crucial tools used in this article is the Relative Nielsen Realisation theorem(Theorem 7.5). In this section we will look at the necessary definitions, the statement of the theorem, and then a special case thereof.

Definition 3.1. Let A and H be groups, and let $\varphi\colon H \to \mathrm{Out}(A)$ be a homomorphism. We say that a metric space X, on which H acts by isometries, *realises* the action φ if and only if X is marked by A, and the action of H on conjugacy classes in $\pi_1(X) \cong A$ is equal to the one induced by φ.

If X is a (metric) graph or a cube complex, we require the action to respect the combinatorial structure as well.

Remark 3.2. The action of H on X induces a group extension

$$A \to \overline{A} \to H$$

When A is centre-free, this extension carries precisely the same information as the action $H \to \mathrm{Out}(A)$; when A has non-trivial centre however, then the extension contains more information.

Definition 3.3. Suppose that A contains a subgroup A_1 which is invariant under the (outer) action of H (up to conjugation). Then the extension $\overline{A}$ contains a subgroup, such that the map $\overline{A} \to H$ restricted to this subgroup is onto H, and its kernel is equal to the subgroup. We call this subgroup $\overline{A_1}$.

Theorem 3.4 (Relative Nielsen Realisation [**HK2**, Theorem 5.4])**.** *Let*

$$\varphi\colon H \to \mathrm{Out}(A)$$

be a homomorphism with a finite domain, and let

$$A = A_1 * \cdots * A_n * B$$

be a free-product decomposition, with B a (possibly trivial) finitely generated free group, such that H preserves the conjugacy class of each A_i.

Suppose that for each $i \in \{1, \ldots, m\}$ we are given an NPC space X_i marked by A_i, on which H acts in such a way that the associated extension of A_i by H is isomorphic (as an extension) to the extension $\overline{A_i}$ coming from $\overline{A}$. Then there exists an NPC space X realising the action φ, and such that for each $i \in \{1, \ldots, m\}$ we have an H-equivariant embedding $\iota_i\colon X_i \to X$ which preserves the marking.

Moreover, the images of the spaces X_i are disjoint, and collapsing each X_i individually to a point yields a graph with fundamental group abstractly isomorphic to the free group B.

Note that the above is a slightly restricted version of the Relative Nielsen Realisation, since in its most general form H is allowed to permute the factors A_i.

Remark 3.5 ([**HK2**, Remark 5.6])**.** When each X_i is a cube complex, we may take X to be a cube complex too; the embeddings $X_i \hookrightarrow X$ are now maps of cube complexes, provided that we allow ourselves to cubically barycentrically subdivide the complexes X_i.

Now let us state a version of the above theorem stated for graphs.

Throughout, we consider only graphs without vertices of valence 1, that is without leaves.

Theorem 3.6 (Adapted Realisation)**.** *Let H be finite and $\varphi\colon H \to \mathrm{Out}(F_n)$ be a homomorphism. Suppose that*

$$F_n = A_1 * \cdots * A_k * B$$

is a free splitting such that $\varphi(H)$ preserves the conjugacy class of A_i for each $1 \leqslant i \leqslant k$. Let $\varphi_i\colon H \to \mathrm{Out}(A_i)$ denote the induced actions. For each i let X_i be a marked metric graph realising φ_i.

Then there is a marked metric graph X with the following properties.

i) H acts on X by combinatorial isometries, realising φ.
ii) There are isometric embeddings $\iota_i\colon X_i \to X$ which respect the markings via maps $\widetilde{\iota_i}$.
iii) The embedded subgraphs X_i are preserved by H, and the restricted action induces the action φ_i on X_i up to homotopy.

When $\pi_1(X_i) = A_i \not\cong \mathbb{Z}$, the map ι_i is actually H-equivariant (and not just H-equivariant up to homotopy).

We also arrange for the images $\iota_i(X_i)$ to be pairwise disjoint, unless we have $F_n = A_1 * A_2$, in which case we require $X = \mathrm{im}(\iota_1) \cup \mathrm{im}(\iota_2)$, and $\mathrm{im}(\widetilde{\iota_1}) \cap \mathrm{im}(\widetilde{\iota_2})$ to be a single point.

Note that since we only require the embeddings ι_i to be isometric, the presence of vertices of valence 2 is completely irrelevant. They might appear in the graphs X_i as well as X to make the action of H combinatorial, but we do not require these appearances to agree under ι_i.

Proof. Without loss of generality let us assume that the factors A_i are non-trivial.

When $n = 1$ we either have $F_n = B$, in which case we take X to be the circle, or we have $F_n = A_1$, in which case we take $X = X_1$.

When $n \geqslant 2$, the group F_n has trivial centre, and so the action $H \to \mathrm{Out}(F_n)$ yields a finite extension

$$F_n \to \overline{F_n} \to H$$

For any $A_i \not\cong \mathbb{Z}$, the extension yielded by the action of H on X_i and the extension $\overline{A_i}$ are isomorphic as extensions. For convenience let us set $Y_i = X_i$ in this case.

When $A_i \cong \mathbb{Z}$, there exists an action of H on a circle S_1 such that the induced extension agrees with $\overline{A_i}$. We define Y_i to be precisely this circle with the H-action. Note that the action of H on Y_i and X_i agree up to homotopy.

We now apply Theorem 3.4, using the graphs Y_i.

When $F_n = A_1 * A_2$, we know that collapsing the images of Y_1 and Y_2 yields a graph with trivial fundamental group, and hence a tree. The preimage of this tree lifts is a forest in X in such a way that each connected component of the forest intersects each Y_i in at most one point. Collapsing each of these components individually to a point yields the result. □

Lemma 3.7. *In the context of Theorem 3.6, the projection of the intersection point*

$$\mathrm{im}(\widetilde{\iota_1}) \cap \mathrm{im}(\widetilde{\iota_2})$$

in X is H-fixed (if it exists).

Proof. The Seifert–van Kampen Theorem tells us that the intersection

$$\mathrm{im}(\iota_1) \cap \mathrm{im}(\iota_2)$$

is simply-connected, and so in particular path connected. But this means that it can be lifted to the universal cover in such a way that the lift lies within

$$\mathrm{im}(\widetilde{\iota_1}) \cap \mathrm{im}(\widetilde{\iota_2})$$

which is just a singleton. Hence so is $\mathrm{im}(\iota_1) \cap \mathrm{im}(\iota_2)$. Now this point is the intersection of two H-invariant subspaces, and so is itself H-invariant, and thus H-fixed. □

4. Systems of subgraphs and invariance

Definition 4.1. Given a homomorphism $\varphi\colon H \to \mathrm{Out}(A_\Gamma)$ we define *the system of invariant subgraphs* $\mathcal{L}^\varphi$ to be the set of induced subgraphs $\Delta \subseteq \Gamma$ (including the empty one) such that $\varphi(H)$ preserves the conjugacy class of $A_\Delta \leqslant A_\Gamma$. Note that L^φ is partially ordered by inclusion.

We next show that L^φ is closed under taking intersections and certain unions.

Lemma 4.2. *Let* $\Delta, \Sigma \in \mathrm{L}^\varphi$. *Then*

i) $\Delta \cap \Sigma \in \mathrm{L}^\varphi$.
ii) If $\mathrm{lk}(\Delta \cap \Sigma) \subseteq \mathrm{st}(\Delta)$ *then* $\Delta \cup \Sigma \in \mathrm{L}^\varphi$.

Proof. i) Pick $h \in H$. Since $\Delta \in \mathrm{L}^\varphi$, there exists a representative

$$h_1 \in \mathrm{Aut}(A_\Gamma)$$

of $\varphi(h)$ such that $h_1(A_\Delta) = A_\Delta$. Analogously, there exists $h_2 \in \mathrm{Aut}(A_\Gamma)$ representing $\varphi(h)$ such that $h_2(A_\Sigma) = A_\Sigma$.

Since h_1 and h_2 represent the same element in $\mathrm{Out}(A_\Gamma)$, we have

$$h_1^{-1} h_2 = c(r)$$

with $r \in A_\Gamma$.

Let $\Theta = \Delta \cap \Sigma$, and take $x \in A_\Theta$. Then $h_1(x) \in A_\Delta$, and so any cyclically reduced word in the alphabet Γ representing the conjugacy class of $h_1(x)$ lies in A_Δ (by Lemma 2.8). Now

$$c(r)h_1(x) = h_2(x) \in A_\Sigma$$

and thus any reduced word representing the conjugacy class of $h_1(x)$ lies in A_Σ. Hence such a word lies in A_Θ. But this implies that $h_1(x)$ is a conjugate of an element of A_Θ for each $x \in A_\Theta$, and so we apply Proposition 2.9 and conclude that

$$h_1(A_\Theta) \leqslant y^{-1} A_\Theta y$$

for some $y \in A_\Gamma$.

We repeat the argument for h_1^{-1}, which is a representative of $\varphi(h^{-1})$, and obtain

$$A_\Theta = h_1^{-1} h_1(A_\Theta) \leqslant h_1^{-1}(y^{-1} A_\Theta y) \leqslant y'^{-1} A_\Theta y'$$

for some $y' \in A_\Gamma$. Now Proposition 2.5 tells us that $y'^{-1} A_\Theta y' = A_\Theta$ and so both inequalities in the expression above are in fact equalities. Thus

$$h_1(A_\Theta) = y^{-1} A_\Theta y$$

and so $\Theta \in \mathrm{L}^\varphi$ as claimed.

ii) Now suppose that $\mathrm{lk}(\Theta) \subseteq \mathrm{st}(\Delta)$. Since $\Theta \in \mathrm{L}^\varphi$, there exists a representative h_3 of $\varphi(h)$ which fixes A_Θ. Since $\Delta \in \mathrm{L}^\varphi$, the subgroup $h_3(A_\Delta)$ is a conjugate of A_Δ by an element $r \in A_\Gamma$. Now we have

$$r A_\Theta r^{-1} \leqslant A_\Delta$$

and so by Proposition 2.5 we know that

$$r \in N(A_\Theta) N(A_\Delta) = A_{\mathrm{st}(\Theta)} A_{\mathrm{st}(\Delta)} = A_{\mathrm{st}(\Delta)}$$

since $\mathrm{lk}(\Theta) \subseteq \mathrm{st}(\Delta)$ by assumption. But then $r^{-1} A_\Delta r = A_\Delta$, and so $h_3(A_\Delta) = A_\Delta$.

Let us apply the same argument to $h_3(A_\Sigma)$ – it must be a conjugate of A_Σ by some $s \in A_\Gamma$, and we conclude as above that

$$s \in N(A_\Theta)N(A_\Sigma) \leqslant A_{\mathrm{st}(\Delta)}A_{\mathrm{st}(\Sigma)}$$

We take a new representative h_4 of $\varphi(h)$, which differs from h_3 by the conjugation by the $A_{\mathrm{st}(\Delta)}$-factor of s. This way we get $h_4(A_\Delta) = A_\Delta$, and $h_4(A_\Sigma)$ equal to a conjugate of A_Σ by an element of $A_{\mathrm{st}(\Sigma)} = N(A_\Sigma)$. Hence we have $h_4(A_\Sigma) = A_\Sigma$, and the result follows. □

The following lemma is very much motivated by the work of Charney–Crisp–Vogtmann [**CSV**].

Lemma 4.3. *Suppose that $\varphi(H) \leqslant \mathrm{Out}^0(A_\Gamma)$. Then $\mathcal{L}^\varphi$ contains*

1. *each connected component of Γ which contains at least one edge;*
2. *the extended star of each induced subgraph;*
3. *the link of each subgraph Δ, such that Δ is not a cone;*
4. *the star of each subgraph in L^φ.*

Proof. We will prove the first three points on our list for $\mathrm{Out}^0(A_\Gamma)$ (and therefore for any subgroup). It is enough to verify that each type of generator of $\mathrm{Out}^0(A_\Gamma)$ preserves the listed subgroups up to conjugacy. It is certainly true for all inversions, and thus we only need to verify it for transvections and partial conjugations.

We will make a rather liberal use of Lemma 2.3.

Transvections. Take two vertices in Γ, say v and w, such that

$$\mathrm{lk}(w) \subseteq \mathrm{st}(v)$$

In this case we have a transvection $w \mapsto wv$. To prove our assertion we need to check that whenever w belongs to the subgraph defining our subgroup, so does v.

1. If w belongs to a connected component of Γ which is not a singleton, then $\mathrm{lk}(w) \neq \emptyset$ lies in the same component, and so our assumption $\mathrm{lk}(w) \subseteq \mathrm{st}(v)$ forces v to lie in the component as well.
2. Take $\Delta \subseteq \Gamma$ with $w \in \widehat{\mathrm{st}}(\Delta)$. If $w \in \mathrm{lk}(\mathrm{lk}(\Delta))$ then
$$\mathrm{lk}(\Delta) \subseteq \mathrm{lk}(w) \subseteq \mathrm{st}(v)$$
and so $v \in \widehat{\mathrm{st}}(\Delta)$. If $w \in \mathrm{lk}(\Delta)$ then $\Delta \subseteq \mathrm{lk}(w) \subseteq \mathrm{st}(v)$ and so $v \in \mathrm{st}(\Delta)$.
3. Take $\Delta \subseteq \Gamma$ which is not a cone, and such that $w \in \mathrm{lk}(\Delta)$. Then $\Delta \subseteq \mathrm{st}(v)$, and so $v \in \mathrm{st}(\Delta)$ as above. However, if $v \in \Delta$, then $\Delta \smallsetminus \{v\} \subseteq \mathrm{st}(v) \smallsetminus \{v\} = \mathrm{lk}(v)$, and so Δ is a cone over v, which is a contradiction. Thus $v \in \mathrm{lk}(\Delta)$.

Note that in the last part the assumption of Δ not being a cone is used only to guarantee that $v \notin \Delta$. If the transvection under consideration was of type I, we would now this immediately since $\Delta \subseteq \mathrm{lk}(w) \subseteq \mathrm{lk}(v)$ in this case, and so the assumption of Δ not being a cone would be unnecessary.

Partial conjugations. Take a vertex v in Γ, such that its star disconnects Γ. In this case we have a partial conjugation of the subgroup A_Θ by v, with $\Theta \subseteq \Gamma \smallsetminus \mathrm{st}(v)$ being a connected component. Let $\Sigma = \Gamma \smallsetminus (\Theta \cup \mathrm{st}(v))$. To show that this automorphism preserves the desired subgroups up to conjugation, we need to show that if the subgraphs defining the subgroups do not contain v, then they cannot intersect both Θ and Σ.

1. If a connected component intersects (and so contains) Θ but does not contain v, then it is in fact equal to Θ, and so intersects Σ trivially.

(2) Take $\Delta \subseteq \Gamma$ with $v \notin \widehat{\mathrm{st}}(\Delta)$, and such that $\widehat{\mathrm{st}}(\Delta)$ intersects Θ and Σ non-trivially. If $\mathrm{lk}(\Delta) \subseteq \mathrm{st}(v)$, then $v \in \widehat{\mathrm{st}}(\Delta)$, which is a contradiction. Hence $\mathrm{lk}(\mathrm{lk}(\Delta))$ cannot intersect both Θ and Σ, since if it did, then we would have $\mathrm{lk}(\Delta) \subseteq \mathrm{st}(v)$. Similarly, if $\Delta \subseteq \mathrm{st}(v)$, then $v \in \mathrm{st}(\Delta)$. This is again a contradiction, and again we conclude that $\mathrm{lk}(\Delta)$ cannot intersect both Θ and Σ. Hence, without loss of generality, we have $\mathrm{lk}(\mathrm{lk}(\Delta))$ intersecting Θ and $\mathrm{lk}(\Delta)$ intersecting Σ. But then the two cannot form a join. This is a contradiction.

(3) Take $\Delta \subseteq \Gamma$, with $\mathrm{lk}(\Delta)$ intersecting both Θ and Σ non-trivially. This condition forces $\Delta \subseteq \mathrm{st}(v)$. In fact we see that $\Delta \subseteq \mathrm{lk}(v)$, since otherwise we would have $\mathrm{lk}(\Delta) \subseteq \mathrm{st}(v)$, and thus the link would intersect neither Θ nor Σ. Now $\Delta \subseteq \mathrm{lk}(v)$ gives $v \in \mathrm{lk}(\Delta)$.

Note that in the last part we did not use the assumption on Δ not being a cone.

Now we need to prove (4). Take $\Delta \in \mathrm{L}^\varphi$. Pick $h \in H$ and let

$$h_1 \in \mathrm{Aut}(A_\Gamma)$$

be a representative of $\varphi(h)$ such that $h_1(A_\Delta) = A_\Delta$. Then also

$$h_1(N(A_\Delta)) = N(A_\Delta) = A_{\mathrm{st}(\Delta)}$$ □

Definition 4.4. We say that $\varphi\colon H \to \mathrm{Out}(A_\Gamma)$ is *link-preserving* if and only if L^φ contains links of all induced subgraphs of Γ.

Note that if $\dim \Gamma = 1$, then every action φ is link-preserving.

Remark 4.5. Note that when φ is link-preserving, then so is every induced action $H \to \mathrm{Out}(A_\Sigma)$ for each $\Sigma \in \mathrm{L}^\varphi$.

Lemma 4.6. *If $\varphi(H) \leqslant \mathrm{U}^0(A_\Gamma)$ then φ is link-preserving.*

Proof. Let $\Delta \subseteq \Gamma$ be given. Again we will show that each generator of $\mathrm{U}^0(A_\Gamma)$ sends $A_{\mathrm{lk}(\Delta)}$ to a conjugate of itself. The inversions clearly have the desired property; so do partial conjugations and transvections of type I, since the proof of (3) above did not use the assumption on Δ not being a cone (as remarked). But these are the generators of $\mathrm{U}^0(A_\Gamma)$ and so we are done. □

Corollary 4.7. *If $\dim \Gamma = 2$ and Γ has no leaves, then any $\varphi\colon H \to \mathrm{Out}^0(A_\Gamma)$ is link-preserving.*

Proof. Since $\dim \Gamma = 2$, the links of vertices are discrete graphs; since Γ has no leaves, they contain at least 2 vertices. Hence such links are never cones, and Lemma 2.10 tells us that

$$\mathrm{Out}^0(A_\Gamma) = \mathrm{U}^0(A_\Gamma)$$

Lemma 4.6 completes the proof. □

Definition 4.8. i) Any subset $\mathcal{S}$ of L^φ closed under taking intersections of its elements will be called a *subsystem of invariant subgraphs.*

ii) Given such a subsystem $\mathcal{S}$, and any induced subgraph $\Theta \in \mathrm{L}^\varphi$, we define
- $\mathcal{S}_\Theta = \{\Sigma \cap \Theta \mid \Sigma \in \mathcal{S}\}$
- $\bigcup \mathcal{S} = \bigcup_{\Sigma \in \mathcal{S}} \Sigma$
- $\bigcap \mathcal{S} = \bigcap_{\Sigma \in \mathcal{S}} \Sigma$

Lemma 4.9. *Let $\mathbb{P} \subseteq \mathrm{L}^\varphi$ be a subsystem of invariant graphs, and let $\Theta \in \mathbb{P}$. Then*

$$\mathbb{P}_\Theta = \{\Delta \in \mathbb{P} \mid \Delta \subseteq \Theta\}$$

Proof. This follows directly from the fact that $\mathbb{P}$ is closed under taking intersections. □

Note that given $\Delta \in \mathrm{L}^\varphi$, we get an induced action $\psi\colon H \to \mathrm{Out}(A_\Delta)$. This follows from the fact that the normaliser $N(A_\Delta)$ satisfies

$$N(A_\Delta) = A_{\mathrm{st}(\Delta)} = A_\Delta \times A_{\mathrm{lk}(\Delta)}$$

and $A_{\mathrm{lk}(\Delta)}$ centralises A_Δ.

It is immediate that $\mathrm{L}^\psi = \mathrm{L}^\varphi_\Delta$.

4.1. Boundaries of subgraphs. Let us record here a useful fact about boundaries of subgraphs.

Definition 4.10. Let $\Sigma \subseteq \Gamma$ be a subgraph. We define its *boundary* $\partial\Sigma$ to be the set of all vertices of Σ whose link is not contained in Σ.

Lemma 4.11. *Suppose that Σ is a maximal (with respect to inclusion) proper subgraph of Γ such that $\Sigma \in \mathrm{L}^\varphi$. Then for every $w \in \partial\Sigma$ we have $\Gamma \smallsetminus \Sigma \subseteq \mathrm{lk}(w)$.*

Proof. Let $\Theta = \Gamma \smallsetminus \Sigma$, and let $w \in \partial\Sigma$. We have $\widehat{\mathrm{st}}(w) \in \mathrm{L}^\varphi$. Now

$$\mathrm{lk}(\widehat{\mathrm{st}}(w) \cap \Sigma) \subseteq \mathrm{lk}(w) \subseteq \widehat{\mathrm{st}}(w)$$

and so $\Sigma \cup \widehat{\mathrm{st}}(w) \in \mathrm{L}^\varphi$ by Lemma 4.2. Thus the maximality of Σ implies that $\Sigma \cup \widehat{\mathrm{st}}(w) = \Gamma$. Hence $\Theta \subseteq \widehat{\mathrm{st}}(w)$. Let $\Theta_1 = \Theta \cap \mathrm{lk}(w)$ and $\Theta_2 = \Theta \cap \mathrm{lk}(\mathrm{lk}(w))$. If the latter is empty then we are done, so let us assume that it is not empty.

If $\widehat{\mathrm{st}}(w) = \mathrm{st}(w)$ then $\Theta_2 = \Theta \cap \{w\} = \emptyset$ and we are done. Otherwise $\mathrm{lk}(\mathrm{lk}(w))$ is not a cone (since w is isolated in it), and thus the link of $\mathrm{lk}(\mathrm{lk}(w))$ is contained in L^φ. But this triple link is in fact equal to $\mathrm{lk}(w)$, and so we have $\mathrm{lk}(w) \in \mathrm{L}^\varphi$. Since L^φ is closed under taking intersections, we also have $\Sigma \cap \mathrm{lk}(w) \in \mathrm{L}^\varphi$ and thus also $\mathrm{st}(\Sigma \cap \mathrm{lk}(w)) \in \mathrm{L}^\varphi$. We have

$$\Theta_2 = \Theta \cap \mathrm{lk}(\mathrm{lk}(w)) \subseteq \mathrm{lk}(\mathrm{lk}(w)) \subseteq \mathrm{lk}(\Sigma \cap \mathrm{lk}(w)) \subseteq \mathrm{st}(\Sigma \cap \mathrm{lk}(w))$$

and so $\mathrm{st}(\Sigma \cap \mathrm{lk}(w))$ is not contained in Σ. We also have

$$\Sigma \cup \mathrm{st}(\Sigma \cap \mathrm{lk}(w)) \in \mathrm{L}^\varphi$$

since

$$\mathrm{lk}(\Sigma \cap \mathrm{st}(\Sigma \cap \mathrm{lk}(w))) \subseteq \mathrm{lk}(\Sigma \cap \mathrm{lk}(w)) \subseteq \mathrm{st}(\Sigma \cap \mathrm{lk}(w))$$

as before. The graph $\Sigma \cup \mathrm{st}(\Sigma \cap \mathrm{lk}(w))$ must contain Θ_1 as well, by the maximality of Σ. Thus $\Theta_1 \subseteq \mathrm{st}(\Sigma \cap \mathrm{lk}(w))$, and so

$$\Theta_1 \subseteq \mathrm{lk}(\mathrm{lk}_\Sigma(w))$$

We have $\Theta_1 \subseteq \mathrm{lk}(w) = \mathrm{lk}(\mathrm{lk}(\mathrm{lk}(w)))$ and so, combining this with the previous observation, we get

$$\begin{aligned}
\Theta_1 &\subseteq \mathrm{lk}\left(\mathrm{lk}_\Sigma(w)\right) \cap \mathrm{lk}\left(\mathrm{lk}(\mathrm{lk}(w))\right) \\
&= \mathrm{lk}\Big(\mathrm{lk}(\mathrm{lk}(w)) \cup (\mathrm{lk}(w) \cap \Sigma)\Big) \\
&\subseteq \mathrm{lk}\Big(\big(\mathrm{lk}(\mathrm{lk}(w)) \cup \mathrm{lk}(w)\big) \cap \Sigma\Big) \\
&= \mathrm{lk}\left(\widehat{\mathrm{st}}(w) \cap \Sigma\right) \\
&\subseteq \mathrm{st}\left(\widehat{\mathrm{st}}(w) \cap \Sigma\right)
\end{aligned}$$

The last subgraph is a star of a subgraph in L^φ, and so is itself in L^φ. As before we have

$$\mathrm{lk}\Big(\mathrm{st}\left(\widehat{\mathrm{st}}(w) \cap \Sigma\right) \cap \Sigma\Big) \subseteq \mathrm{lk}(w) = \mathrm{lk}_\Sigma(w) \cup \Theta_1 \subseteq \mathrm{st}\left(\widehat{\mathrm{st}}(w) \cap \Sigma\right)$$

and so $\mathrm{st}\left(\widehat{\mathrm{st}}(w) \cap \Sigma\right) \cup \Sigma \in \mathrm{L}^\varphi$. It contains Σ and Θ_1, and therefore it must contain all of Θ. But then $\Theta \subseteq \mathrm{st}\left(\widehat{\mathrm{st}}(w) \cap \Sigma\right)$, which is only possible when

$$\Theta \subseteq \mathrm{lk}\left(\widehat{\mathrm{st}}(w) \cap \Sigma\right) \subseteq \mathrm{lk}(w)$$

which is the desired statement. □

5. Cubical systems

In this section we give the definition of the most fundamental object in the paper. We then begin proving some central properties that will be used throughout.

Definition 5.1 (Metric cube complex)**.** i) A *metric cube complex* is a (realisation of a) combinatorial cube complex, which comes equipped with a metric such that every n-cube in X (for each n) is isometric to a Cartesian product of n closed intervals in $\mathbb{R}$.

ii) Given two such complexes X and Y, we say that Y is a *subdivision* of X if and only if the combinatorial cube complex underlying Y is a subdivision of X, and the induced map (on the realisations) $Y \to X$ is an isometry.

iii) We say that $Z \subseteq X$ is a *subcomplex* of X if and only if there exists a subdivision Y of X such that, under the identification $X = Y$, the subspace $Z \subseteq Y$ is a subcomplex in the combinatorial sense.

iv) A connected metric cube complex is called *non-positively curved*, or NPC, whenever its universal cover is CAT(0).

At this point we want to warn the reader that our notion of subcomplex is more relaxed than the usual definition.

Note further that we will make no distinction between a metric cube complex X and its subdivisions; as metric spaces they are isomorphic; moreover given any group action (by isometries) $H \curvearrowright X$ which respects the combinatorial structure of X and any subdivision Y of X, there exists a further subdivision Z of Y such that the inherited action of H on Z respects the combinatorial structure of Z.

Definition 5.2 (Cubical systems)**.** Suppose we have a subsystem of invariant subgraphs $\mathbb{P}$, such that $\mathbb{P}$ is closed under taking restricted links, that is for all $\Delta, E \in \mathbb{P}$ with $\Delta \subseteq E$ we have $\mathrm{lk}_E(\Delta) \in \mathbb{P}$. A *cubical system* $\mathcal{X}$ (for $\mathbb{P}$) consists of the following data.

(1) For each $\Delta \in \mathbb{P}$ a marked metric NPC cube complex X_Δ, of the same dimension as A_Δ, realising $H \to \mathrm{Out}(A_\Delta)$. We additionally require X_Δ not to have leaves when Δ is 1-dimensional.

(2) For each pair $\Delta, \Theta \in \mathbb{P}$ with $\Delta \subseteq \Theta$, an H-equivariant isometric embedding

$$\iota_{\Delta,\Theta} \colon X_\Delta \times X_{\mathrm{lk}_\Theta(\Delta)} \to X_\Theta$$

whose image is a subcomplex, and which respects the markings via a map $\widetilde{\iota}_{\Delta,\Theta}$, where the product is given the product marking. We set $\iota_{\Delta,\Delta}$ and $\widetilde{\iota}_{\Delta,\Delta}$ to be the respective identity maps.

Given $\Delta, \Theta \in \mathbb{P}$ we will refer to the map $\widetilde{\iota}_{\Delta,\Theta}|_{\widetilde{X}_\Delta \times \{x\}}$ for any $x \in \widetilde{X}_{\mathrm{lk}_\Theta(\Delta)}$ as the *standard copy of* $\widetilde{X}_\Delta$ *in* $\widetilde{X}_\Theta$ *determined by* x, or simply *a standard copy of* $\widetilde{X}_\Delta$ *in* $\widetilde{X}_\Theta$.

We say that such a standard copy is *fixed* if and only if its projection in X_Θ is H-invariant.

We require our maps to satisfy four conditions.

Product Axiom. Given $\Delta, \Theta \in \mathbb{P}$ such that $\Theta = \Delta * \mathrm{lk}_\Theta(\Delta)$, we require $\widetilde{\iota}_{\Delta,\Theta}$ to be surjective.

Orthogonal Axiom. Given $\Delta \subseteq \Theta$, both in $\mathbb{P}$, for each $\widetilde{x} \in \widetilde{X}_\Delta$ we require $\widetilde{\iota}_{\Delta,\Theta}(\{\widetilde{x}\} \times \widetilde{X}_{\mathrm{lk}_\Theta(\Delta)})$ to be equal to the image of some standard copy of $\widetilde{X}_{\mathrm{lk}_\Theta(\Delta)}$ in $\widetilde{X}_\Theta$.

Intersection Axiom. Let $\Sigma_1, \Sigma_2, \Theta \in \mathbb{P}$ be such that $\Sigma_i \subseteq \Theta$ for both values of i. Suppose that we are given standard copies of $\widetilde{X}_{\Sigma_i}$ in $\widetilde{X}_\Theta$ whose images intersect non-trivially. Then the intersection of the images is equal to the image of a standard copy of $\widetilde{X}_{\Sigma_1 \cap \Sigma_2}$ in $\widetilde{X}_\Theta$. Moreover, this intersection is also the image of a standard copy of $\widetilde{X}_{\Sigma_1 \cap \Sigma_2}$ in $\widetilde{X}_{\Sigma_i}$, under the given standard copy $\widetilde{X}_{\Sigma_i} \to \widetilde{X}_\Theta$, for both values of i

System Intersection Axiom. Let $\mathcal{S} \subseteq \mathbb{P}$ be a subsystem of invariant graphs closed under taking unions of its elements. Then for each $\Sigma \in \mathcal{S}$ there exists a standard copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_{\bigcup \mathcal{S}}$ such that the images of all of these copies intersect non-trivially.

Remark 5.3. We will often make no distinction between a standard copy and its image. Let us remark here that any standard copy is a subcomplex (with our non-standard definition of a subcomplex; see above).

Now we will list some implications of the definition.

Remark 5.4. Suppose we are given a cubical system $\mathcal{X}$ for L^φ, and let $\Delta \in \mathrm{L}^\varphi$. Then the subsystem $\mathcal{X}_\Delta$, consisting of all complexes $X_\Sigma \in \mathcal{X}$ with $\Sigma \subseteq \Delta$ together with all relevant maps, is a cubical system for $\mathrm{L}^\varphi_\Delta$.

Lemma 5.5. *Let $\mathcal{X}$ be a cubical system for L^φ, and let $\mathbb{P} \subseteq \mathrm{L}^\varphi$ be a subsystem of invariant graphs which is closed under taking unions. Suppose that $\mathbb{P}$ contains another subsystem $\mathbb{P}'$, also closed under taking unions, such that $\Sigma \cup \bigcap \mathbb{P}' \in \mathbb{P}'$ for each $\Sigma \in \mathbb{P}$. Suppose further that for each $\Sigma' \in \mathbb{P}'$ we are given a standard copy $\widetilde{Y}_{\Sigma'}$ of $\widetilde{X}_{\Sigma'}$ in $\widetilde{X}_{\bigcup \mathbb{P}}$ such that*

$$\bigcap_{\Sigma' \in \mathbb{P}'} \widetilde{Y}_{\Sigma'} \neq \emptyset$$

Then for each $\Sigma \in \mathbb{P} \smallsetminus \mathbb{P}'$ there exists a standard copy $\widetilde{Y}_\Sigma$ of $\widetilde{X}_\Sigma$ in $\widetilde{X}_{\bigcup \mathbb{P}}$ such that

$$\bigcap_{\Sigma \in \mathbb{P}} \widetilde{Y}_\Sigma \neq \emptyset$$

Proof. Let us first set some notation: we let $\Delta = \bigcap \mathbb{P}'$, and for each $\Sigma \in \mathbb{P}$ we set $\Sigma' = \Sigma \cup \Delta \in \mathbb{P}'$.

The System Intersection Axiom gives us for each $\Sigma \in \mathbb{P}$ a standard copy $\widetilde{Z}_\Sigma$ of $\widetilde{X}_\Sigma$ in $\widetilde{X}_{\bigcup \mathbb{P}}$ such that there is a point $\widetilde{b} \in \widetilde{X}_{\bigcup \mathbb{P}}$ with

$$\widetilde{b} \in \bigcap_{\Sigma \in \mathbb{P}} \widetilde{Z}_\Sigma.$$

In particular $\widetilde{b} \in \widetilde{Z}_\Delta$, and thus there exists a point $\widetilde{b}' \in \widetilde{Y}_\Delta$ such that both $\widetilde{b}$ and $\widetilde{b}'$ lie in the same standard copy $\widetilde{W}$ of $\widetilde{X}_{\mathrm{lk}_{\bigcup \mathbb{P}}(\Delta)}$ (due to the Orthogonal Axiom). Note that the Intersection Axiom guarantees that $\widetilde{b}' \in \widetilde{Y}_{\Sigma'}$ for each $\Sigma' \in \mathbb{P}'$.

We want to construct standard copies $\widetilde{Y}_\Sigma$ for each $\Sigma \in \mathbb{P}$ which contain $\widetilde{b}'$.

Let δ denote the geodesic in the complete CAT(0) space $\widetilde{X}_{\bigcup \mathbb{P}}$ connecting $\widetilde{b}$ to $\widetilde{b}'$. Since all maps $\widetilde{\iota}$ are isometric embeddings, the geodesic δ lies in $\widetilde{W}$.

Take $\Sigma \in \mathbb{P}$. Then δ connects two standard copies of $\widetilde{X}_{\Sigma'}$, namely $\widetilde{Y}_{\Sigma'}$ and $\widetilde{Z}_{\Sigma'}$, and hence it must lie in a standard copy of $\widetilde{X}_{\mathrm{st}_{\bigcup \mathbb{P}}(\Sigma')}$; this copy is unique since the link of $\mathrm{st}_{\bigcup \mathbb{P}}(\Sigma')$ in $\bigcup \mathbb{P}$ is trivial. Thus the geodesic lies in the intersection of this copy and $\widetilde{W}$, which itself is a standard copy of

$$\widetilde{X}_{\mathrm{lk}_{\bigcup \mathbb{P}}(\Delta) \cap \mathrm{st}_{\bigcup \mathbb{P}}(\Sigma')} = \widetilde{X}_{\mathrm{lk}_{\Sigma'}(\Delta) * \mathrm{lk}_{\bigcup \mathbb{P}}(\Sigma')}$$

by the Intersection Axiom. In particular, this implies that $\widetilde{b}$ lies in this standard copy. But $\widetilde{b}$ also lies in $\widetilde{Z}_\Sigma$, and hence in the unique standard copy of $\widetilde{X}_{\mathrm{st}_{\bigcup \mathbb{P}}(\Sigma)}$.

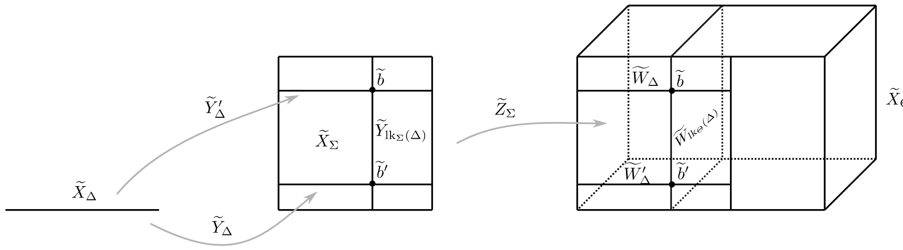

FIGURE 5.1. A schematic for the proof of Lemma 5.7

Therefore $\widetilde{b}$ lies in the intersection of this copy and the copy of $\widetilde{X}_{\mathrm{lk}_{\Sigma'}(\Delta)*\mathrm{lk}_{\bigcup\mathbb{P}}(\Sigma')}$. However

$$\mathrm{lk}_{\Sigma'}(\Delta) * \mathrm{lk}_{\bigcup\mathbb{P}}(\Sigma') \subseteq \mathrm{st}_{\bigcup\mathbb{P}}(\Sigma)$$

(since $\Sigma \subsetneq \Sigma'$ and $\Sigma' = \Sigma \cup \Delta$), and thus the Intersection Axiom implies that the copy of $\widetilde{X}_{\mathrm{lk}_{\Sigma'}(\Delta)*\mathrm{lk}_{\bigcup\mathbb{P}}(\Sigma')}$ which contains δ lies within the unique copy of $\widetilde{X}_{\mathrm{st}_{\bigcup\mathbb{P}}(\Sigma)}$. This must be also true for δ itself, and so there exists a standard copy $\widetilde{Y}_\Sigma$ of $\widetilde{X}_\Sigma$ which contains $\widetilde{b}'$, the other endpoint of δ. □

Lemma 5.6 (Matching Property)**.** *Let $\mathcal{X}$ be a cubical system for L^φ, and let $\Sigma_1, \Sigma_2 \in \mathrm{L}^\varphi$ be such that $\Theta = \Sigma_1 \cup \Sigma_2 \in \mathrm{L}^\varphi$ as well. Let $\widetilde{Y}_i$ be a given standard copy of $\widetilde{X}_{\Sigma_i}$ in $\widetilde{X}_\Theta$. Then their images intersect in a standard copy of $\widetilde{X}_\Delta$, with $\Delta = \Sigma_1 \cap \Sigma_2$.*

Proof. Let

$$\mathbb{P} = \{\Delta, \Sigma_1, \Sigma_2, \Theta\}$$

and

$$\mathbb{P}' = \{\Sigma_1, \Theta\} \subseteq \mathbb{P}$$

Note that $\mathbb{P}'$ satisfies the assumptions of the previous lemma, and so there exists a standard copy $\widetilde{Z}_{\Sigma_2}$ of $\widetilde{X}_{\Sigma_2}$ in $\widetilde{X}_\Theta$ such that $\widetilde{Y}_{\Sigma_1}$ and $\widetilde{Z}_{\Sigma_2}$ intersect.

Now there exists a geodesic δ connecting standard copies $\widetilde{Z}_{\Sigma_2}$ and $\widetilde{Y}_{\Sigma_2}$, such that δ lies in a standard copy of $\widetilde{X}_{\mathrm{lk}_\Theta(\Sigma_2)}$, and one of the endpoints of δ lies in $\widetilde{Y}_{\Sigma_1} \cap \widetilde{Z}_{\Sigma_2}$. But

$$\mathrm{lk}_\Theta(\Sigma_2) \subseteq \Sigma_1$$

and so δ lies in $\widetilde{Y}_{\Sigma_1}$ entirely (by the Intersection Axiom). Thus $\widetilde{Y}_{\Sigma_1}$ and $\widetilde{Y}_{\Sigma_2}$ intersect non-trivially (at the other endpoint of δ). □

Lemma 5.7 (Composition Property)**.** *Let $\mathcal{X}$ be a cubical system for $\mathbb{P}$, and let $\Delta, \Sigma, \Theta \in \mathbb{P}$ satisfy $\Delta \subseteq \Sigma \subseteq \Theta$. Suppose that we are given a standard copy $\widetilde{Y}_\Delta$ of $\widetilde{X}_\Delta$ in $\widetilde{X}_\Sigma$, and a standard copy $\widetilde{Z}_\Sigma$ of $\widetilde{X}_\Sigma$ in $\widetilde{X}_\Theta$. Then there exists a standard copy of $\widetilde{X}_\Delta$ in $\widetilde{X}_\Theta$, whose image is equal to the image of $\widetilde{Y}_\Delta$ in $\widetilde{Z}_\Sigma$.*

Proof. The System Intersection Axiom, Lemma 5.5, and the Intersection Axiom give us a standard copy $\widetilde{W}_\Delta$ of $\widetilde{X}_\Delta$ in $\widetilde{X}_\Theta$ contained in $\widetilde{Z}_\Sigma$, which is an image of a standard copy $\widetilde{Y}'_\Delta$ of $\widetilde{X}_\Delta$ in $\widetilde{X}_\Sigma$. Now the Orthogonal Axiom gives us a standard copy $\widetilde{Y}_{\mathrm{lk}_\Sigma(\Delta)}$ of $\widetilde{X}_{\mathrm{lk}_\Sigma(\Delta)}$ in $\widetilde{X}_\Sigma$ which intersects both $\widetilde{Y}_\Delta$ and $\widetilde{Y}'_\Delta$, the latter in a point $\widetilde{b}$.

The Orthogonal Axiom also gives us a standard copy $\widetilde{W}_{\mathrm{lk}_\Theta(\Delta)}$ of $\widetilde{X}_{\mathrm{lk}_\Theta(\Delta)}$ in $\widetilde{X}_\Theta$ which intersects $\widetilde{W}_\Delta$ at the image of point $\widetilde{b}$. The Intersection Axiom implies that $\widetilde{W}_{\mathrm{lk}_\Theta(\Delta)}$ intersects $\widetilde{Z}_\Sigma$ in a copy of $\widetilde{X}_{\mathrm{lk}_\Sigma(\Delta)}$, which is also the image of a copy of $\widetilde{X}_{\mathrm{lk}_\Sigma(\Delta)}$ in $\widetilde{X}_\Sigma$. But this copy contains $\widetilde{b}$, and two standard copies of a given complex intersect if and only if they coincide. Therefore it is equal to $\widetilde{Y}_{\mathrm{lk}_\Sigma(\Delta)}$. Thus it intersects $\widetilde{Y}'_\Delta$ in a point $\widetilde{b}'$, and the image of this point in $\widetilde{X}_\Theta$ lies in the image of $\widetilde{Y}_\Delta$ in $\widetilde{Z}_\Sigma$ and $\widetilde{W}_{\mathrm{lk}_\Theta(\Delta)}$.

The Orthogonal Axiom gives us a copy $\widetilde{W}'_\Delta$ of $\widetilde{X}_\Delta$ in $\widetilde{X}_\Theta$ which contains $\widetilde{b}'$. Hence $\widetilde{W}'_\Delta$ intersects $\widetilde{Z}_\Sigma$, and the Intersection Axiom implies that this intersection is a standard copy equal to the image of a standard copy of $\widetilde{X}_\Delta$ in $\widetilde{X}_\Sigma$. But this standard copy intersects $\widetilde{Y}_\Delta$ in $\widetilde{b}'$, and so coincides with $\widetilde{Y}_\Delta$. This finishes the proof. □

Definition 5.8. Let $\Delta \subseteq \Sigma$ be two elements of L^φ, and suppose that $\mathcal{X}'$ is a cubical system for $\mathrm{L}^\varphi_\Delta$. Let $\Delta = \Delta_1 * \cdots * \Delta_k$ be its join decomposition. We say that a cubical system $\mathcal{X}$ for $\mathrm{L}^\varphi_\Sigma$ *extends* $\mathcal{X}'$ if and only if

- when $|\Delta_i| \geqslant 2$, for every $E \in \mathrm{L}^\varphi_{\Delta_i}$ we have an H-equivariant isometry $j_E \colon X'_E \to X_E$ which preserves the markings via a map $\widetilde{j}_E$;
- when Δ_i is a singleton we have an isometry $j_{\Delta_i} \colon X'_{\Delta_i} \to X_{\Delta_i}$, which preserves the markings via a map $\widetilde{j}_{\Delta_i}$, and is H-equivariant up to homotopy;
- the maps $\widetilde{\iota}', \widetilde{\iota}$ and $\widetilde{j}$ make the obvious diagrams commute.

We say that $\mathcal{X}$ *strongly extends* $\mathcal{X}'$ if and only if all the maps j are H-equivariant.

Proposition 5.9. *Suppose that $\Gamma = \Gamma_1 * \Gamma_2$ with $\Gamma_i \in \mathrm{L}^\varphi$ for both values of i. Let $\mathcal{X}^i$ be a cubical system for $\mathrm{L}^\varphi_{\Gamma_i}$. Then there exists a cubical system $\mathcal{X}$ for L^φ which extends both $\mathcal{X}^1$ and $\mathcal{X}^2$ strongly.*

Proof. This construction is fairly straightforward. Given $\Sigma \in \mathrm{L}^\varphi$ we define $\Sigma_i = \Sigma \cap \Gamma_i$. Note that we have $\Sigma = \Sigma_1 * \Sigma_2$.

Take such a Σ. We define $\widetilde{X}_\Sigma = \widetilde{X}^1_{\Sigma_1} \times \widetilde{X}^2_{\Sigma_2}$. We mark it with the product marking, and immediately see that $\widetilde{X}_\Sigma$ and its projection X_Σ are of the form required in the definition of a cubical system.

Now let us also take $\Theta \in \mathrm{L}^\varphi$ such that $\Sigma \subseteq \Theta$. Crucially,

$$\mathrm{lk}_\Theta(\Sigma) = \mathrm{lk}_{\Theta_1}(\Sigma_1) * \mathrm{lk}_{\Theta_2}(\Sigma_2)$$

We define $\widetilde{\iota}_{\Sigma,\Theta} = \widetilde{\iota}_{\Sigma_1,\Theta_1} \times \widetilde{\iota}_{\Sigma_2,\Theta_2}$. Again, it is clear that these maps have the form required in the definition.

The four axioms are immediate; they all follow from the observation that any standard copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_\Theta$ for any $\Sigma, \Theta \in \mathrm{L}^\varphi$ with $\Sigma \subseteq \Theta$ is equal to a product of standard copies of $\widetilde{X}_{\Sigma_i}$ in $\widetilde{X}_{\Theta_i}$ for both values of i, and vice-versa – taking a product of two such copies yields a copy of the former kind. □

We will say that the cubical system $\mathcal{X}$ obtained above is the *product* of systems $\mathcal{X}^1$ and $\mathcal{X}^2$.

Lemma 5.10. *Suppose that Γ is a cone over $s \in \Gamma$, with $\{s\}, \Gamma \smallsetminus \{s\} \in \mathrm{L}^\varphi$. Let $\mathcal{X}'$ be a cubical system for $\mathrm{L}^\varphi_{\{s\}}$, and let $\mathcal{X}$ be a cubical system for L^φ which extends $\mathcal{X}'$. Then there exists a cubical system $\mathcal{X}''$ for L^φ which extends $\mathcal{X}$ and extends $\mathcal{X}'$ strongly.*

Proof. We define $\mathcal{X}''$ to be the product of $\mathcal{X}_{\Gamma \smallsetminus \{s\}}$ and $\mathcal{X}'$. It is then clear that $\mathcal{X}''$ extends $\mathcal{X}'$ strongly. To verify that $\mathcal{X}''$ extends $\mathcal{X}$ we only need to observe that $X''_{\{s\}} = X'_{\{s\}}$, which in turn is isometric to $X_{\{s\}}$ in a way H-equivariant up to homotopy since $\mathcal{X}$ extends $\mathcal{X}'$. □

Lemma 5.11. *Let $\mathcal{X}$ be a cubical system for L^φ, and suppose that the action φ is link-preserving, and its domain H is finite. Suppose that we have a group extension*

$$A_\Gamma \to \overline{H} \to H$$

yielding the given action φ. Then there exists a cubical system $\mathcal{Y}$ for L^φ which extends $\mathcal{X}$, and such that the extension of A_Γ by H given by the action of H on X_Γ is isomorphic to $\overline{H}$ (as an extension).

Proof. Let $\Gamma = \Gamma_1 * \cdots * \Gamma_m$ be a join decomposition of Γ, and suppose that Γ_i is a singleton if and only if $i \leqslant k$. Then $E = \Gamma_{i+1} * \cdots * \Gamma_m \in \mathrm{L}^\varphi$, since φ is link-preserving, and so we have a cubical system $\mathcal{Y}^1 = \mathcal{X}_E$ for L^φ_E.

We also have $\Delta = \Gamma_1 * \cdots * \Gamma_k \in \mathrm{L}^\varphi$, and thus the extension $\overline{H}$ yields an extension

$$A_\Delta \to \overline{A_\Delta} \to H$$

We easily build a cubical system $\mathcal{Y}^2$ for $\mathrm{L}^\varphi_\Delta$ such that the action of H on Y^2_Δ gives $\overline{A_\Delta}$ – it is enough to do it for Δ being a singleton, in which case we only need to observe that any finite extension of $\mathbb{Z}$ acts on $\mathbb{R}$ by isometries preserving the integer points and without a global fixed point, and taking the quotient by $\mathbb{Z}$ yields an action of H on a circle which gives the desired extension.

Note that $\mathcal{Y}^2$ extends $\mathcal{X}_\Delta$. Proposition 5.9 finishes the proof. □

6. Cubical systems for free products

In this section we use Relative Nielsen Realisation (Theorem 3.4) to build cubical systems for RAAGs whose defininig graph Γ is disconnected. The construction will assume that we have already built cubical systems for some unions of connected components of Γ.

Proposition 6.1. *Suppose that Γ is the disjoint union*

$$\Gamma = \Sigma_1 \sqcup \cdots \sqcup \Sigma_n \sqcup \Theta$$

where $\Sigma_i \in \mathrm{L}^\varphi$, each Σ_i is a union of connected components, and Θ is discrete. Let $\mathcal{X}^i$ be a cubical system for $\mathrm{L}^\varphi_{\Sigma_i}$ for each i. Then there exists a cubical system $\mathcal{X}$ for $\mathrm{L}^\varphi_{\Sigma_1} \cup \cdots \cup \mathrm{L}^\varphi_{\Sigma_n} \cup \{\Gamma\}$ extending each $\mathcal{X}_i$.

Moreover, the (unique) standard copies in $\widetilde{X}_\Gamma$ of $\widetilde{X}_{\Sigma_i}$ and $\widetilde{X}_{\Sigma_j}$ are disjoint when $i \neq j$, unless $\Gamma = \Sigma_1 \sqcup \Sigma_2$, in which case their standard copies intersect in a single point, and the union of their projections covers X_Γ.

Proof. If each Σ_i is empty, then the classical Nielsen Realisation for graphs yields the result. We may thus assume that Σ_1 is not empty. We may also assume that $\Gamma \neq \Sigma_1$. Therefore we may assume that A_Γ has no centre, and hence the action $H \to \mathrm{Out}(A_\Gamma)$ gives us an extension

$$A_\Gamma \to \overline{H} \to H$$

This extension in turn gives us extensions

$$A_{\Sigma_i} \to \overline{A_{\Sigma_i}} \to H$$

for each i.

We use Lemma 5.11 and modify each $\mathcal{X}^i$ so that these extensions agree with the ones given by the action of H on $X^i_{\Sigma_i}$.

We now apply Theorem 3.4, together with Remark 3.5, using the cube complexes $X^i_{\Sigma_i}$ as input. This way we construct the cube complex X_Γ. Now we define the cubical system $\mathcal{X}$.

Given $\Sigma \in \mathrm{L}^\varphi_{\Sigma_i}$ we define $X_\Sigma = X^i_\Sigma$. We have already defined X_Γ. Since the subgraphs Σ_i, Σ_j are disjoint when $i \neq j$, there are no choices involved in this definition.

Given $\Sigma, \Sigma' \in \mathrm{L}^{\varphi}_{\Sigma_i}$ we set $\widetilde{\iota}_{\Sigma,\Sigma'}$ to be the corresponding map in $\mathcal{X}^i$. Given $\Sigma \in \mathrm{L}^{\varphi}_{\Sigma_i}$ we define

$$\widetilde{\iota}_{\Sigma,\Gamma} = \widetilde{\iota}_{\Sigma_i,\Gamma} \circ \widetilde{\iota}_{\Sigma,\Sigma_i}$$

This map is of the desired kind since $\mathrm{lk}(\Sigma) \subseteq \Sigma_i$, as this last subgraph is a union of components of Γ. We also set $\widetilde{\iota}_{\Gamma,\Gamma}$ to be the identity.

What remains is the verification of the four axioms.

Product and Orthogonal Axioms. Suppose that $\Sigma' = \mathrm{st}_{\Sigma'}(\Sigma)$, with

$$\Sigma, \Sigma' \in \mathrm{L} = \mathrm{L}^{\varphi}_{\Sigma_1} \cup \cdots \cup \mathrm{L}^{\varphi}_{\Sigma_n} \cup \{\Gamma\}$$

Then either $\Sigma = \Sigma'$, which is the trivial case, or $\Sigma \subseteq \Sigma_i$ for some i. But then also $\Sigma' \subseteq \Sigma_i$, since Σ_i is a union of connected components. In this latter case we only need the Product or Orthogonal Axiom in $\mathcal{X}^i$.

Intersection Axiom. Take $\Sigma, \Sigma', \Omega \in \mathrm{L}$ such that $\Sigma \subseteq \Omega$ and $\Sigma' \subseteq \Omega$, and let $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ be standard copies of, respectively, $\widetilde{X}_\Sigma$ and $\widetilde{X}_{\Sigma'}$ in $\widetilde{X}_\Omega$ with non-empty intersection. We need to show that the intersection is the image of a standard copy of $\Sigma \cap \Sigma'$ in each.

The situation becomes trivial when $\Sigma = \Omega$ or $\Sigma' = \Omega$, so let us assume neither of these situations occurs.

If there exists an i such that $\Sigma, \Sigma' \subseteq \Sigma_i$, then either $\Omega \subseteq \Sigma_i$, in which case we use the Intersection Axiom of $\mathcal{X}^i$, or $\Omega = \Gamma$, in which case we use the Intersection Axiom for the triple $\Sigma, \Sigma', \Sigma_i$, noting that the standard copies $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ are just images of standard respective standard copies in $\widetilde{X}_{\Sigma_i}$.

The remaining case occurs when $\Sigma \subseteq \Sigma_i$ and $\Sigma' \subseteq \Sigma_j$ for $i \neq j$. Then $\Sigma \cap \Sigma' = \emptyset$, and so we need to show that $\widetilde{Y}_\Sigma \cap \widetilde{Y}_{\Sigma'}$ is a single point. Since the standard copies intersect non-trivially, it means that the images of Y_i and Y_j intersected. Thus we must have had $\Gamma = \Sigma_i \sqcup \Sigma_j$, and $\widetilde{Y}_\Sigma \cap \widetilde{Y}_{\Sigma'}$ is precisely the unique point at which the images of $\widetilde{Y}_i$ and $\widetilde{Y}_j$ intersect.

System Intersection Axiom. Take a subsystem $\mathbb{P} \subseteq \mathrm{L}$ closed under taking unions. Suppose that we have $\Sigma, \Sigma' \in \mathbb{P}$ with $\Sigma \subseteq \Sigma_i$ and $\Sigma' \subseteq \Sigma_j$ with $i \neq j$. Then $\Sigma \cup \Sigma' \in \mathbb{P} \subseteq \mathrm{L}$, and so $\Sigma \cup \Sigma' = \Gamma$. An analogous reasoning immediately implies that $\mathbb{P} = \{\emptyset, \Sigma_i, \Sigma_j, \Gamma\}$. Now the last three subgraphs have unique corresponding standard copies, and these intersect at a single point; we take this point to be the standard copy of $\widetilde{X}_\emptyset$ and the argument is finished.

If all subgraphs in $\mathbb{P} \smallsetminus \{\Gamma\}$ lie in Σ_i for some i, then we are done by the System Intersection Axiom of $\mathcal{X}^i$ applied to $\mathbb{P} \smallsetminus \{\Gamma\}$ – either this covers all of $\mathbb{P}$, or we need to observe that standard copies in $\widetilde{X}_\Gamma$ are images of copies in $\widetilde{X}_{\Sigma_i}$. □

7. Gluing

In this section we deal with the situation in which we are given cube complexes realising the induced action on A_Σ and A_Δ, where $\Sigma \cap \Delta \neq \emptyset$, and we build a cube complex realising the action on $A_{\Sigma \cup \Delta}$.

Suppose that $\varphi \colon H \to \mathrm{Out}(A_\Gamma)$ is link-preserving, and that H is finite. Let $\Sigma, \Theta \in \mathrm{L}^\varphi$ be such that $\Sigma \cup \Theta = \Gamma$. Take $E = \Sigma \cap \Theta$. Let us set $E' = E \smallsetminus Z(E)$, and $Z(E) = \{s_1, \ldots, s_k\}$. Since φ is link-preserving, we have $E' \in \mathrm{L}^\varphi$ and $\{s_i\} \in \mathrm{L}^\varphi$ for each i.

Suppose that we have cubical systems $\mathcal{X}$ and $\mathcal{X}'$ for $\mathrm{L}^{\varphi}_{\Sigma}$ and $\mathrm{L}^{\varphi}_{\Theta}$ respectively, such that $\mathcal{X}'$ extends $\mathcal{X}_E$.

The main goal of this section is to show that (under mild assumptions) one can equivariantly glue X_Σ to X'_Θ so that the result realises the correct action $\varphi : H \to \mathrm{Out}(A_\Gamma)$. This will be done in two steps; first showing that the loops

X_{s_i} and X'_{s_i} are actually equal (and hence that $\mathcal{X}'$ extends $\mathcal{X}_E$ strongly), and then constructing a gluing.

We will repeatedly use the following construction.

Definition 7.1. Given a cube complex X_Γ realising an action

$$\varphi\colon H \to \mathrm{Out}(A_\Gamma)$$

we say that an element $h_p \in \mathrm{Aut}(A_\Gamma)$ is a *geometric representative* of $h \in H$ if and only if it is obtained by the following procedure: take a basepoint $\widetilde{p} \in \widetilde{X}_\Gamma$ with a projection $p \in X_\Gamma$, and a path γ from p to $h.p$. The choice of $\widetilde{p}$ induces an identification $\pi(X_\Gamma, p) = A_\Gamma$. We now take $h_p \in \mathrm{Aut}(\pi_1(X_\Gamma, p))$ to be the automorphism induced on the fundamental group by first applying h to X_Γ, and then pushing the basepoint back to p via γ.

Suppose that H acts on a graph (without leaves) of rank 1 (i.e. on a subdivided circle). Let us fix an orientation on the circle.

Definition 7.2. Given an element $h \in H$ we say that it *flips* the circle if and only if it reverses the circle's orientation; otherwise we say that it *rotates* the circle. In the latter case we say that it rotates by k if and only if the simple path from some vertex to its image under h, going along the orientation of the circle, has combinatorial length k.

Lemma 7.3. *Suppose that $\Sigma \neq \mathrm{st}_\Sigma(s_i)$ and $\Theta \neq \mathrm{st}_\Theta(s_i)$ for some i. Then $X_{\{s_i\}}$ and $X'_{\{s_i\}}$ are H-equivariantly isometric.*

Proof. To simplify notation set $s = s_i$, $Y = X_{s_i}$ and $Z = X'_{s_i}$. Note that Y and Z are isometric. Let us subdivide the edges of Y and Z so that both actions of H are combinatorial, and so that the two loops can be made combinatorially isomorphic.

Let m denote the number of vertices in the subdivided loop Y. Fix an orientation on both Y and Z so that going around the loops once in the positive direction yields $s \in A_\Gamma$.

We first focus on those $h \in H$ which map the conjugacy class of $s \in A_\Gamma$ to itself. Then h acts on Y and Z as a rotation. We claim that the two actions of h rotate by the same number of vertices.

Consider a representative $h_0 \in \mathrm{Aut}(A_\Gamma)$ of $\varphi(h)$ which preserves s. For any such representative we have $h_0^{\mathrm{ord}(h)}$ equal to a conjugation which fixes s. Hence

$$h_0^{\mathrm{ord}(h)} = c(s^{K(h_0)} t_0)$$

where

$$t_0 \in A_{\mathrm{lk}(s)}$$

We know that $s \notin Z(\Gamma)$ (otherwise $\Sigma = \mathrm{st}_\Sigma(s)$ which contradicts our assumption). Thus, the integer $K(h_0)$ is unique.

Since $\mathrm{lk}(s) \in \mathrm{L}^\varphi$, the subgroup $h_0(A_{\mathrm{lk}(s)})$ is conjugate to $A_{\mathrm{lk}(s)}$. But the former subgroup must centralise $h_0(s) = s$, and so

$$h_0(A_{\mathrm{lk}(s)}) \leqslant A_{\mathrm{st}(s)}$$

The only subgroup of $A_{\mathrm{st}(s)} = A_s \times A_{\mathrm{lk}(s)}$ conjugate to $A_{\mathrm{lk}(s)}$ is $A_{\mathrm{lk}(s)}$ itself, and therefore $h_0(A_{\mathrm{lk}(s)}) = A_{\mathrm{lk}(s)}$.

Let h_1 and h_2 be representatives of $\varphi(h)$. There exists a unique integer l such that

$$h_2 = c(s^l t) \circ h_1$$

with $t \in A_{\mathrm{lk}(s)}$. Note that

$$c(x) \circ h_1 = h_1 \circ c(h_1^{-1}(x))$$

and so, in particular, using $h_1(s) = s$ and $h_1(A_{\mathrm{lk}(s)}) = A_{\mathrm{lk}(s)}$, we get

$$(c(s^l t) \circ h_1)^{\mathrm{ord}(h)} = c(s^{\mathrm{ord}(h)l} t') \circ h_1^{\mathrm{ord}(h)}$$

where $t' \in A_{\mathrm{lk}(s)}$. Thus

$$\begin{aligned} c(s^{K(h_2)} t_2) &= h_2^{\mathrm{ord}(h)} \\ &= (c(s^l t) \circ h_1)^{\mathrm{ord}(h)} \\ &= c(s^{\mathrm{ord}(h)l} t') \circ h_1^{\mathrm{ord}(h)} \\ &= c(s^{\mathrm{ord}(h)l + K(h_1)} t' t_1) \end{aligned}$$

This shows that $K(h_1) \bmod \mathrm{ord}(h)$ is independent of the representative, and so we can define $K(h) \in \mathbb{Z}/\mathrm{ord}(h)\mathbb{Z}$ in the obvious way. This algebraic invariant will be the main tool in showing that Y and Z are H-equivariantly isometric.

Fix a basepoint p in X_Σ lying in $\mathrm{im}(\iota_{E,\Sigma})$, and a basepoint q in X_Θ lying in $\mathrm{im}(\iota_{E,\Theta})$.

Let $h_p \in \mathrm{Aut}(A_\Sigma)$ be the geometric representative of h (using the action of H on X_Σ), obtained by taking the basepoint p and a path γ inside

$$\mathrm{im}(\iota_{E,\Sigma}) \cong \prod_{i=1}^{k} X_{\{s_i\}} \times X_{E'}$$

which first travels orthogonally to $Y = X_{\{s_i\}}$, and then along the copy of Y containing p (in the negative direction).

If h rotates Y by μ vertices (in the positive direction), then $(h_p)^{\mathrm{ord}(h)}$ is equal to the conjugation by $s^{\mathrm{ord}(h)\mu/m} t$ for some $t \in A_{\mathrm{lk}_\Sigma(s)}$. Since $\Sigma \neq \mathrm{st}_\Sigma(s)$, the number $\mathrm{ord}(h)\mu/m$ is unique. Hence, by taking any representative of $\varphi(h)$ in $\mathrm{Aut}(A_\Gamma)$ which restricts to h_p on A_Σ, we see that

$$K(h) = \mathrm{ord}(h)\mu/m \bmod \mathrm{ord}(h)$$

Now we define h_q in the analogous manner using X_Θ instead of X_Σ. Since h_p and h_q represent the same element h, the computation above shows that they rotate by the same number of vertices. We have thus dealt with elements $h \in H$ which fix the conjugacy class of s.

If h maps the conjugacy class of s to the conjugacy class of s^{-1}, then h must flip both Y and Z, and therefore must have two fixed points on each loop. If another element $g \in H$ flips Y, then hg rotates Y, and by the above rotates Z by the same number of vertices. This implies that the fixed points of h and g on Y differ by the same number of vertices as the respective fixed points on Z.

Hence there exists an identification between Y and Z which is H-equivariant. □

Now suppose that $\mathcal{X}'$ extends $\mathcal{X}_E$ strongly. Suppose further that we have fixed standard copies of $\widetilde{X}_E$: one in $\widetilde{X}_\Sigma$, called $\widetilde{P}$, and one in $\widetilde{X}'_\Theta$, called $\widetilde{Q}$. We can from a cube complex marked by A_Γ from $\widetilde{X}_\Sigma$ and $\widetilde{X}'_\Theta$ by gluing $\widetilde{P}$ and $\widetilde{Q}$. Note that this is in general not unique, as there might be more than one marking-respecting isometry of $\widetilde{P}$ and $\widetilde{Q}$ such that the projections P and Q become H-equivariantly isomorphic.

Let $\widetilde{Y}$ denote the glued-up complex, and let Y denote its projection. Our gluing gives us an action of H on the projection Y. This induces an action $H \to \mathrm{Out}(A_\Gamma)$ in the obvious way; but this action is in general not equal to φ. We are now going to measure the difference of these two actions.

Let us choose a point $\widetilde{p} \in \widetilde{P}$ as a basepoint. For each $h \in H$ we choose a path $\gamma(h)$ in P connecting p to $h.p$. Since P and Q are standard copies of the same complex, they are isomorphic via a fixed isomorphism. This gives us a copy of $\widetilde{p}$ and $\gamma(h)$ in $\widetilde{Q}$ and Q respectively; let us denote the former by $\widetilde{q}$ and the latter by $\gamma'(h)$. The points p and q are naturally points in Y.

Now let h_p and h_q denote the respective geometric representatives of h. The former restricts to the same automorphism as $\varphi(h)$ on A_Σ, the latter on A_Θ. They represent the same outer automorphism, and agree on A_E. Thus we have

$$h_p h_q^{-1} = c(x(h))$$

for some $x(h) \in C(A_E)$.

Definition 7.4. We say that the gluing above is *faulty within* $G \leqslant C(A_E)$ if and only if $x(h) \in G$ for all h (with our choices of $\widetilde{p}$ and $\gamma(h)$).

Proposition 7.5 (Gluing Lemma)**.** *Suppose that* $\widetilde{Y}$ *is faulty within* $Z(A_E)$. *Then there exists another gluing as above,* $\widetilde{X}$, *such that its projection* X *realises* φ.

Proof. First let us note that when $Z(\Gamma) = Z(E)$ then all the conjugations $c(x(h))$ are trivial, and $\widetilde{Y}$ is already as desired. We will henceforth assume that $Z(\Gamma) \neq Z(E)$.

Take $h \in H$. By assumption, the gluing $\widetilde{Y}$ gives us geometric representatives h_p and h_q such that

$$h_p h_q^{-1} = c(x(h))$$

with $x(h) \in Z(A_E) = A_{Z(E)}$.

We assume that $x(h) \in A_{Z(E) \smallsetminus Z(\Gamma)}$ unless $x(h)$ is the identity; we can always do this since conjugating by elements in $A_{Z(\Gamma)} = Z(A_\Gamma)$ is trivial. Now we define $h'_p \in \mathrm{Aut}(A_\Gamma)$ so that $h'_p h_p^{-1}$ is a conjugation by an element of $Z(A_\Sigma)$ and that $h'_p h_q^{-1}$ is equal to a conjugation by an element in $A_{Z(E) \smallsetminus (Z(\Gamma) \cup Z(\Sigma))}$; we further define $h'_q \in \mathrm{Aut}(A_\Gamma)$ so that $h'_q h_q^{-1}$ is a conjugation by an element of $Z(A_\Theta)$ and that

$$h'_p h'^{-1}_q = c(x'(h))$$

with $x'(h) \in A_{Z(E) \smallsetminus (Z(\Gamma) \cup Z(\Sigma) \cup Z(\Theta))}$. Since $Z(\Sigma) \cap Z(\Theta) = Z(\Gamma)$, the elements h'_p and h'_q are unique.

Note that h_p and h'_p are identical when restricted to A_Σ; the analogous statement holds for h_q and h'_q restricted to A_Θ.

Consider now ${h'_p}^{\mathrm{ord}(h)}$. It is equal to a conjugation $c(y_p)$ where

$$y_p \in N(A_\Sigma)$$

by construction. Let $y_q \in N(A_\Theta)$ be the corresponding element for h'_q. Now

$$c(y_p) = {h'_p}^{\mathrm{ord}(h)} = \left(c(x'(h))h'_q\right)^{\mathrm{ord}(h)} = c(x'')\,{h'_q}^{\mathrm{ord}(h)} = c(x'')c(y_q)$$

where

$$x'' = \prod_{i=0}^{\mathrm{ord}(h)} (h'_q)^i(x'(h)) \in A_{Z(E) \smallsetminus (Z(\Sigma) \cup Z(\Theta))}$$

since $Z(E) \smallsetminus (Z(\Sigma) \cup Z(\Theta)) \in \mathrm{L}^\varphi$ and $h'_q(A_E) = A_E$ by construction.

The element y_p is determined up to $C(A_\Sigma)$ by its action by conjugation on A_Σ. Here however we immediately see that $c(y_p)$ is equal to conjugation by the element given by the loop obtained from concatenating images of our path $\gamma(h)$ under successive iterations of h.

We repeat the argument for y_q and conclude that $y_p = y_q$ up to $C(A_\Sigma)C(A_\Theta)$. But we have already shown that they differ by $x'' \in A_{Z(E) \smallsetminus (Z(\Sigma) \cup Z(\Theta))}$ up to

$Z(A_\Gamma)$. Hence $c(x'') = 1$ as $Z(E)$ intersects $\mathrm{lk}(\Sigma)$ and $\mathrm{lk}(\Theta)$ trivially. Thus $x'' \in Z(A_\Gamma) \cap A_{Z(E) \smallsetminus (Z(\Sigma) \cup Z(\Theta))} = \{1\}$. We obtain

$$1 = x'' = \prod_{i=0}^{\mathrm{ord}(h)} (h'_q)^i (x'(h))$$

and so $x'(h)$ must lie in the subgroup of $A_{Z(E)}$ generated by all vertices $s_i \in Z(E)$ such that h flips the corresponding loop $P_i = X_{s_i}$, since any other generator s_j satisfies

$$\prod_{i=0}^{\mathrm{ord}(h)} (h'_q)^i (s_j) = {s_j}^{\mathrm{ord}(h)} \neq 1$$

The purpose of the proof so far was exactly to establish that $x'(h)$ lies in the subgroup of $A_{Z(E)}$ generated by all vertices $s_i \in Z(E)$ such that h flips the loop P_i.

Now we are going to construct a whole family of gluings, and show that one of them is as desired.

To analyse the situation we need to look more closely at $\widetilde{P} = \widetilde{Q}$ (with the equality coming from the fact that $\mathcal{X}'$ extends $\mathcal{X}_E$ strongly). By the Product Axiom, we have $\widetilde{P} = \widetilde{P}_0 \times \widetilde{P}_1 \times \cdots \times \widetilde{P}_k$, with $\widetilde{P}_0 = \widetilde{X}_{E'}$, and $\widetilde{P}_i = \widetilde{X}_{s_i}$ for $i \geqslant 1$. Our basepoint $\widetilde{p}$ satisfies

$$\widetilde{p} = (\widetilde{p}_0, \ldots, \widetilde{p}_k) \in \widetilde{P}_0 \times \cdots \times \widetilde{P}_k$$

Let $\widetilde{q} = (\widetilde{q}_0, \ldots, \widetilde{q}_k)$ be the corresponding expression for $\widetilde{q}$.

We construct complexes from $\widetilde{X}_\Sigma$ and $\widetilde{X}'_\Theta$ by gluing $\widetilde{P}$ and $\widetilde{Q}$ in a way respecting the markings, and so that the projections P and Q are glued in an H-equivariant fashion. The resulting space is determined by the relative position of $\widetilde{p}$ and $\widetilde{q}$, now both seen as points in the glued-up complex (so in particular they do not need to coincide). We glue so that the images of $\widetilde{p}$ and $\widetilde{q}$ coincide if we project $\widetilde{X}_E$ onto $\widetilde{X}_{E'}$ – this is in fact forced on us since $\widetilde{X}_{E'}$ can be glued to itself only in one way. Hence any such gluing will give us a geodesic from $\widetilde{p}$ to $\widetilde{q}$, which will lie in a subcomplex of $\widetilde{X}_E$ isomorphic to $\widetilde{P}_1 \times \cdots \times \widetilde{P}_k$, and hence isometric to a Euclidean space.

We start by taking $x \in A_{Z(E)}$; we form a complex $\widetilde{X}^x$ (with projection X^x as usual) by gluing as above, so that the geodesic $\widetilde{\delta}$ we just discussed is such that if we extend it to twice its length (which is possible in a Euclidean space), still starting at $\widetilde{p}$, the projection in X^x becomes a closed loop equal to $x \in \pi_1(X^x, p)$.

Note that given $x \in Z(E)$ such a geodesic (and hence gluing) always exist: the Euclidean space we discussed is marked by $A_{Z(E)}$, and so the point $x(\widetilde{p})$ lies therein. We take the unique geodesic from $\widetilde{p}$ to $x(\widetilde{p})$, cut it in half, and declare the first half to be $\widetilde{\delta}$.

Let us now calculate the action on (conjugacy classes in) A_Γ induced from the action of h on X^x. The element h maps the local geodesic δ (the projection of $\widetilde{\delta}$) to a local geodesic $h.\delta$ connecting $h.p$ to $h.q$ such that the loop obtained by following $\gamma(h)$ (starting at p), $h.\delta$, the inverse of $\gamma'(h)$, and the inverse of δ, gives an element

$$\overline{x}^h \in \pi_1(P_1 \times \cdots \times P_k, (p_1, \ldots, p_k))$$

which is the projection of x onto the subgroup of $Z(E)$ generated by all the generators s_i such that $h'_p(s_i) = s_i^{-1}$. Hence the action of h on X^x, followed by pushing the basepoint p via $\gamma(h)$ as before, gives us an automorphism equal to h'_p on the subgroup $\pi_1(X_\Sigma, p) = A_\Sigma$, and to $c(\overline{x}^h)h'_q$ on the subgroup $\pi_1(X_\Theta, p) = A_\Theta$.

It follows that the action of h on $X^{x'(h)}$ is the correct one: by the observation above, $x'(h)$ lies in the subgroup of $A_{Z(E)}$ generated by all the vertices in $Z(E)$

which are mapped to their own inverse under h'_p. Hence $\overline{x'(h)}^h = x'(h)$, and it is enough to observe that $c(x'(h))h'_q = h_p$.

We now need to specify a single x that will work for all elements $h \in H$; we will denote such an x by $x'(H)$. If there is a vertex in $Z(E)$ which is preserved by all elements h'_p (for all $h \in H$), then we set the corresponding coordinate of $x'(H)$ to 0. If a generator is mapped to its inverse by some h'_p, then we set the corresponding coordinate of $x'(H)$ to be equal to the relevant coordinate in $x'(h)$.

Since this definition involves making choices (of elements h that flip a generator), we need to make sure that we indeed obtain the desired action. Let $g \in H$ be any element. We need to confirm that $\overline{x'(g)}^g = \overline{x'(H)}^g$.

Suppose that this is not the case; then there exists a generator s_i such that $g'(s_i) = {s_i}^{-1}$, and $x'(g)$ and $x'(H)$ disagree on the s_i-coordinate. This means that the geometric representative g'_p obtained from the action on $X^{x'(H)}$ is not a representative of $\varphi(g)$. To make it such a representative we need to postcompose it with a partial conjugation of A_Θ by $\overline{x'(g)}^g \left(\overline{x'(H)}^g\right)^{-1}$, which has a non-trivial s_i-coordinate.

The construction of $x'(H)$ required an element $h \in H$ such that h'_p flips the loop P_i. We have $g = fh$ with f acting trivially on the conjugacy class of s_i. Hence, even though f might not act correctly on $X^{x'(H)}$, the geometric representative f'_p can be made into a representative of $\varphi(f)$ by postcomposing it with a partial conjugation of A_Θ by an element of $Z(E)$ with a trivial s_i-coordinate. The exact same statement is true for h. Using the fact that f'_p maps each s_j either to itself or its inverse, we deduce that $f'_p h'_p$ can be made into a representative of $\varphi(fh) = \varphi(g)$ by postcomposing it with a partial conjugation of A_Θ by an element of $Z(E)$ with a trivial s_i-coordinate. But $f'_p h'_p$ differs from $(fh)'_p = g'_p$ by a conjugation, and hence $\overline{x'(g)}^g \left(\overline{x'(H)}^g\right)^{-1}$ cannot have a non-trivial s_i-coordinate (as $\Gamma \neq Z(E)$). □

In particular we have

Corollary 7.6. *If* $\mathrm{lk}(E) = \emptyset$ *then there exists a complex obtained from* $\widetilde{X}_\Sigma$ *and* $\widetilde{X}'_\Theta$ *by gluing* $\widetilde{P}$ *to* $\widetilde{Q}$ *such that its projection realises* φ.

Proof. When $\mathrm{lk}(E) = \emptyset$ we have $C(A_E) = Z(A_E)$. Hence any gluing will be faulty within $Z(A_E)$, and then the existence of a desired gluing follows from the previous proposition. □

Let us record here a (very standard and) very useful result.

Lemma 7.7. *Let* X *be a complete NPC space realising the action of a finite group* $\varphi\colon H \to \mathrm{Out}(A)$, *where* A *is a group. Suppose that there exists a lift* $\psi\colon H \to \mathrm{Aut}(A)$ *of* φ. *Then the action of* H *on* X *has a fixed point* r; *moreover there exists a lift* $\widetilde{r} \in \widetilde{X}$ *of* r *such that under the identification* $\pi_1(X, r) = A$ *induced by choosing* $\widetilde{r}$ *as a basepoint, the induced action oh* H *on* $\pi_1(X, r)$ *is* ψ.

Proof. Consider the action $A \curvearrowright \widetilde{X}$ on the universal cover by deck transformations. The action of each $h \in H$ can be lifted from X to $\widetilde{X}$. The fact that $H \to \mathrm{Out}(A)$ lifts to $H \to \mathrm{Aut}(A)$ tells us that we can lift each h in such a way that in fact the group H acts on $\widetilde{X}$. But $\widetilde{X}$ is CAT(0), and hence every finite group of isometries has a fixed point. Thus in particular H does; let $\widetilde{r}$ denote this fixed point. Letting r be the projection of $\widetilde{r}$ finishes the proof. □

8. Proof of the main theorem: preliminaries

Our proof of the main theorem has an inductive character – we will induct on the dimension of the defining graph Γ. To emphasize this (and to simplify statements) let us introduce the following definitions.

Definition 8.1. We say that *Relative Nielsen Realisation holds* for an action $\varphi\colon H \to \mathrm{Out}(A_\Gamma)$ if and only if given $\Delta \in \mathrm{L}^\varphi$ and any cubical system $\mathcal{Y}$ for $\mathrm{L}^\varphi_\Delta$, there exists a cubical system $\mathcal{X}$ for L^φ extending $\mathcal{Y}$.

The rest of the paper is devoted to proving

Theorem 8.2. *Relative Nielsen Realisation holds for all link-preserving actions $\varphi\colon H \to \mathrm{Out}(A_\Gamma)$ with H finite.*

Before proceeding to the proof, let us record the following corollaries.

Corollary 8.3. *Let $\varphi\colon H \to \mathrm{Out}^0(A_\Gamma)$ be a homomorphism with a finite domain. Then there exists a metric NPC cube complex X realising φ, provided that for any vertex $v \in \Gamma$, its link $\mathrm{lk}(v)$ is not a cone.*

Proof. By Corollary 4.7 φ is link-preserving. We take $X = X_\Gamma \in \mathcal{X}$ obtained by an application of the previous theorem. □

Note that in particular the statement above holds for all $\varphi\colon H \to \mathrm{Out}(A_\Gamma)$, provided that in addition Γ has no symmetries.

Corollary 8.4. *Let $\varphi\colon H \to \mathcal{U}^0(A_\Gamma)$ be a homomorphism with a finite domain. Then there exists a metric NPC cube complex X realising φ.*

Proof. By Lemma 4.6 the action φ is link-preserving. We take $X = X_\Gamma \in \mathcal{X}$ obtained by an application of the previous theorem. □

Lemma 8.5. *Assume that φ is link-preserving, and that Relative Nielsen Realisation holds for all link-preserving actions $\psi\colon H \to \mathrm{Out}(A_\Sigma)$ with $\dim \Sigma < \dim \Gamma$. Suppose that $\Gamma = \Delta * (E \cup \Theta)$, where Δ and Θ are non-empty. Suppose that $\Delta * E \in \mathrm{L}^\varphi$, and that we are given a cubical system $\mathcal{X}'$ for $\mathrm{L}^\varphi_{\Delta * E}$. Then there exists a cubical system $\mathcal{X}$ for L^φ extending $\mathcal{X}'$, which furthermore extends $\mathcal{X}'_\Delta$ strongly.*

Proof. Since L^φ contains all links, we have $E \cup \Theta \in \mathrm{L}^\varphi$. Note that $E \cup \Theta$ has lower dimension than Γ. Since

$$E = (E \cup \Theta) \cap (\Delta * E) \in \mathrm{L}^\varphi$$

we can apply the assumption to the induced action on $A_{E\cup\Theta}$, and obtain a cubical system $\mathcal{X}_{E\cup\Theta}$ extending $\mathcal{X}'_E$, the subsystem of $\mathcal{X}'$ corresponding to L^φ_E. This last system also contains a subsystem $\mathcal{X}'_\Delta$ corresponding to L_Δ.

We now define $\mathcal{X}$ to be the product of $\mathcal{X}_{E\cup\Theta}$ and $\mathcal{X}'_\Delta$ (compare Proposition 5.9). □

Proof of Theorem 8.2 . We proceed by induction on the dimension of the defining graph Γ. Since we need to prove Relative Nielsen Realisation, let us fix $\Xi \in \mathrm{L}^\varphi$, and a cubical system $\mathcal{Y}$ for L^φ_Ξ. Our aim is to construct a cubical system $\mathcal{X}$ for L^φ extending $\mathcal{Y}$. If $\Xi = \Gamma$ then there is nothing to prove, so let us assume that $\Xi \subset \Gamma$ is a proper subgraph.

First we consider the case of Γ being a join $\Gamma_1 * \Gamma_2$ for some non-empty subgraphs Γ_1 and Γ_2. The dimension of Γ_1 is strictly smaller than that of Γ, and

$$\Gamma_1 = \mathrm{lk}(\Gamma_2) \in \mathrm{L}^\varphi$$

since φ is link-preserving. Thus, by the inductive assumption, Relative Nielsen Realisation holds for Γ_1 and yields a cubical system $\mathcal{X}_{\Gamma_1}$ for $\mathrm{L}^\varphi_{\Gamma_1}$ extending the system $\mathcal{Y}_{\Xi\cap\Gamma_1}$. The same applies to Γ_2 and yields a cubical system $\mathrm{L}^{\varphi}_{\Gamma_2}$ extending $\mathcal{Y}_{\Xi\cap\Gamma_2}$. In this setting Proposition 5.9 yields a cubical system $\mathcal{X}$ as required; we only need to observe that $\mathcal{Y}$ extends the product of $\mathcal{Y}_{\Xi\cap\Gamma_1}$ and $\mathcal{Y}_{\Xi\cap\Gamma_2}$.

For the rest of the proof we assume that Γ is not a join. We now proceed by induction on the *depth* k of Γ. The depth is the length k of a maximal chain of proper inclusions

$$\emptyset = \Sigma_0 \subset \Sigma_1 \subset \cdots \subset \Sigma_k = \Gamma$$

where each $\Sigma_i \in \mathrm{L}^{\varphi}$.

If $k = 1$ then $\mathrm{L}^{\varphi} = \{\emptyset, \Gamma\}$. This in particular implies that Γ is discrete: if it is not, let v be a vertex of Γ with non-empty link. Now $\widehat{\mathrm{st}}(v) \in \mathrm{L}^{\varphi}$ is a join of two non-empty graphs, and hence cannot be equal to Γ, which is not a join. But $\widehat{\mathrm{st}}(v) \neq \emptyset$, which contradicts the assumption on depth. The depth k being 1 also implies that $\Xi = \emptyset$. Thus our theorem is reduced to the classical Nielsen Realisation for free groups [**Cul, Khr1, Zim1, HOP**]. Note that we require our graphs to be leaf-free, so we might need to prune the leaves.

Suppose that Γ is of depth $k \geqslant 2$, and that the theorem holds for all graphs of smaller depth. Let Γ' denote the maximal (with respect to inclusion) proper subgraph of Γ such that

$$\Xi \subseteq \Gamma' \in \mathrm{L}^{\varphi}$$

Since Γ' is of smaller depth, our inductive assumption gives us a cubical system $\mathcal{X}'$ for $\mathrm{L}^{\varphi}_{\Gamma'}$ extending $\mathcal{Y}$. The remainder of the proof will consist of a construction of a cubical system $\mathcal{X}$ for L^{φ}, which extends $\mathcal{X}'$ (and thus $\mathcal{Y}$).

Whenever we speak about maximal subgraphs of Γ, we will always mean them to be maximal proper subgraphs in L^{φ} (in particular not Γ itself).

Claim 1. Γ' is either a union of connected components of Γ, or it properly contains the union of all but one connected components of Γ.

Proof. Suppose that Γ' is not a union of connected components. Then there exists such a component, Γ_0 say, which intersect Γ' in a non-empty proper subgraph of itself. In particular this implies that Γ_0 is not a singleton. This in turn implies that $\Gamma_0 \in \mathrm{L}^{\varphi}$, and hence $\Gamma' \cup \Gamma_0 \in \mathrm{L}^{\varphi}$. The maximality of Γ' informs us that $\Gamma' \cup \Gamma_0 = \Gamma$, and so Γ' contains all but one component of Γ properly. □

Note that the above holds for any maximal subgraph of Γ.

We need to investigate two main cases, depending on whether Γ' is a union of components (part I) or Γ' properly contains the union of all but one components (part II). Note that this second case always occurs if Γ is connected.

9. Proof of the main theorem: part I

In this section we consider the case that Γ' is a union of connected components.

Let us suppose first that any two maximal subgraphs of Γ either coincide or are disjoint. In this case we have

$$\Gamma = \bigsqcup_{\Sigma \text{ maximal}} \Sigma \sqcup \Theta$$

with Θ discrete, since each vertex with non-trivial link lies in its extended star, which is preserved (and not all of Γ, since the latter is not a join), and thus is contained in a maximal subgraph. Hence we also have

$$\mathrm{L}^{\varphi} = \bigsqcup_{\Sigma \text{ maximal}} \mathrm{L}^{\varphi}_{\Sigma} \sqcup \{\Gamma\}$$

Observe that each maximal Σ is a union of components, as otherwise it would have to intersect Γ' non-trivially (by Claim 1). We conclude by applying Proposition 6.1 to cubical systems for $\mathrm{L}^{\varphi}_{\Sigma}$ (with Σ maximal) obtained by the inductive assumption (taking $\mathcal{X}'$ for $\mathrm{L}^{\varphi}_{\Gamma'}$).

We are left with the much more involved case, in which there exist maximal subgraphs of Γ which are neither equal nor disjoint.

Let $\Theta = \Gamma \smallsetminus \Gamma'$. Note that Θ is also a union of connected components. We define

$$\mathcal{S} = \{\Sigma \in \mathrm{L}^{\varphi} \mid \Theta \subseteq \Sigma\}$$

Claim 2. If $\Sigma \in \mathrm{L}^{\varphi}$ intersects both Θ and Γ' non-trivially then $\Sigma \in \mathcal{S}$. In particular, any maximal Σ different from Γ' contains Θ.

Proof. Take Σ as specified. Then $\Sigma \cup \Gamma' \in \mathrm{L}^{\varphi}$ since $\mathrm{st}(\Sigma \cap \Gamma') \subseteq \Gamma'$ as Γ' is a union of connected components. Thus $\Sigma \cup \Gamma' = \Gamma$ by maximality of Γ', and so $\Theta \subseteq \Sigma$. □

Claim 3. The system $\mathcal{S}$ is non-empty.

Proof. Suppose first that there exists a maximal $\Sigma \in \mathrm{L}^{\varphi}$ which is not a union of components. Then $\Sigma \cap \Gamma'$ is non-empty (by Claim 1), and $\Sigma \neq \Gamma'$. Hence $\Sigma \in \mathcal{S}$.

Now suppose that all maximal subgraphs Σ are unions of components. By our assumption there exist maximal Σ, Σ' which intersect non-trivially and do not coincide. But then, in particular, $\mathrm{st}(\Sigma \cap \Sigma') \subseteq \Sigma$, and so

$$\Sigma \cup \Sigma' \in \mathrm{L}^{\varphi}$$

Maximality now yields $\Sigma \cup \Sigma' = \Gamma$, and so, without loss of generality,

$$\Sigma \cap \Theta \neq \emptyset$$

If $\Sigma \cap \Gamma' = \emptyset$, then we must have $\Gamma' \subseteq \Sigma'$, and so $\Gamma' = \Sigma'$. But then $\Sigma \cap \Gamma' = \Sigma \cap \Sigma' \neq \emptyset$ by assumption, and so $\Sigma \in \mathcal{S}$. □

Since Θ is a non-empty union of components, the link of Θ is empty, and thus it follows by Lemma 4.2 that $\mathcal{S}$ is closed under taking unions and intersections. Hence so is

$$\mathcal{S}_{\Gamma'} = \{\Sigma \cap \Gamma' \mid \Sigma \in \mathcal{S}\}$$

Let us define

$$\mathcal{S}' = \{\mathrm{st}(\Sigma \cap \Gamma') \mid \Sigma \in \mathcal{S}\}$$

Note that $\mathcal{S}' \subseteq \mathcal{S}_{\Gamma'}$, since for every $\Sigma \in \mathcal{S}$ we have $\Sigma \cup \mathrm{st}(\Sigma \cap \Gamma') \in \mathcal{S}$ (see the proof of Claim 4 below).

Let

$$\Delta = \bigcap \mathcal{S}_{\Gamma'}$$

and

$$\Delta' = \bigcap \mathcal{S}'$$

Apply the System Intersection Axiom in $\mathcal{X}'$ to the collection $\mathcal{S}_{\Gamma'}$. We obtain a collection of standard copies $\widetilde{Y}_{\Sigma}$ of $\widetilde{X}'_{\Sigma}$ in $\widetilde{X}'_{\Gamma'}$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$, such that the copies $\widetilde{Y}_{\Sigma}$ intersect in $\widetilde{Y}_{\Delta}$ (which is a point when $\Delta = \emptyset$). Note that for each $\Sigma \in \mathcal{S}'$, the copy $\widetilde{Y}_{\Sigma}$ is unique (as $\mathrm{lk}(\Sigma) = \emptyset$), and hence fixed. Thus the intersection of all such copies is also fixed. By the Intersection Axiom, this intersection is $\widetilde{Y}_{\Delta'}$.

Let us note that

$$\Delta \cup \Theta = \bigcap \mathcal{S} \in \mathrm{L}^{\varphi}$$

Claim 4. $\Delta' \cup \Theta \in \mathrm{L}^{\varphi}$.

Proof. For each $\Sigma \in \mathcal{S}$ we have

$$\mathrm{st}(\Sigma \cap \Gamma') \cup \Theta = \mathrm{st}(\Sigma \cap \Gamma') \cup \Sigma \in \mathrm{L}^{\varphi}$$

since $\mathrm{lk}(\mathrm{st}(\Sigma\cap\Gamma')\cap\Sigma) \subseteq \mathrm{lk}(\Sigma\cap\Gamma') \subseteq \mathrm{st}(\Sigma\cap\Gamma')$, and due to part ii) of Lemma 4.2. Thus

$$\Delta' \cup \Theta = \bigcap_{\Sigma\in\mathcal{S}} (\mathrm{st}(\Sigma \cap \Gamma') \cup \Theta) \in \mathrm{L}^{\varphi}$$

by part i) of Lemma 4.2. □

Claim 5. There exists a cubical system $\mathcal{X}''$ for $\mathrm{L}^{\varphi}_{\Delta'\cup\Theta}$ which extends the subsystem $\mathcal{X}'_{\Delta'}$ of $\mathcal{X}'$ strongly.

Proof. Since Relative Nielsen Realisation holds for $\Delta'\cup\Theta$, as it is a proper subgraph of Γ and hence has lower depth, we obtain a complex $\mathcal{X}''$ for $\mathrm{L}^{\varphi}_{\Delta'\cup\Theta}$ extending $\mathcal{X}'_{\Delta'}$.

We now need to look at vertices of Δ' which are singletons in the join decomposition of Γ'; in other words we are looking at vertices in $\Delta' \cap Z(\Gamma')$. If there are no such, then Lemma 7.3 implies that $\mathcal{X}''$ extends $\mathcal{X}'_{\Delta'}$ strongly, since $Z(\Theta \cup \Delta') = \emptyset$.

Now suppose that $s \in \Delta' \cap Z(\Gamma')$ exists. Note that $\{s\} \in \mathrm{L}^{\varphi}$ since φ is link-preserving. Then we use Lemma 5.10 and replace $\mathcal{X}'$ by another cubical system which extends it (and so still extends the cubical system $\mathcal{Y}$ for $\mathrm{L}^{\varphi}_{\Xi}$), and such that it extends $\mathcal{X}''_{\{s\}}$ strongly. Repeat this procedure for each vertex in $\Delta' \cap Z(\Gamma')$. □

Now we are ready to start the construction.

9.1. Constructing the complexes.

Let us begin by building up the necessary cube complexes.

The basic idea is to glue the desired complexes from pieces lying in Γ' (where, by induction, we already have them in $\mathcal{X}'$) and pieces overlapping with Θ. The details will be different depending on how the invariant graph in question intersects Γ' and Θ, and thus the construction has several steps.

Step 1: constructing X_Γ. In this step we will construct various objects at once. First, we will build $\widetilde{X}_\Gamma$ from $\widetilde{\mathcal{X}}'_{\Gamma'}$ and $\widetilde{\mathcal{X}}''_{\Delta'\cup\Theta}$ by gluing $\widetilde{Y}_{\Delta'}$ and $\widetilde{R}_{\Delta'}$, a fixed standard copy of $\widetilde{\mathcal{X}}''_{\Delta'}$ in $\widetilde{X}''_{\Delta'\cup\Theta}$.

Since $\widetilde{X}_\Gamma$ is obtained from a gluing procedure, we obtain at the same time maps

$$\widetilde{\iota}_{\Gamma',\Gamma} \colon \widetilde{X}'_{\Gamma'} \to \widetilde{X}_\Gamma$$

and

$$\widetilde{\iota}_{\Delta'\cup\Theta,\Gamma} \colon \widetilde{X}''_{\Delta'\cup\Theta} \to \widetilde{X}_\Gamma$$

which are as required, since $\mathrm{lk}(\Gamma') = \mathrm{lk}(\Delta' \cup \Theta) = \emptyset$. We will also construct a fixed standard copy $\widetilde{R}_\Delta$ of $\widetilde{X}''_\Delta$ and $\widetilde{R}_{\Theta\cup\Delta}$ of $\widetilde{X}''_{\Theta\cup\Delta}$ in $\widetilde{X}''_{\Delta'\cup\Theta}$, with $\widetilde{R}_\Delta \subseteq \widetilde{R}_{\Theta\cup\Delta}$. Furthermore, we will construct a collection of standard copies $\widetilde{Z}_\Sigma$ of $\widetilde{X}'_\Sigma$ in $\widetilde{X}'_{\Gamma'}$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$, such that the copies $\widetilde{Z}_\Sigma$ intersect in $\widetilde{Z}_\Delta$, and $\widetilde{Z}_\Delta$ is identified with $\widetilde{R}_\Delta$ under our gluing.

We need to consider two cases.

Case 1: Δ' is empty.

In this case we have $\Theta = \Theta \cup \Delta' \in \mathrm{L}^{\varphi}$.

We apply Proposition 6.1 to $\mathcal{X}'$ and $\mathcal{X}''$ and obtain a cubical system for $\mathrm{L}^{\varphi}_{\Gamma'} \cup \mathrm{L}^{\varphi}_{\Theta} \cup \{\Gamma\}$. We are only interested in the cube complex associated to Γ in this system – let us call it C, and its universal cover $\widetilde{C}$. The complex $\widetilde{C}$ is obtained from $\widetilde{X}'_{\Gamma'}$ and $\widetilde{X}''_{\Theta}$ by identifying a single point in each; let $\widetilde{q}$ denote the point in $\widetilde{X}''_{\Theta}$.

In particular q, the projection of $\widetilde{q}$ in C, is fixed. Hence, picking $\widetilde{q}$ as a basepoint of $\widetilde{C}$ (and a trivial path), we get a geometric representative $h_q \in \mathrm{Aut}(A_\Gamma)$ of $\varphi(h)$ for each $h \in H$. This representative satisfies

$$h_q(A_{\Gamma'}) = A_{\Gamma'} \text{ and } h_q(A_\Theta) = A_\Theta$$

by construction.

Since $\Delta' = \emptyset$, the fixed standard copy $\widetilde{Y}_{\Delta'}$ is simply a point $\widetilde{p} \in \widetilde{X}'_{\Gamma'}$ with a fixed projection $p \in X'_{\Gamma'}$. The point $\widetilde{p}$ can also be viewed as a point in $\widetilde{C}$. Similarly p becomes a point in C which is also H-fixed. Thus we can use $\widetilde{p}$ as a basepoint, and for each $\varphi(h)$ obtain a geometric representative $h_p \in \mathrm{Aut}(A_\Gamma)$ such that $h_p(A_{\Gamma'}) = A_{\Gamma'}$.

We have $h_q h_p^{-1} = c(x)$ since they are representatives of the same outer automorphism, with $x \in A_\Gamma$.

Take $\Sigma \in \mathcal{S}$. Now

$$A_\Theta = h_q(A_\Theta) \leqslant h_q(A_\Sigma) = {A_\Sigma}^y$$

for some $y \in A_\Gamma$, since $\Theta \in \mathrm{L}^\varphi$. Proposition 2.5 implies that

$$y \in N(A_\Sigma)N(A_\Theta) = A_\Sigma$$

and so $h_q(A_\Sigma) = A_\Sigma$.

Since $h_q(A_\Sigma) = A_\Sigma$ and $h_q(A_{\Gamma'}) = A_{\Gamma'}$, we have $h_q(A_{\Sigma\cap\Gamma'}) = A_{\Sigma\cap\Gamma'}$ and so $h_q(A_{\mathrm{st}(\Sigma\cap\Gamma')}) = A_{\mathrm{st}(\Sigma\cap\Gamma')}$ (since $A_{\mathrm{st}(\Sigma\cap\Gamma')} = N(A_{\mathrm{st}(\Sigma\cap\Gamma')})$). But we also have $h_p(A_{\mathrm{st}(\Sigma\cap\Gamma')}) = A_{\mathrm{st}(\Sigma\cap\Gamma')}$, since the basepoint $\widetilde{p}$ lies in $\widetilde{Y}_{\mathrm{st}(\Sigma\cap\Gamma')}$ for each $\Sigma \in \mathcal{S}$. Thus

$$A_{\mathrm{st}(\Sigma\cap\Gamma')} = h_q h_p^{-1}(A_{\mathrm{st}(\Sigma\cap\Gamma')}) = {A_{\mathrm{st}(\Sigma\cap\Gamma')}}^x$$

and thus $x \in N(A_{\mathrm{st}(\Sigma\cap\Gamma')}) = A_{\mathrm{st}(\Sigma\cap\Gamma')}$ for each $\Sigma \in \mathcal{S}$. Hence

$$x \in A_{\Delta'} = \{1\}$$

and so $h_q = h_p$. This equality means that the inherited action of H on X_Γ, the cube complex obtained from $\widetilde{X}'_{\Gamma'}$ and $\widetilde{X}''_\Theta$ by gluing $\widetilde{p}$ and $\widetilde{q}$, induces φ.

We now put $\widetilde{R}_\Delta = \widetilde{R}_{\Delta'} = \widetilde{q}$, and set $\widetilde{Z}_\Sigma = \widetilde{Y}_\Sigma$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$. We also put $\widetilde{R}_{\Theta\cup\Delta} = \widetilde{X}''_\Theta$.

Case 2: Δ' is non-empty.

Let us observe the crucial property of Δ'.

Claim 6. Let $\Sigma \in \mathcal{S}$, and let $\widetilde{p} \in \widetilde{Y}_{\Delta'}$. Then there exists a standard copy $\widetilde{W}_{\Sigma\cap\Gamma'}$ of $\widetilde{X}'_{\Sigma\cap\Gamma'}$ in $\widetilde{X}'_{\Gamma'}$ containing $\widetilde{p}$

Proof. When $\Sigma = \Theta$ the result follows trivially.

Suppose that $\Theta \subset \Sigma$ is a proper subgraph. Note that we have a standard copy $\widetilde{Y}_{\mathrm{st}(\Sigma\cap\Gamma')}$ which contains $\widetilde{Y}_{\Delta'}$, and hence $\widetilde{p}$. We have

$$\mathrm{st}(\Sigma\cap\Gamma') = \Sigma\cap\Gamma' * \mathrm{lk}_{\Gamma'}(\Sigma\cap\Gamma')$$

and so the Product Axiom and Composition Axiom imply that there exists a standard copy $\widetilde{W}_{\Sigma\cap\Gamma'}$ as required. □

We have $\mathrm{lk}(\Delta')_{\Delta'\cup\Theta} = \emptyset$, as Θ is a union of components, and so there is a unique (and hence fixed) standard copy of $X''_{\Delta'} = X'_{\Delta'}$ in $X''_{\Delta'\cup\Theta}$; we will denote it by $\widetilde{R}_{\Delta'}$. For the same reason there is a unique (and so fixed) standard copy $\widetilde{R}_{\Delta\cup\Theta}$ of $\widetilde{X}''_{\Delta\cup\Theta}$ in $\widetilde{X}''_{\Theta\cup\Delta'}$. The Matching Property tells us that these intersect in a standard copy $\widetilde{R}_\Delta$ of $\widetilde{X}''_\Delta$; this copy is fixed, since it is the intersection of two fixed standard copies.

We now obtain a complex $\widetilde{C}$ from $\widetilde{X}'_{\Gamma'}$ and $\widetilde{X}''_{\Delta'\cup\Theta}$ by gluing $\widetilde{Y}_{\Delta'} = \widetilde{R}_{\Delta'}$. However, the projection C of our complex $\widetilde{C}$ might not yet realise the desired action of H.

Pick $h \in H$. Take a point $\widetilde{r} \in \widetilde{R}_\Delta$ as a basepoint, let r be its projection as usual. Take a path $\gamma(h)$ in R_Δ from r to $h.r$. Note that $\widetilde{r}$ is a point in $\widetilde{R}_{\Delta'}$, and hence in $\widetilde{C}$ (after the gluing). Let $\widetilde{p} \in \widetilde{Y}_{\Delta'} = \widetilde{R}_{\Delta'}$ be the corresponding point; we view it as a point in $\widetilde{C}$ as well, and denote its projection by p as usual. Let $h_r, h_p \in \mathrm{Aut}(A_\Gamma)$ denote the corresponding geometric representatives.

Claim 7. The gluing $\widetilde{C}$ is faulty within $Z(A_{\Delta'})$.

Proof. Let us first remark that $\mathrm{lk}(\Sigma) \subseteq \mathrm{lk}(\Theta) = \emptyset$ for all $\Sigma \in \mathcal{S}$.

By construction we see that

$$h_r(A_{\Delta\cup\Theta}) = A_{\Delta\cup\Theta}$$

Take $\Sigma \in \mathcal{S}$. Since Σ lies in L^φ, we have

$$h_r(A_\Sigma) = A_\Sigma^y$$

for some element $y \in A_\Gamma$. But then

$$A_{\Delta\cup\Theta} = h_r(A_{\Delta\cup\Theta}) \leqslant h_r(A_\Sigma) = A_\Sigma^y$$

Proposition 2.5 implies that $y \in N(A_\Sigma)N(A_{\Delta\cup\Theta}) = A_\Sigma A_{\Delta\cup\Theta} = A_\Sigma$, and thus

$$h_r(A_\Sigma) = A_\Sigma$$

Now $\Sigma \cap \Gamma' \in \mathrm{L}^\varphi$ and so

$$h_r(A_{\Sigma\cap\Gamma'}) = A_{\Sigma\cap\Gamma'}^z$$

We have $A_{\Sigma\cap\Gamma'}^z \subseteq A_\Sigma$, and so (using Proposition 2.5 again) we get, without loss of generality, $z \in N(A_\Sigma) = A_\Sigma$ since $\mathrm{lk}(\Sigma) = \emptyset$ (as $\Sigma \in \mathcal{S}$).

Let us now focus on h_p. By Claim 6, there exists a standard copy of $\widetilde{X}'_{\Sigma\cap\Gamma'}$ containing $\widetilde{p}$; it is clear that it will also contain the copy of the path $\gamma(h)$. Hence

$$h_p(A_{\Sigma\cap\Gamma'}) = A_{\Sigma\cap\Gamma'}$$

and so

$$A_{\Sigma\cap\Gamma'}^z = h_r(A_{\Sigma\cap\Gamma'}) = c(x(h)^{-1})h_p(A_{\Sigma\cap\Gamma'}) = A_{\Sigma\cap\Gamma'}^{x(h)^{-1}}$$

where

$$c(x(h)) = h_p h_r^{-1}$$

and $x(h) \in C(A_{\Delta'})$.

Now

$$z = z\,x(h)\,x(h)^{-1} \in N(A_{\Sigma\cap\Gamma'})C(A_{\Delta'}) \leqslant A_{\Gamma'}$$

and thus

$$z \in A_\Sigma \cap A_{\Gamma'} = A_{\Sigma\cap\Gamma'}$$

and so

$$h_r(A_{\Sigma\cap\Gamma'}) = {A_{\Sigma\cap\Gamma'}}^z = A_{\Sigma\cap\Gamma'}$$

Finally

$$c(x(h)) = h_p h_r^{-1}(A_{\Sigma\cap\Gamma'}) = A_{\Sigma\cap\Gamma'}$$

for all $\Sigma \in \mathcal{S}$, which in turn implies

$$x(h) \in \bigcap_{\Sigma\in\mathcal{S}} N(A_{\Sigma\cap\Gamma'}) = \bigcap_{\Sigma\in\mathcal{S}} A_{\mathrm{st}(\Sigma\cap\Gamma')} = A_{\Delta'}$$

Therefore $x(h) \in C(A_{\Delta'}) \cap A_{\Delta'} = Z(A_{\Delta'})$.

This statement holds for each h, and thus the fault of our gluing satisfies the claim. ∎

Now we are in a position to apply Proposition 7.5 and obtain a new glued up complex, which we call $\widetilde{X}_\Gamma$, which realises our action φ.

Recall that we have a standard copy $\widetilde{R}_\Delta$ in $\widetilde{X}''_{\Delta' \cup \Theta}$. The gluing sends $\widetilde{R}_\Delta$ to some standard copy of $\widetilde{X}'_\Delta$ in $\widetilde{X}'_{\Gamma'}$, which lies within $\widetilde{Y}_{\Delta'}$ (we are using the Composition Property here); let us denote this standard copy by $\widetilde{Z}_\Delta$. It is fixed since $\widetilde{R}_\Delta$ is.

By Claim 6, we may pick a standard copy $\widetilde{Z}_\Sigma$ of $\widetilde{X}'_\Sigma$ in $\widetilde{X}'_{\Gamma'}$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$ such that they will all intersect in $\widetilde{Z}_\Delta$. Let us choose such a family of standard copies.

Step 2: Constructing the remaining complexes. Let $\Sigma \in \mathrm{L}^\varphi$ be a proper subgraph of Γ. When $\Sigma \subseteq \Gamma'$ we set $\widetilde{X}_\Sigma = \mathcal{X}'_\Sigma$; when $\Sigma \subseteq \Theta \cup \Delta$, we set $\widetilde{X}_\Sigma = \widetilde{X}''_\Sigma$. Note that this is consistent for graphs $\Sigma \subseteq \Delta$ since $\mathcal{X}''$ extends $\mathcal{X}'_{\Delta'}$ (and thus $\mathcal{X}'_\Delta$) strongly.

Now let $\Sigma \in \mathcal{S}$. We have a standard copy $\widetilde{Z}_{\Sigma\cap\Gamma'}$ in $\widetilde{X}_{\Gamma'} = \widetilde{X}'_{\Gamma'}$. This last complex is embedded in $\widetilde{X}_\Gamma$, and so we may think of $\widetilde{Z}_{\Sigma\cap\Gamma'}$ as being embedded therein as well.

By construction $\widetilde{Z}_{\Sigma\cap\Gamma'}$ contains $\widetilde{Z}_\Delta$; this last copy is glued to $\widetilde{R}_\Delta$, which in turn lies within $\widetilde{R}_{\Theta\cup\Delta}$. We define X_Σ to be the union (taken in X_Γ) of $Z_{\Sigma\cap\Gamma'}$ and $R_{\Delta\cup\Theta}$.

Observe that when $\Sigma = \Theta \cup \Delta$ we have given two ways of constructing $\widetilde{X}_\Sigma$; it is however immediate that the outcome of both methods is identical.

The projection X_Σ carries the desired marking by construction. The action of H is also the desired one; taking any $h \in H$, and looking at a geometric representative (in $\mathrm{Aut}(A_\Gamma)$) obtained by choosing a basepoint and a path in the subcomplex X_Σ, we get an automorphism of A_Γ which preserves A_Σ. This automorphism is a representative of $\varphi(h)$, and so the restriction to A_Σ is the desired one. But this is equal to the geometric representative of the action of h on X_Σ obtained using the same basepoint and path.

We now embed $\widetilde{X}_\Sigma$ into $\widetilde{X}_\Gamma$ in such a way that the image contains both $\widetilde{Z}_{\Sigma\cap\Gamma'}$ and $\widetilde{R}_{\Theta\cup\Delta}$, and so that this embedding gives the inclusion $X_\Sigma \subseteq X_\Gamma$ when we take the quotients.

Note that the construction gives us an embeddings $\widetilde{\iota}_{\Sigma,\Gamma} \colon \widetilde{X}_\Sigma \to \widetilde{X}_\Gamma$, $\widetilde{\iota}_{\Sigma\cap\Gamma',\Sigma} \colon \widetilde{X}_{\Sigma\cap\Gamma'} \to \widetilde{X}_\Sigma$ and $\widetilde{\iota}_{\Delta\cup\Theta,\Sigma} \colon \widetilde{X}_{\Delta\cup\Theta} \to \widetilde{X}_\Sigma$, which are as required since $\mathrm{lk}(\Sigma) = \mathrm{lk}_\Sigma(\Sigma \cap \Gamma') = \mathrm{lk}_\Sigma(\Delta \cup \Theta) = \emptyset$.

9.2. Constructing the maps. We now need to specify the maps $\widetilde{\iota}$. We take $\Sigma, \Sigma' \in \mathrm{L}^\varphi$ with $\Sigma \subseteq \Sigma'$. When the two graphs are identical, we set $\widetilde{\iota}_{\Sigma,\Sigma'}$ to be the identity. Let us now suppose that $\Sigma \neq \Sigma'$.

(1) $\Sigma' \subseteq \Gamma'$ **or** $\Sigma' \subseteq \Delta \cup \Theta$

In this case the cube complexes X_Σ and $X_{\Sigma'}$ are obtained directly from another cubical system ($\mathcal{X}'$ or $\mathcal{X}''_\Delta$), and we take $\widetilde{\iota}_{\Sigma,\Sigma'}$ to be the map coming from that system.

For the remaining cases, we will assume that the hypothesis of (1) is not satisfied, which implies $\Sigma' \in \mathcal{S}$.

(2) $\Sigma \subseteq \Gamma'$

In this case we have $\mathrm{st}(\Sigma) = \mathrm{st}(\Sigma) \cap \Gamma'$, since Γ' is a union of components. Therefore $\mathrm{lk}_{\Sigma'}(\Sigma) = \mathrm{lk}_{\Sigma'}(\Sigma) \cap \Gamma'$ as well. We define

$$\widetilde{\iota}_{\Sigma,\Sigma'} = \widetilde{\iota}_{\Sigma'\cap\Gamma',\Sigma'} \circ \widetilde{\iota}_{\Sigma,\Sigma'\cap\Gamma'}$$

where the last map was defined in (1) above, and the map $\widetilde{\iota}_{\Sigma'\cap\Gamma',\Sigma'}$ was constructed together with the complex $\widetilde{X}_{\Sigma'}$ in Step 2 of Subsection 9.1.

(3) $\Sigma \subseteq \Theta \cup \Delta$

In this case we have $\mathrm{st}(\Sigma) = \mathrm{st}(\Sigma) \cap (\Theta \cup \Delta)$, since Θ is a union of components. Therefore $\mathrm{lk}_{\Sigma'}(\Sigma) = \mathrm{lk}_{\Sigma'}(\Sigma) \cap \Theta$ as before. We define

$$\widetilde{\iota}_{\Sigma,\Sigma'} = \widetilde{\iota}_{\Delta\cup\Theta,\Sigma'} \circ \widetilde{\iota}_{\Sigma,\Delta\cup\Theta}$$

where the last map was defined in (1) above, and the map $\widetilde{\iota}_{\Sigma'\cap\Gamma',\Sigma'}$ was constructed together with the complex $\widetilde{X}_{\Sigma'}$ in Step 2 of Subsection 9.1.

(4) $\Sigma \in \mathcal{S}$

Observe that both $\widetilde{X}_\Sigma$ and $\widetilde{X}_{\Sigma'}$ are defined as subcomplexes of $\widetilde{X}_\Gamma$. Since $\widetilde{Z}_\Sigma \subset \widetilde{Z}_{\Sigma'}$ (by the Intersection Axiom, as they both contain $\widetilde{Z}_\Delta$), we see that $\widetilde{X}_\Sigma \subset \widetilde{X}_{\Sigma'}$; we define this embedding to be $\widetilde{\iota}_{\Sigma,\Sigma'}$. Note that the map is as required since $\mathrm{lk}_{\Sigma'}(\Sigma) = \emptyset$ as $\Theta \subseteq \Sigma$, and Θ is a union of components.

Note that we have (again) given two constructions when

$$\Sigma = \Delta \cup \Theta$$

but again they are easily seen to be identical.

9.3. Verifying the axioms. The first two axioms depend only on two subgraphs $\Sigma, \Sigma' \in \mathrm{L}^\varphi$ with $\Sigma \subseteq \Sigma'$. This is the same assumption as in 9.2, and hence the verification of the two axioms will follow the same structure as the construction of the maps – we will consider four cases, and the assumption in each will be identical to the assumptions of the corresponding case above.

Product Axiom. Suppose that $\Sigma' = \mathrm{st}_{\Sigma'}(\Sigma)$.

(1) In this case the Product Axiom follows from the Product Axiom in $\mathcal{X}'$ or $\mathcal{X}''$.
(2) As in the analogous case of Subsection 9.2, we see that $\mathrm{st}_{\Sigma'}(\Sigma) \subseteq \Gamma'$. But here $\Sigma' = \mathrm{st}_{\Sigma'}(\Sigma)$ and so we are in case (1) above.
(3) As in the analogous case of Subsection 9.2, we see that

$$\mathrm{st}_{\Sigma'}(\Sigma) \subseteq \Delta \cup \Theta$$

But here $\Sigma' = \mathrm{st}_{\Sigma'}(\Sigma)$ and so we are in case (1) above.
(4) In this case we have $\mathrm{lk}(\Sigma) = \emptyset$, and so $\Sigma = \Sigma'$. But in such a case we defined $\widetilde{\iota}_{\Sigma,\Sigma'}$ to be the identity.

Orthogonal Axiom. Let $\Lambda = \mathrm{lk}_{\Sigma'}(\Sigma)$.

(1) Follows from the axiom in $\mathcal{X}'$ or $\mathcal{X}''$.
(2) As before we have $\Lambda \subseteq \Gamma'$. The construction of the maps $\widetilde{\iota}_{\Sigma,\Sigma'}$ and $\widetilde{\iota}_{\Lambda,\Sigma'}$ immediately tells us that it is enough to verify the axiom for the pair Σ, Γ'. But this is covered by (1).
(3) As before we have $\Lambda \subseteq \Delta \cup \Theta$. The construction of the maps $\widetilde{\iota}_{\Sigma,\Sigma'}$ and $\widetilde{\iota}_{\Lambda,\Sigma'}$ immediately tells us that it is enough to verify the axiom for the pair $\Sigma, \Delta \cup \Theta$. But this is covered by (1).
(4) In this case we have $\Lambda = \emptyset$, and hence the axiom follows trivially, since $\widetilde{\iota}_{\Lambda,\Sigma'}$ is onto (by the Product Axiom, which we have already shown).

Intersection Axiom. Let us now verify that $\mathcal{X}$ satisfies the Intersection Axiom. Take $\Sigma, \Sigma', \Omega \in \mathrm{L}^\varphi$ such that $\Sigma \subseteq \Omega$ and $\Sigma' \subseteq \Omega$, and let $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ be standard copies of, respectively, $\widetilde{X}_\Sigma$ and $\widetilde{X}_{\Sigma'}$ in $\widetilde{X}_\Omega$ with non-empty intersection. We need to show that the intersection is the image of a standard copy of $\Sigma \cap \Sigma'$ in each.

As in the first two cases, the details depend on the inclusions $\Sigma, \Sigma' \subseteq \Omega$. The cases will thus be labeled by pairs of integers (n, m), the first determining in which case the map $\widetilde{\iota}_{\Sigma,\Omega}$ was constructed, and the second playing the same role for $\widetilde{\iota}_{\Sigma',\Omega}$. By symmetry we only need to consider $n \leqslant m$.

(1,1) In this case the axiom follows from the Intersection Axiom in $\mathcal{X}'$ or $\mathcal{X}''$.

In what follows we can assume that $\Omega \in \mathcal{S}$, and hence $\widetilde{X}_\Omega$ is obtained by gluing $\widetilde{Z}_{\Omega\cap\Gamma'}$ and $\widetilde{R}_{\Delta\cup\Theta}$ along $\widetilde{Z}_\Delta$.

(2,2) Here the problem is reduced to checking the axiom for the triple $\Sigma, \Sigma', \Omega \cap \Gamma'$, for which it holds by (1,1).

(2,3) By construction, the given standard copies $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ must lie within the standard copies $\widetilde{Z}_{\Omega\cap\Gamma'}$ and $\widetilde{R}_{\Delta\cup\Theta}$ respectively, and hence intersect within $\widetilde{Z}_\Delta = \widetilde{R}_\Delta$.

We use the Intersection Axiom of $\mathcal{X}'$ for $\widetilde{Z}_\Delta$ and $\widetilde{Y}_\Sigma$ inside $\widetilde{Z}_{\Omega\cap\Gamma'}$ and see that the two copies intersect in a copy of $\widetilde{X}_{\Delta\cap\Sigma}$, which is also the image of a standard copy of $\widetilde{X}_{\Delta\cap\Sigma}$ in $\widetilde{Z}_{\Delta\cap\Omega}$.

We repeat the argument for Σ' and obtain a standard copy of $\widetilde{X}_{\Delta\cap\Sigma'}$ in $\widetilde{Z}_\Delta$. Now this copy intersects the one of $\widetilde{X}_{\Delta\cap\Sigma}$, and hence, applying the Intersection Axiom again, they intersect in a copy of $\widetilde{X}_{\Delta\cap\Sigma\cap\Sigma'}$ in $\widetilde{Z}_\Delta$. But $\Delta \cap \Sigma \cap \Sigma' = \Sigma \cap \Sigma'$, and so we have found the desired standard copy in $\widetilde{Z}_\Delta$. Now the Composition Property (Lemma 5.7) implies that this is also a standard copy in $\widetilde{X}_\Sigma$, $\widetilde{X}_{\Sigma'}$ and $\widetilde{X}_{\Omega\cap\Gamma'}$, and thus this is also a standard copy in $\widetilde{X}_\Omega$ by construction.

(2,4) The non-trivial intersection of any standard copy of $\widetilde{X}_\Sigma$ and any standard copy of $\widetilde{X}_{\Sigma'}$ in $\widetilde{X}_\Omega$ is in fact contained in the standard copy of $\widetilde{X}_{\Omega\cap\Gamma'}$, since any copy of $\widetilde{X}_\Sigma$ is contained therein. Therefore the intersection is also contained in a standard copy of $\widetilde{X}_{\Sigma'\cap\Gamma'}$ by construction of $\widetilde{\iota}_{\Sigma',\Omega}$. We apply the Intersection Axiom in $\widetilde{X}_{\Omega\cap\Gamma'}$, and observe that the standard copy of $\widetilde{X}_{\Sigma\cap\Sigma'}$ in $\widetilde{X}_{\Omega\cap\Gamma'}$ obtained this way is also a standard copy in $\widetilde{X}_\Omega$ by construction.

(3,3) Here the problem is reduced to checking the axiom for the triple $\Sigma, \Sigma', \Delta \cup \Theta$, for which it holds by (1,1).

(3,4) The non-trivial intersection of any standard copy of $\widetilde{X}_{\Sigma'}$ and any standard copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_\Omega$ is in fact contained in the standard copy of $\widetilde{X}_{\Delta\cup\Theta}$, since any copy of $\widetilde{X}_{\Sigma'}$ is contained therein. Therefore the intersection is also contained in a standard copy of $\widetilde{X}_{\Delta\cup\Theta}$ by construction of $\widetilde{\iota}_{\Sigma',\Omega}$. We apply the Intersection Axiom in $\widetilde{X}_{\Delta\cup\Theta}$, and observe that the standard copy of $\widetilde{X}_{\Sigma\cap\Sigma'}$ in $\widetilde{X}_{\Delta\cup\Theta}$ obtained this way is also a standard copy in $\widetilde{X}_\Omega$ by construction.

(4,4) In this case $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ are obtained from $\widetilde{Z}_{\Sigma\cap\Gamma'}$, $\widetilde{R}_{\Delta\cup\Theta}$, and $\widetilde{Z}_{\Sigma'\cap\Gamma'}$, $\widetilde{R}_{\Delta\cup\Theta}$ respectively. Hence $\widetilde{Z}_{\Sigma\cap\Gamma'} \cap \widetilde{Z}_{\Sigma'\cap\Gamma'}$ contains $\widetilde{Z}_\Delta$, and so $\widetilde{Z}_{\Sigma\cap\Gamma'}$ and $\widetilde{Z}_{\Sigma'\cap\Gamma'}$ intersect non-trivially. Applying the Intersection Axiom in $\widetilde{X}_{\Omega\cap\Gamma'}$ tells us that in fact $\widetilde{Z}_{\Sigma\cap\Gamma'} \cap \widetilde{Z}_{\Sigma'\cap\Gamma'} = \widetilde{Z}_{\Sigma\cap\Sigma'\cap\Gamma'}$, which contains $\widetilde{Z}_\Delta$. Thus $\widetilde{Y}_\Sigma \cap \widetilde{Y}_{\Sigma'}$ is equal to the unique standard copy of $\widetilde{X}_{\Sigma\cap\Sigma'}$ in $\widetilde{X}_\Gamma$, which lies within the unique standard copy of $\widetilde{X}_\Omega$ by construction. It is also a standard copy in $\widetilde{X}_\Omega$, again by construction.

System Intersection Axiom. Take a subsystem $\mathbb{P} \subseteq \mathrm{L}^\varphi$ closed under taking unions. If all elements of $\mathbb{P}$ lie in Γ' or in $\Delta \cup \Theta$, then we are done (from the System Intersection Axiom of $\mathcal{X}'$ or $\mathcal{X}''$). So let us suppose this is not the case, that is, suppose that there exists $\Sigma \in \mathbb{P} \cap \mathcal{S}$. Hence $\bigcup \mathbb{P} \in \mathcal{S}$.

Define

$$\mathbb{P}' = \{\Sigma \in \mathbb{P} \mid \Theta \subseteq \Sigma\}$$

Note that each $\Sigma \in \mathbb{P}'$ has $\mathrm{lk}(\Sigma) = \emptyset$, and thus there exists a unique standard copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_{\bigcup \mathbb{P}}$ for each $\Sigma \in \mathbb{P}'$, and all these copies intersect non-trivially (they all contain $\widetilde{R}_{\Delta\cup\Theta}$).

Note that $\mathbb{P}'$ is closed under taking unions, and that $\Theta \subseteq \bigcap \mathbb{P}'$. Thus we may apply Lemma 5.5. The construction is now finished by the application of the Intersection Axiom (which we have already proved).

10. Proof of the main theorem: part II

In this section we deal with the case that Γ' properly contains the union of all but one connected components of Γ. For notational convenience let Γ_0 denote the connected component not contained in Γ'.

Claim 8. $\mathrm{lk}(\Gamma') = \emptyset$

Proof. Note that $\mathrm{st}(\Gamma') \in \mathrm{L}^\varphi$ since $\Gamma' \in \mathrm{L}^\varphi$. If $\mathrm{lk}(\Gamma') \neq \emptyset$ we have $\Gamma' \subset \mathrm{st}(\Gamma')$ and so $\mathrm{st}(\Gamma') = \Gamma$. But Γ is not a join, and $\mathrm{st}(\Gamma')$ is. This is a contradiction. □

We let Θ be the complement

$$\Theta = \Gamma \smallsetminus \Gamma' \subseteq \Gamma_0$$

Note that at this point we do not claim that Θ is contained in L^φ.

Lemma 4.11 (applied to Γ_0 and $\Gamma' \cap \Gamma_0$) implies that for all $w \in \partial\Gamma'$ we have $\Theta \subseteq \mathrm{lk}(w)$. Thus we have

$$\partial\Gamma' = \mathrm{lk}(\Theta)$$

10.1. Constructing the complexes. We are now ready to define the cube complexes forming $\mathcal{X}$.

To describe the cases, we need the following additional graphs. We let

$$\overline{\Theta} = \Theta * \partial\Gamma'$$

and we consider the subsystem of all invariant graphs which contain $\overline{\Theta}$:

$$\mathcal{S} = \{\Sigma \in \mathrm{L}^\varphi \mid \overline{\Theta} \subseteq \Sigma\}$$

Since $\mathrm{lk}(\overline{\Theta}) = \emptyset$, the link of every element in $\mathcal{S}$ is empty as well. Thus, $\mathcal{S}$ is closed under taking unions by Lemma 4.2, Hence

$$\mathcal{S}_{\Gamma'} = \{\Sigma \cap \Gamma' \mid \Sigma \in \mathcal{S}\}$$

is also closed under taking unions.

We abbreviate

$$\Delta = \bigcap \mathcal{S}_{\Gamma'}$$

For any $\Sigma \in \mathrm{L}^\varphi$ we let

$$\Sigma_1 = \Sigma \cap \Gamma' \quad \text{and} \quad \Sigma_2 = \Sigma \cap (\Delta \cup \Theta)$$

In other words, these are the part of Σ which lie inside Γ' and inside $\Theta \cup \Delta$.

Step 1: Constructing X_Γ. In this step we will actually do a little more. First we will construct a suitable cubical system $\mathcal{X}''$ for $\mathrm{L}^\varphi_{\mathrm{st}(\partial\Gamma')}$ extending $\mathcal{X}'_\Delta$ strongly. Then we will find a fixed standard copy $\widetilde{R}_\Delta$ of $\widetilde{X}''_\Delta$ in $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$. We will also construct a family of standard copies $\widetilde{Z}_\Sigma$ of $\widetilde{X}'_\Sigma$ in $\widetilde{X}'_{\Gamma'}$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$, such that they all intersect in the copy $\widetilde{Z}_\Delta$. We will show that this last copy is fixed, and obtain $\widetilde{X}_\Gamma$ from $\widetilde{X}'_{\Gamma'}$ and $\widetilde{X}''_{\Delta\cup\Theta}$ (which will be shown to contain $\widetilde{R}_\Delta$) by gluing $\widetilde{R}_\Delta$ and $\widetilde{Z}_\Delta$.

Note that the construction of a gluing then also gives us maps

$$\widetilde{\iota}_{\Gamma',\Gamma} \colon \widetilde{X}'_{\Gamma'} \to \widetilde{X}_\Gamma$$

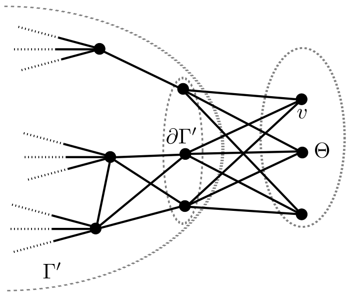

FIGURE 10.1. The relevant subgraphs of Γ in case 1

and

$$\widetilde{\iota}_{\Delta\cup\Theta,\Gamma} \colon \widetilde{X}''_{\Delta\cup\Theta} \to \widetilde{X}_\Gamma$$

which will be as desired since $\mathrm{lk}(\Gamma') = \emptyset = \mathrm{lk}(\Delta \cup \Theta)$.

We need to consider two cases depending on whether $\mathrm{lk}(\partial\Gamma')$ intersects Γ' or not.

Case 1: $\mathrm{lk}(\partial\Gamma') \cap \Gamma' = \emptyset$. Recall that $\partial\Gamma' = \mathrm{lk}(\Theta)$. By the case assumption we therefore have

$$\Theta = \mathrm{lk}(\partial\Gamma') \in \mathrm{L}^\varphi$$

Since Θ has lower depth than Γ, our inductive assumption gives us a cubical system for $\mathrm{L}^\varphi_\Theta$. By Proposition 5.9 there is a system $\mathcal{X}''$ for $\mathrm{L}^\varphi_{\overline{\Theta}}$, strongly extending both this cubical system and the subsystem $\mathcal{X}'_{\partial\Gamma'}$ of $\mathcal{X}'$.

Take any $v \in \Theta$. We have $\partial\Gamma' = \widehat{\mathrm{st}}(v)_1$ and thus $\partial\Gamma' \in \mathrm{L}^\varphi$. Hence in particular $\Delta = \partial\Gamma'$. We also have $\mathrm{st}(\partial\Gamma') = \partial\Gamma' * \Theta \in \mathrm{L}^\varphi$.

Claim 9. The induced action $H \to \mathrm{Out}(A_\Theta)$ lifts to an action

$$H \to \mathrm{Aut}(A_\Theta)$$

Proof. Since $\partial\Gamma' \in \mathrm{L}^\varphi$, for each $h \in H$ there exists a representative

$$h_1 \in \mathrm{Aut}(A_\Gamma)$$

of $\varphi(h)$ such that $h_1(A_{\partial\Gamma'}) = A_{\partial\Gamma'}$. Now

$$h_1(A_{\overline{\Theta}}) = A_{\overline{\Theta}}{}^x$$

since $\overline{\Theta} \in \mathrm{L}^\varphi$, with $x \in A_\Gamma$. Proposition 2.5 tells us that

$$x \in N(A_{\overline{\Theta}})N(A_{\partial\Gamma'}) = A_{\overline{\Theta}}$$

and so

$$h_1(A_{\overline{\Theta}}) = A_{\overline{\Theta}}$$

for each h. Since $\Theta \in \mathrm{L}^\varphi$ we also have

$$h_1(A_\Theta) = A_\Theta{}^y$$

with $y \in N(A_\Theta)N(A_{\overline{\Theta}}) = A_{\overline{\Theta}}$, again by Proposition 2.5. Hence

$$h_1(A_\Theta) = A_\Theta$$

Finally, we have

$$h_1(A_{\Gamma'}) = A_{\Gamma'}{}^z$$

with $z \in N(A_{\Gamma'})N(A_\Theta) = A_{\partial\Gamma'}A_\Theta$. Without loss of generality we take $z \in A_\Theta$, and so we construct new representatives $h_2 \in \mathrm{Aut}(A_\Gamma)$ for $\varphi(h)$ (by multiplying h_1 with the conjugation by z^{-1} in the appropriate way) which satisfy

$$h_2(A_{\partial\Gamma'}) = A_{\partial\Gamma'}, \ h_2(A_\Theta) = A_\Theta, \text{ and } h_2(A_{\Gamma'}) = A_{\Gamma'}$$

Now, given $h, g \in H$, we know that $h_2 g_2 {(hg)_2}^{-1}$ is equal to a conjugation which preserves $A_{\Gamma'}$ and A_Θ. The former fact implies that the conjugating element lies in $A_{\Gamma'}$, and the latter that it lies in $A_\Theta \times A_{\partial\Gamma'}$. Hence the conjugating element lies in $A_{\partial\Gamma'}$. But such a conjugation is trivial when we restrict the action to $\mathrm{Out}(A_\Theta)$, and hence our chosen representatives lift $H \to \mathrm{Out}(A_\Theta)$ to $H \to \mathrm{Aut}(A_\Theta)$ as required. □

Claim 9 allows us to use Lemma 7.7, and conclude that the cube complex $X''_{\overline{\Theta}}$ contains an H-fixed point r, with a chosen lift $\widetilde{r}$ in the universal cover.

Let $\widetilde{R}$ denote the standard copy of $\widetilde{X}''_{\partial\Gamma'}$ in $\widetilde{X}''_{\overline{\Theta}}$ determined by $\widetilde{r}$. Its projection R in $X_{\overline{\Theta}}$ is preserved by H, since r is.

On the other hand, let $\widetilde{Z}$ be the standard copy of $\widetilde{X}'_{\partial\Gamma'}$ in $\widetilde{X}'_{\Gamma'}$; this is unique since $\mathrm{lk}_{\Gamma'}(\partial\Gamma') = \emptyset$ by the assumption of Case 1. The standard copies $\widetilde{R}$ and $\widetilde{Z}$ are naturally isometric since both are standard copies of the same cube complex. We now form $\widetilde{X}_\Gamma$ by gluing $\widetilde{X}'_{\Gamma'}$ to $\widetilde{X}''_{\overline{\Theta}}$ via the natural isomorphism $\widetilde{Z} = \widetilde{R}$.

Claim 10. The H-action on X_Γ inherited from the gluing is the correct one.

Proof. Take $h \in H$. Let us pick a basepoint $\widetilde{p} \in \widetilde{Z} = \widetilde{R}$, and let p denote its projection in $Z = R$. Note that choosing $\widetilde{p}$ fixes an isomorphism between fundamental groups of $Z, X'_{\Gamma'}$, and $X''_{\overline{\Theta}}$ (based at p), and the groups $A_{\partial\Gamma'}, A_{\Gamma'}$, and $A_{\overline{\Theta}}$ respectively.

Choose a path $\gamma(h)$ in Z connecting p to $h.p$. Let $h_{\partial\Gamma'} \in \mathrm{Aut}(A_{\partial\Gamma'})$ be the geometric representative of the restriction of $\varphi(h)$. Note that we can construct any representative of the restriction of $\varphi(h)$ this way, so let us choose the path so that $h_{\partial\Gamma'}$ is equal to the restriction of $h_2 \in \mathrm{Aut}(A_\Gamma)$ from the previous claim.

Using the same basepoint and path we obtain geometric representatives $h_{\overline{\Theta}} \in \mathrm{Aut}(A_{\overline{\Theta}})$ and $h_{\Gamma'} \in \mathrm{Aut}(A_{\Gamma'})$. Each of these representatives can be (algebraically) extended to a representative of $\varphi(h)$ in $\mathrm{Aut}(A_\Gamma)$, but this is in general not unique; two such extensions will differ by conjugation by an element of the centraliser of $A_{\overline{\Theta}}$ and $A_{\Gamma'}$ respectively.

Since $\mathrm{lk}(\partial\Gamma') \cap \Gamma' = \emptyset$ (which is the assumption of Case 1 we are in), $Z(\Gamma') \subseteq \partial\Gamma'$. Therefore each vertex in Θ is connected to each vertex in $Z(\Gamma')$ by a single edge. Hence, if $Z(\Gamma')$ is non-empty, then Γ is a cone (and hence a join) over any vertex of it. As this would be a contradiction, we see that Γ' has trivial centre. It also has a trivial link, as we have shown in Claim 8. Therefore there is a unique way of extending $h_{\Gamma'}$ to an automorphism of A_Γ; we will continue to denote this extension by $h_{\Gamma'}$.

Note that $h_{\Gamma'}$ and h_2 both preserve $A_{\Gamma'}$, and by construction they agree when restricted to $\mathrm{Aut}(A_{\partial\Gamma'})$. These two facts imply that $h_{\Gamma'} {h_2}^{-1} \in \mathrm{Aut}(A_\Gamma)$ is equal to the conjugation by an element of $A_{\Gamma'} \cap C(A_{\partial\Gamma'}) = Z(A_{\partial\Gamma'})$.

Now let us look more closely at $h_{\overline{\Theta}}$. When restricted to $A_{\partial\Gamma'}$, it is equal to the restriction of h_2. The same is true when restricted to A_Θ, since in this case this is exactly the geometric representative constructed using $\widetilde{r}$ and no path at all. Hence, picking any representative $h_3 \in \mathrm{Aut}(A_\Gamma)$ of $h_{\overline{\Theta}}$, we see that $h_3 {h_2}^{-1}$ is equal to the conjugation by an element of $Z(A_{\overline{\Theta}}) = A_{Z(\partial\Gamma')} \times A_{Z(\Theta)}$. Hence the identical statement holds for $h_{\Gamma'} {h_3}^{-1}$.

Let us now go back to the action of h on the glued-up complex X_Γ. Again we use $\widetilde{p}$ and $\gamma(h)$ to obtain a geometric representative in $\mathrm{Aut}(A_\Gamma)$. On the subgroup $A_{\Gamma'}$ this automorphism is equal to $h_{\Gamma'}$. On the subgroup $A_{\overline{\Theta}}$, it is equal to h_3, but also to $h_{\Gamma'}$, since conjugation by elements in $Z(A_{\overline{\Theta}})$ is trivial here. Thus h acts as the outer automorphism $\varphi(h)$, as required. □

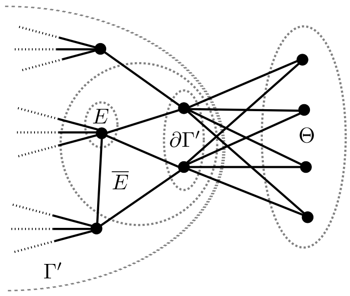

FIGURE 10.2. The relevant subgraphs of Γ in case 2

The System Intersection Axiom gives a standard copy $\widetilde{Z}_\Sigma$ of $\widetilde{X}'_\Sigma$ in $\widetilde{X}'_{\Gamma'}$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$. Note that this includes $\Delta \in \mathcal{S}_{\Gamma'}$; we have $\widetilde{Z} = \widetilde{Z}_\Delta$ since $\mathrm{lk}_{\Gamma'}(\Delta) \subseteq \mathrm{lk}_{\Gamma'}(\partial\Gamma') = \emptyset$ (by the assumption of Case 1 we are in), and so there is only one standard copy of $\widetilde{X}'_\Delta$ in $\widetilde{X}'_{\Gamma'}$.

We have thus completed the construction in Case 1.

Case 2: $\mathrm{lk}(\partial\Gamma') \cap \Gamma' \neq \emptyset$. Let

$$E = \mathrm{lk}(\partial\Gamma')_1$$

be the part of the link of $\partial\Gamma$ lying in Γ'. We define

$$\overline{E} = E * \partial\Gamma'$$

analogous to the definition of $\overline{\Theta} = \Theta * \partial\Gamma'$. In general we put

$$\overline{\Sigma} = \Sigma \cup \partial\Gamma'$$

for any subgraph $\Sigma \subseteq \Gamma$.

Note that we may assume that $\overline{E} \neq \Gamma'$, as otherwise Γ is a join, namely $\mathrm{st}(\partial\Gamma')$. We have $\partial\Gamma' = \mathrm{lk}(\Theta \cup E)$ and thus $\partial\Gamma' \in \mathrm{L}^\varphi$, since it is a link of a non-cone. Thus $\overline{E} = \mathrm{st}(\partial\Gamma')_1 \in \mathrm{L}^\varphi$, by part (4) of Lemma 4.3 and by Lemma 4.2.

Claim 11. There exists a cubical system $\mathcal{X}''$ for $\mathrm{L}^\varphi_{\mathrm{st}(\partial\Gamma')}$ which extends the subsystem $\mathcal{X}'_{\overline{E}}$ of $\mathcal{X}'$ strongly.

Proof. We begin by applying Lemma 8.5, using the decomposition

$$\mathrm{st}(\partial\Gamma') = \partial\Gamma' * (E \cup \Theta)$$

to obtain a cubical system $\mathcal{X}''$ for $\mathrm{L}^\varphi_{\mathrm{st}(\partial\Gamma')}$ which extends the subsystem $\mathcal{X}'_{\overline{E}}$ of $\mathcal{X}'$, in such a way that it extends $\mathcal{X}'_{\partial\Gamma'}$ strongly. If the join decomposition (compare Definition 2.4) of E does not contain singletons, then in fact $\mathcal{X}''$ extends $\mathcal{X}'_{\overline{E}}$ strongly by definition.

Suppose that the decomposition does contain singletons; note that each such singleton belongs to L^φ, since φ is link-preserving. Moreover, such a singleton is also a singleton in the join decomposition of $\overline{E} = E * \partial\Gamma'$.

Let s be such a singleton which is not connected to every other vertex in Γ'. Then Lemma 7.3 implies that the action on the associated loop in $\mathcal{X}'$ and $\mathcal{X}''$ is the same.

The remaining case occurs when E contains a singleton s in its join decomposition, and $\mathrm{st}(s) = \Gamma'$. In this situation we will replace $\mathcal{X}'$ by the system given by Lemma 5.10, so that $\mathcal{X}'_{\{s\}}$ and $\mathcal{X}''_{\{s\}}$ are H-equivariantly isometric; we will continue to denote this slightly altered system by $\mathcal{X}'$. Repeating this operation for each singleton as described guarantees that $\mathcal{X}''$ extends $\mathcal{X}'_{\overline{E}}$ strongly. ☐

We need to introduce one more graph. Let

$$\Delta' = \bigcap_{\Sigma \in \mathcal{S}} \mathrm{st}(\Sigma_1)_1$$

Note that $\Delta' \subseteq \Gamma'$. The significance of Δ' will be explained shortly. For now let us observe the following.

Claim 12. $\Delta' \cup \Theta \in \mathrm{L}^{\varphi}$.

Proof. We have

$$\mathrm{st}(\Sigma_1)_1 \cup \Theta = \mathrm{st}(\Sigma_1) \cup \Sigma \in \mathrm{L}^{\varphi}$$

since $\mathrm{lk}(\mathrm{st}(\Sigma_1) \cap \Sigma) \subseteq \mathrm{lk}(\Sigma_1) \subseteq \mathrm{st}(\Sigma_1) = \mathrm{st}(\mathrm{st}(\Sigma_1))$, and by part ii) of Lemma 4.2. Thus

$$\Delta' \cup \Theta = \bigcap_{\Sigma \in \mathcal{S}} (\mathrm{st}(\Sigma_1)_1 \cup \Theta) \in \mathrm{L}^{\varphi}$$

by part i) of Lemma 4.2. □

By construction we have $\overline{\Theta} \subseteq \Delta' \cup \Theta$ and so $\Delta' \in \mathcal{S}_{\Gamma'}$. We also have $\Delta' \subseteq \overline{E}$ since $\overline{E} \in \mathcal{S}_{\Gamma'}$ satisfies $\mathrm{st}(\overline{E})_1 = \overline{E}$.

Claim 13. There exist fixed standard copies $\widetilde{R}_{\Delta}$ and $\widetilde{R}_{\Delta'}$ of, respectively, $\widetilde{X}''_{\Delta}$ and $\widetilde{X}''_{\Delta'}$ in $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$ such that $\widetilde{R}_{\Delta} \subseteq \widetilde{R}_{\Delta'}$.

Proof. We have

$$\mathrm{lk}(\Delta' \cup \Theta) \subseteq \mathrm{lk}(\Theta) \smallsetminus \Delta' \subseteq \partial\Gamma' \smallsetminus \partial\Gamma' = \emptyset$$

Thus there is a unique (and therefore fixed) standard copy of $\widetilde{X}''_{\Delta' \cup \Theta}$ in $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$. Since

$$\mathrm{lk}(\overline{E}) = \mathrm{lk}(\partial\Gamma') \cap \mathrm{lk}(E) \subseteq \mathrm{lk}(\partial\Gamma') \smallsetminus (E \cup \Theta) = \emptyset$$

there is a unique (and therefore again fixed) standard copy of $\widetilde{X}''_{\overline{E}}$ in $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$. These two copies intersect in a standard copy of $\widetilde{X}''_{\Delta'}$ (by the Matching Property); let us denote it by $\widetilde{R}_{\Delta'}$. Since the two copies are fixed, so is $\widetilde{R}_{\Delta'}$.

By noting that

$$\mathrm{lk}(\Delta \cup \Theta) \subseteq \mathrm{lk}(\Theta) \smallsetminus \Delta \subseteq \partial\Gamma' \smallsetminus \partial\Gamma' = \emptyset$$

and applying the same procedure we obtain a fixed standard copy $\widetilde{R}_{\Delta}$ of $\widetilde{X}''_{\Delta}$ in $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$.

Since the unique standard copy of $\widetilde{X}''_{\Delta \cup \Theta}$ must be contained in the unique standard copy of $\widetilde{X}''_{\Delta' \cup \Theta}$ (by the Composition Property), we have

$$\widetilde{R}_{\Delta} \subseteq \widetilde{R}_{\Delta'}$$

as claimed. □

The System Intersection Axiom gives a standard copy $\widetilde{Y}_{\Sigma}$ of $\widetilde{X}'_{\Sigma}$ in $\widetilde{X}'_{\Gamma'}$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$ such that all these copies intersect in $\widetilde{Y}_{\Delta}$. For each $\Sigma \in \mathcal{S}_{\Gamma'}$ we have

$$\mathrm{lk}_{\Gamma'}(\mathrm{st}(\Sigma)_1) \subseteq \mathrm{lk}_{\Gamma'}(\Sigma) \smallsetminus \mathrm{st}(\Sigma)_1 \subseteq \mathrm{lk}(\Sigma)_1 \smallsetminus \mathrm{lk}(\Sigma)_1 = \emptyset$$

and thus the copy $\widetilde{Y}_{\mathrm{st}(\Sigma)_1}$ is fixed. Note that $\Delta' \in \mathcal{S}_{\Gamma'}$, and by the Intersection Axiom we have

$$\widetilde{Y}_{\Delta'} = \bigcap_{\Sigma \in \mathcal{S}_{\Gamma'}} \widetilde{Y}_{\mathrm{st}(\Sigma)_1}$$

Thus, the standard copy $\widetilde{Y}_{\Delta'}$ is fixed.

Let us now observe the crucial property of Δ'.

Claim 14. Let $\Sigma \in \mathcal{S}$, and let $\widetilde{p} \in \widetilde{Y}_{\Delta'}$. Then there exists a standard copy $\widetilde{W}_{\Sigma_1}$ of $\widetilde{X}'_{\Sigma_1}$ in $\widetilde{X}'_{\Gamma'}$ such that $\widetilde{p} \in \widetilde{W}_{\Sigma_1}$.

Proof. Note that we have a standard copy $\widetilde{Y}_{\mathrm{st}(\Sigma_1)_1}$ which contains $\widetilde{Y}_{\Delta'}$, and hence $\widetilde{p}$. We have

$$\mathrm{st}(\Sigma_1)_1 = \Sigma_1 * \mathrm{lk}_{\Gamma'}(\Sigma_1)$$

and so the Product Axiom and Composition Axiom tell us that there exists a standard copy $\widetilde{W}_{\Sigma_1}$ as required. □

We obtain a complex $\widetilde{C}$ from $\widetilde{X}'_{\Gamma'}$ and $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$ by gluing $\widetilde{Y}_{\Delta'} = \widetilde{R}_{\Delta'}$. However, the projection C of our complex $\widetilde{C}$ might not yet realise the desired action of H.

Pick $h \in H$. Take a point $\widetilde{r} \in \widetilde{R}_\Delta$ as a basepoint, let r be its projection as usual. Take a path $\gamma(h)$ in R_Δ from r to $h.r$. Note that $\widetilde{r}$ is a point in $\widetilde{R}_{\Delta'}$, and hence in $\widetilde{C}$ (after the gluing). Let $\widetilde{p} \in \widetilde{Y}_{\Delta'} = \widetilde{R}_{\Delta'}$ be the corresponding point; we view it as a point in $\widetilde{C}$ as well, and denote its projection by p as usual.

We obtain two geometric representatives of $\varphi(h)$ this way, h_r obtained using r and $\gamma(h)$, and h_p obtained using p and the corresponding path.

Claim 15. The gluing $\widetilde{C}$ is faulty within $Z(A_{\Delta'})$.

Proof. Let us first remark that $\mathrm{lk}(\Sigma) \subseteq \mathrm{lk}(\Theta) \smallsetminus \partial\Gamma' = \emptyset$ for all $\Sigma \in \mathcal{S}$.

By construction we see that

$$h_r(A_{\Delta\cup\Theta}) = A_{\Delta\cup\Theta}$$

Take $\Sigma \in \mathcal{S}$. Since Σ lies in L^φ, we have

$$h_r(A_\Sigma) = {A_\Sigma}^y$$

for some element $y \in A_\Gamma$. But then

$$A_{\Delta\cup\Theta} = h_r(A_{\Delta\cup\Theta}) \leqslant h_r(A_\Sigma) = {A_\Sigma}^y$$

Proposition 2.5 implies that $y \in N(A_\Sigma)N(A_{\Delta\cup\Theta}) = A_\Sigma A_{\Delta\cup\Theta} = A_\Sigma$, and thus

$$h_r(A_\Sigma) = A_\Sigma$$

Now $\Sigma_1 \in \mathrm{L}^\varphi$ and so

$$h_r(A_{\Sigma_1}) = {A_{\Sigma_1}}^z$$

We have ${A_{\Sigma_1}}^z \subseteq A_\Sigma$, and so (using Proposition 2.5 again) we get, without loss of generality, $z \in N(A_\Sigma) = A_\Sigma$ since $\mathrm{lk}(\Sigma) = \emptyset$ (as $\Sigma \in \mathcal{S}$).

Let us now focus on h_p. Claim 14 tells us that there exists a standard copy of $\widetilde{X}'_{\Sigma_1}$ containing $\widetilde{p}$; it is clear that it will also contain the copy of the path $\gamma(h)$. Hence

$$h_p(A_{\Sigma_1}) = A_{\Sigma_1}$$

and so

$${A_{\Sigma_1}}^z = h_r(A_{\Sigma_1}) = c(x(h)^{-1})h_p(A_{\Sigma_1}) = {A_{\Sigma_1}}^{x(h)^{-1}}$$

where

$$c(x(h)) = h_p h_r^{-1}$$

and $x(h) \in C(A_{\Delta'})$.

The construction of h_r also tells us that $A_{\overline{E}} = h_r(A_{\overline{E}})$, since $\widetilde{r}$ and $\gamma(h)$ lie in the fixed standard copy of $\widetilde{X}''_{\overline{E}}$ in $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$. We have $A_{\overline{E}} = h_p(A_{\overline{E}})$ as well, since $\overline{E} \in \mathcal{S}_{\Gamma'}$. So

$$x(h) \in N(A_{\overline{E}}) \leqslant A_{\Gamma'}$$

Thus, by Proposition 2.5, $z \in N(A_{\Sigma_1})A_{\Gamma'}$.

We claim that ${A_{\Sigma_1}}^z = A_{\Sigma_1}$. If $\Theta \cap \mathrm{lk}(\Sigma_1) \neq \emptyset$, then $\Sigma_1 = \partial\Gamma'$, and so

$$z \in A_\Sigma \leqslant N(A_{\Sigma_1})$$

which yields the desired conclusion.

Otherwise we have $\mathrm{st}(\Sigma_1) \subseteq \Gamma'$, and so $N(A_{\Sigma_1}) \leqslant A_{\Gamma'}$, which in turn implies $z \in A_{\Gamma'}$. Now $z \in A_\Sigma \cap A_{\Gamma'} = A_{\Sigma_1}$ and so ${A_{\Sigma_1}}^z = A_{\Sigma_1}$ as claimed.

We have

$${A_{\Sigma_1}}^{x(h)^{-1}} = {A_{\Sigma_1}}^z = A_{\Sigma_1}$$

for all $\Sigma \in \mathcal{S}$. Hence

$$x(h) \in \bigcap_{\Sigma \in \mathcal{S}} A_{\mathrm{st}(\Sigma_1)}$$

We have already shown that $x(h) \in A_{\Gamma'}$ and so

$$x(h) \in \bigcap_{\Sigma \in \mathcal{S}} A_{\mathrm{st}(\Sigma_1)_1} = A_{\Delta'}$$

Therefore $x(h) \in C(A_{\Delta'}) \cap A_{\Delta'} = Z(A'_\Delta)$. This statement holds for each h, and thus the fault of our gluing satisfies the claim. □

Now we are in a position to apply Proposition 7.5 and obtain a new glued up complex, which we call $\widetilde{X}_\Gamma$, which realises our action φ.

Recall that we have a standard copy $\widetilde{R}_\Delta$ in $\widetilde{X}''_{\mathrm{st}(\partial\Gamma')}$. The gluing sends $\widetilde{R}_\Delta$ to some standard copy of $\widetilde{X}'_\Delta$ in $\widetilde{X}'_{\Gamma'}$, which lies within $\widetilde{Y}_{\Delta'}$ (we are using the Composition Property here); let us denote this standard copy by $\widetilde{Z}_\Delta$. It is fixed since $\widetilde{R}_\Delta$ is fixed.

By Claim 14, we may pick a standard copy $\widetilde{Z}_\Sigma$ of $\widetilde{X}'_\Sigma$ in $\widetilde{X}'_{\Gamma'}$ for each $\Sigma \in \mathcal{S}_{\Gamma'}$ such that they will all intersect in $\widetilde{Z}_\Delta$. Let us choose such a family of standard copies.

To finish the construction in this case we need to remark that the complex $\widetilde{X}_\Gamma$ we constructed is equal to a complex obtained from $\widetilde{X}'_{\Gamma'}$ and $\widetilde{X}''_{\Delta\cup\Theta}$ by gluing $\widetilde{R}_\Delta$ and $\widetilde{Z}_\Delta$.

Step 2: Constructing X_Σ for $\Sigma \subseteq \Gamma'$ or $\Sigma \subseteq \Delta \cup \Theta$. Since $\mathcal{X}'$ and $\mathcal{X}''_{\Delta\cup\Theta}$ strongly extend $\mathcal{X}'_\Delta$, and $\Delta = \Gamma' \cap (\Delta \cup \Theta)$, we simply define the complexes in $\mathcal{X}$ for graphs $\Sigma \subseteq \Gamma'$ or $\Sigma \subseteq \Delta \cup \Theta$ to be the ones in $\mathcal{X}'$ or $\mathcal{X}''_{\Delta\cup\Theta}$, respectively.

Interlude. Before we begin Step 3, we record the following

Claim 16. For any graph $\Sigma \in \mathrm{L}^\varphi$ such that $\Sigma \not\subseteq \Gamma'$ and $\Sigma \not\subseteq \overline{\Theta}$ we have $\Theta \subseteq \Sigma$.

Additionally, every graph $\Sigma \in \mathrm{L}^\varphi$ with $\Theta \subseteq \Sigma$ satisfies

$$(*) \qquad \mathrm{lk}(\partial\Gamma') \cap (\Theta \cup \Delta) \subseteq \Sigma$$

In particular, $\overline{\Sigma} = \Sigma \cup \Delta$ for subgraphs Σ with $(*)$, and so $\overline{\Sigma} \in \mathcal{S}$ in this case.

Proof. First assume that $\Sigma \not\subseteq \Gamma'$ and $\Sigma \not\subseteq \overline{\Theta}$. In this case $\mathrm{lk}(\Sigma_1) \subseteq \Gamma'$, since otherwise $\Sigma_1 \subseteq \partial\Gamma'$, which would force $\Sigma \subseteq \overline{\Theta}$, a contradiction. Therefore, by Lemma 4.2, $\overline{\Sigma} = \Sigma \cup \Gamma' \in \mathrm{L}^\varphi$. Thus

$$\Theta \subseteq \Sigma$$

since there is no subgraph in L^φ which is properly contained in Γ and properly contains Γ' (recall that Γ' is maximal among proper subgraphs of Γ in L^φ, and this is why we had to first check that $\overline{\Sigma} \in \mathrm{L}^\varphi$).

Now let us suppose that $\Sigma \in \mathrm{L}^\varphi$ satisfies $\Theta \subseteq \Sigma$. Then

$$\mathrm{lk}(\Sigma \cap (\Delta \cup \Theta)) \subseteq \mathrm{lk}(\Theta) \subseteq \Delta \cup \Theta$$

and so $\Sigma \cup \Delta \in \mathrm{L}^\varphi$ by Lemma 4.2.

We now claim that $\mathrm{lk}(\partial\Gamma') \cap \Delta \subseteq \Sigma$. Note that $E \cup \Theta = \mathrm{lk}(\partial\Gamma') \in \mathrm{L}^\varphi$ since φ is link-preserving. Now $\Sigma \cap (E \cup \Theta) \in \mathrm{L}^\varphi$ and thus

$$\partial\Gamma' * (\Sigma \cap (E \cup \Theta)) = \mathrm{st}(\Sigma \cap (E \cup \Theta)) \in \mathrm{L}^\varphi$$

where the equality follows from the observation that $\Theta \subseteq \Sigma \cap (E \cup \Theta)$. But $\partial\Gamma' * (\Sigma \cap (E \cup \Theta)) \in \mathcal{S}$, and therefore

$$\Delta \subseteq \partial\Gamma' * (\Sigma \cap (E \cup \Theta))$$

This in turn implies that

$$\Delta \subseteq \partial\Gamma' \cup \Sigma$$

and so $\Delta \cap \mathrm{lk}(\partial\Gamma') \subseteq \Sigma$. We assumed that $\Theta \subseteq \Sigma$, and so

$$\mathrm{lk}(\partial\Gamma') \cap (\Theta \cup \Delta) \subseteq \Sigma$$

that is Σ satisfies $(*)$.

Lastly, suppose that Σ satisfies $(*)$. Then

$$\overline{\Sigma} = \Sigma \cup \partial\Gamma' = \Sigma \cup (\Delta \smallsetminus (\Delta \cap \mathrm{lk}(\partial\Gamma')) = \Sigma \cup \Delta$$

as required. □

Since $\Sigma \not\subseteq \Delta \cup \Theta$ implies $\Sigma \not\subseteq \overline{\Theta}$, we see that any $\Sigma \in \mathrm{L}^\varphi$ not covered by Step 2 satisfies $(*)$ and so $\overline{\Sigma} \in \mathcal{S}$. Hence given such a Σ we have

$$\Sigma \subseteq \mathrm{st}_{\overline{\Sigma}}(\Sigma) \subseteq \overline{\Sigma}$$

We will deal with each graph in this chain of inclusions in turn.

Step 3: Constructing X_Σ for $\Sigma \in \mathcal{S} \smallsetminus \{\Delta \cup \Theta\}$. In this case

$$\Sigma_1 = \Sigma \cap \Gamma' \in \mathcal{S}_{\Gamma'}$$

and so (by construction) we have $\widetilde{Z}_{\Sigma_1}$ in $\widetilde{X}_\Gamma$ containing $\widetilde{Z}_\Delta$ (which in particular implies that it is fixed). We also have the image under our gluing map $\widetilde{\iota}_{\Delta\cup\Theta,\Gamma}$ of $\widetilde{X}''_{\Delta\cup\Theta} = \widetilde{X}_{\Delta\cup\Theta}$; let us call it $\widetilde{R}_{\Delta\cup\Theta}$. This subcomplex contains $\widetilde{Z}_\Delta$ by construction. We obtain $\widetilde{X}_\Sigma$ from the two complexes by gluing the two instances of $\widetilde{Z}_\Delta$. Its projection carries the desired marking by construction. The action of H is also the desired one; taking any $h \in H$, and looking at a geometric representative obtained by choosing a basepoint and a path in the subcomplex X_Σ, we get an automorphism of A_Γ which preserves A_Σ. This automorphism is a representative of $\varphi(h)$, and so the restriction to A_Σ is the desired one. But this is equal to the geometric representative of the action of h on X_Σ obtained using the same basepoint and path.

Our construction also gives us maps $\widetilde{\iota}_{\Sigma_1,\Sigma}$ and $\widetilde{\iota}_{\Delta\cup\Theta,\Sigma}$. These maps are as required, since $\Delta \cup \Theta$ and Σ_1 both have trivial links in Σ; the former statement is clear, and the latter follows from the observation that

$$\mathrm{lk}_\Sigma(\Sigma_1) \neq \emptyset$$

implies that $\Sigma_1 \subseteq \partial\Gamma'$ and hence $\Sigma_1 = \Delta = \partial\Gamma'$, which in turn gives $\Sigma = \Delta \cup \Theta$, contradicting our assumption.

We also get a map $\widetilde{\iota}_{\Sigma,\Gamma}$, since we define $\widetilde{X}_\Sigma$ as a subcomplex of $\widetilde{X}_\Gamma$. It is as required since $\mathrm{lk}(\Sigma) = \emptyset$ for all $\Sigma \in \mathcal{S}$.

Step 4: Constructing X_Σ for Σ with $(*)$ and such that $\mathrm{lk}_{\overline{\Sigma}}(\Sigma) = \emptyset$. By the Composition Property there exists a standard copy of $\widetilde{X}_{\partial\Gamma'}$ in $\widetilde{X}_{\Delta\cup\Theta}$ which lies within $\widetilde{R}_\Delta$; let us denote it by $\widetilde{R}_{\partial\Gamma'}$. The gluing $\widetilde{R}_\Delta = \widetilde{Z}_\Delta$ gives us the corresponding standard copy $\widetilde{Z}_{\partial\Gamma'}$ in $\widetilde{Z}_\Delta$.

Note that the assumption implies that $\mathrm{lk}_{\overline{\Sigma}_1}(\Sigma_1) = \emptyset$. Let us take the unique standard copy of $\widetilde{X}_{\Sigma_1}$ in $\widetilde{X}_{\overline{\Sigma}_1}$; we will denote it by $\widetilde{Z}_{\Sigma_1}$. By the Matching Property, it intersects $\widetilde{Z}_\Delta$ in a standard copy of $\widetilde{X}_{\Sigma\cap\Delta}$. Using the Matching Property again, this time in $\widetilde{Z}_\Delta$, we see that this copy of $\widetilde{X}_{\Sigma\cap\Delta}$ intersects $\widetilde{Z}_{\partial\Gamma'}$ in a copy of $\widetilde{X}_{\Sigma\cap\partial\Gamma'}$; we will denote thus copy by $\widetilde{Z}_{\Sigma\cap\partial\Gamma'}$, and the corresponding one in $\widetilde{R}_{\partial\Gamma'}$ by $\widetilde{R}_{\Sigma\cap\partial\Gamma'}$.

Recall that $\Sigma_2 = \Sigma \cap (\Delta \cup \Theta)$. Since $\mathcal{X}''$ is a product of $\mathcal{X}''_{\partial\Gamma'}$ and $\mathcal{X}''_{\mathrm{lk}_{\Delta\cup\Theta}(\partial\Gamma')}$, there exists a standard copy of $\widetilde{X}_{\Sigma_2}$ in $\widetilde{X}_{\Delta\cup\Theta}$ which contains $\widetilde{R}_{\Sigma\cap\partial\Gamma'}$; we will denote it by $\widetilde{R}_{\Sigma_2}$. We define $\widetilde{X}_\Sigma$ to be the subcomplex of $\widetilde{X}_{\overline{\Sigma}}$ obtained by gluing $\widetilde{Z}_{\Sigma_1}$ to $\widetilde{R}_{\Sigma_2}$. Note that these two copies overlap in a copy of $\widetilde{X}_{\Sigma\cap\Delta}$, which is the unique such copy containing $\widetilde{R}_{\Sigma\cap\partial\Gamma'}$.

Note that again our gluing procedure determines maps $\widetilde{\iota}_{\Sigma_1,\Sigma}$ and $\widetilde{\iota}_{\Sigma_2,\Sigma}$ of the required type.

From this construction we also obtain a map $\widetilde{\iota}_{\Sigma,\overline{\Sigma}}$, since we define $\widetilde{X}_\Sigma$ as a subcomplex of $\widetilde{X}_{\overline{\Sigma}}$.

Step 5: Constructing the remaining complexes. As remarked above we are left with graphs $\Sigma \in \mathrm{L}^\varphi$ satisfying $(*)$ and such that $\Sigma \subset \mathrm{st}_{\overline{\Sigma}}(\Sigma)$ is a proper subgraph. Let $\Sigma' = \mathrm{st}_{\overline{\Sigma}}(\Sigma)$, and note that Σ' is covered by the previous step. We need to exhibit a product structure on $\widetilde{X}_{\Sigma'}$, one factor of which will be the desired complex for Σ, the other for $\Lambda = \mathrm{lk}_{\overline{\Sigma}}(\Sigma)$.

The complex $\widetilde{X}_{\Sigma'}$ is obtained from complexes $\widetilde{X}_{\Sigma'_1}$ and $\widetilde{X}_{\Sigma'_2}$ by gluing them along a copy of $\widetilde{X}_{\Sigma'\cap\Delta}$. Since $\Lambda \subseteq \partial\Gamma'$, each of these three complexes is a product of $\widetilde{X}_\Lambda$ and some other complex by the Product Axiom. Moreover, the embeddings $\widetilde{X}_{\Sigma'\cap\Delta} \to \widetilde{X}_{\Sigma'_i}$ with $i \in \{1,2\}$ respect the product structure, that is the image of any standard copy of $\widetilde{X}_\Lambda$ in $\widetilde{X}_{\Sigma'\cap\Delta}$ is still a standard copy in $\widetilde{X}_{\Sigma'_i}$ by the Composition Property. Hence the glued-up complex $\widetilde{X}_{\Sigma'}$ is a product of $\widetilde{X}_\Lambda$ and a complex obtained by gluing some standard copies of $\widetilde{X}_{\Sigma_i}$ in $\widetilde{X}_{\Sigma'_i}$ along $\widetilde{X}_{\Sigma\cap\Delta}$; we call this latter complex $\widetilde{X}_\Sigma$. For notational convenience we pick some such standard copies of $\widetilde{X}_{\Sigma_1}$ and $\widetilde{X}_{\Sigma_2}$, and denote them by $\widetilde{Z}_{\Sigma_1}$ and $\widetilde{R}_{\Sigma_2}$ respectively.

The gluing of the standard copies of $\widetilde{X}_{\Sigma_i}$ gives us maps $\widetilde{\iota}_{\Sigma_i,\Sigma}$ for both values of i, which are as required.

The construction also gives us a map

$$\widetilde{\iota}_{\Sigma,\Sigma'} \colon \widetilde{X}_\Sigma \times \widetilde{X}_\Lambda \to \widetilde{X}_{\Sigma'}$$

which is again as wanted.

10.2. Constructing the maps. Let $\Sigma, \Sigma' \in \mathrm{L}^\varphi$ be such that $\Sigma \subseteq \Sigma'$. We need to construct a map $\widetilde{\iota}_{\Sigma,\Sigma'}$. We will do it in several steps.

(1) $\Sigma' = \Sigma'_i$ **for some** $i \in \{1,2\}$

In this case the cube complexes X_Σ and $X_{\Sigma'}$ are obtained directly from another cubical system (in Step 2), and we take $\widetilde{\iota}_{\Sigma,\Sigma'}$ to be the map coming from that system.

We will now assume that the hypothesis of this step is not satisfied, which implies that Σ' satisfies $(*)$ and that $\Sigma' \neq \Delta \cup \Theta$.

(2) $\Sigma = \Sigma_i$ **and** $\Sigma \neq \Sigma_j$ **with** $\{i,j\} = \{1,2\}$

In this case we have $\mathrm{st}(\Sigma) = \mathrm{st}(\Sigma)_i$, since if $\mathrm{lk}(\Sigma) \neq \mathrm{lk}(\Sigma)_i$ then (knowing that $\Sigma = \Sigma_i$) we must have $\Sigma \subseteq \partial\Gamma'$, and thus $\Sigma = \Sigma_j$ which contradicts the assumption. Therefore $\mathrm{lk}_{\Sigma'}(\Sigma) = \mathrm{lk}_{\Sigma'}(\Sigma)_i$ as well. We define

$$\widetilde{\iota}_{\Sigma,\Sigma'} = \widetilde{\iota}_{\Sigma'_i,\Sigma'} \circ \widetilde{\iota}_{\Sigma,\Sigma'_i}$$

where the last map was defined in the previous step, and the map $\widetilde{\iota}_{\Sigma'_i,\Sigma'}$ was constructed together with the complex $\widetilde{X}_{\Sigma'}$ in Step 3, 4 or 5 of Subsection 10.1.

(3) **Σ satisfies $(*)$ and $\Sigma \neq \Delta \cup \Theta$**

Observe that $\mathrm{st}_{\overline{\Sigma}}(\Sigma) = \mathrm{st}_{\overline{\Sigma'}}(\Sigma)$ since $\mathrm{lk}(\Sigma) \subseteq \partial\Gamma'$. Hence Step 5 above gives us the map $\widetilde{\iota}_{\Sigma,\mathrm{st}_{\overline{\Sigma'}}(\Sigma)}$.

Let $\Omega = \mathrm{st}_{\overline{\Sigma'}}(\Sigma) \cup \mathrm{st}_{\overline{\Sigma'}}(\Sigma')$

Claim 17. $\Omega \in \mathrm{L}^{\varphi}$.

Proof. We use part (3) of Lemma 4.2. Thus we only need to observe that

$$\mathrm{lk}\left(\mathrm{st}_{\overline{\Sigma'}}(\Sigma) \cap \mathrm{st}_{\overline{\Sigma'}}(\Sigma')\right) \subseteq \mathrm{lk}(\Sigma) \subseteq \mathrm{st}(\Sigma) = \mathrm{st}_{\overline{\Sigma'}}(\Sigma)$$

since $\mathrm{lk}(\Sigma) \subseteq \partial\Gamma'$ as Σ satisfies $(*)$. □

We have

$$\mathrm{lk}(\mathrm{st}_{\overline{\Sigma}}(\Sigma)) \subseteq \mathrm{lk}(\Sigma) \smallsetminus \mathrm{st}_{\overline{\Sigma}}(\Sigma) = \emptyset$$

since $\mathrm{lk}(\Sigma) \subseteq \partial\Gamma'$ and so $\mathrm{lk}(\Sigma) = \mathrm{lk}_{\overline{\Sigma}}(\Sigma)$. Analogously we have $\mathrm{lk}(\mathrm{st}_{\overline{\Sigma'}}(\Sigma')) = \emptyset$. It immediately follows that $\mathrm{lk}(\Omega) = \emptyset$ as well. Now the complex $\widetilde{X}_{\mathrm{st}_{\overline{\Sigma}}(\Sigma)}$ is formed by gluing $\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma}}(\Sigma)_1}$ and $\widetilde{R}_{\mathrm{st}_{\overline{\Sigma}}(\Sigma)_2}$; we have similar statements for $\widetilde{X}_{\mathrm{st}_{\overline{\Sigma'}}(\Sigma')}$ and $\widetilde{X}_{\Omega}$. By construction (see Step 4), $\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma}}(\Sigma)_1}$ and $\widetilde{Z}_{\Omega_1}$ are both unique standard copies of the corresponding complexes in $\widetilde{Z}_{\overline{\Sigma}_1}$. The Composition Property in $\mathcal{X}'$ implies that there exists a standard copy of $\widetilde{X}_{\mathrm{st}_{\overline{\Sigma}}(\Sigma)_1}$ contained in $\widetilde{Z}_{\Omega_1}$, and thus

$$\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma}}(\Sigma)_1} \subseteq \widetilde{Z}_{\Omega_1}$$

Similarly

$$\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma'}}(\Sigma')_1} \subseteq \widetilde{Z}_{\Omega_1}$$

Now the Matching Property (in $\mathcal{X}'$) tells us that $\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma'}}(\Sigma')_1}$ and $\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma'}}(\Sigma')_1}$ intersect non-trivially. Since both have a product structure, we find standard copies $\widetilde{Z}_{\Sigma_1}$ and $\widetilde{Z}_{\Sigma'_1}$ of $\widetilde{X}_{\Sigma_1}$ and $\widetilde{X}_{\Sigma'_1}$ respectively (in $\widetilde{Z}_{\Omega_1}$) which also intersect non-trivially.

Let us look more closely at the product structure of $\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma'}}(\Sigma)_1}$. It is isomorphic to a product of $\widetilde{X}_{\Sigma_1}$ and $\widetilde{X}_{\mathrm{lk}_{\overline{\Sigma'}}(\Sigma)_1}$. The latter complex contains a standard copy of $\widetilde{X}_{\mathrm{lk}_{\Sigma'}(\Sigma)_1}$. Hence there is a standard copy $\widetilde{Z}_{\mathrm{st}_{\Sigma'}(\Sigma)_1}$ of $\widetilde{X}_{\mathrm{st}_{\Sigma'}(\Sigma)_1}$ in $\widetilde{Z}_{\mathrm{st}_{\overline{\Sigma'}}(\Sigma)_1}$ containing the chosen standard copy $\widetilde{Z}_{\Sigma_1}$ (by the Intersection Axiom). Thus $\widetilde{Z}_{\mathrm{st}_{\Sigma'}(\Sigma)_1}$ and $\widetilde{Z}_{\Sigma'_1}$ intersect non-trivially, and so the intersection Axiom tells us that

$$\widetilde{Z}_{\mathrm{st}_{\Sigma'}(\Sigma)_1} \subseteq \widetilde{Z}_{\Sigma'_1}$$

After the gluing this yields a map $\widetilde{\iota}_{\mathrm{st}_{\Sigma'}(\Sigma),\Sigma'}$, and since $\widetilde{Z}_{\mathrm{st}_{\Sigma}(\Sigma')_1}$ had a product structure, so does $\widetilde{X}_{\mathrm{st}_{\Sigma'}(\Sigma)}$ (arguing as in Step 5 above). This gives us the desired map $\widetilde{\iota}_{\Sigma,\Sigma'}$.

(4) $\Sigma \subseteq \partial\Gamma'$

In this (last) case we observe that $\mathrm{lk}_{\Sigma'}(\Sigma)$ satisfies $(*)$, and so we have already constructed the map

$$\widetilde{\iota}_{\mathrm{lk}_{\Sigma'}(\Sigma),\Sigma'} \colon \widetilde{X}_{\mathrm{lk}_{\Sigma'}(\Sigma)} \times \widetilde{X}_\Sigma \to \widetilde{X}_{\Sigma'}$$

We define $\widetilde{\iota}_{\Sigma,\Sigma'}$ by reordering the factors in the domain.

10.3. Verifying the axioms. The first two axioms depend only on two subgraphs $\Sigma, \Sigma' \in \mathrm{L}^\varphi$ with $\Sigma \subseteq \Sigma'$. This is the same assumption as in the maps part of our proof, and hence the verification of the two axioms will follow the structure as the construction of maps – we will consider four cases, and the assumption in each will be identical to the assumptions of the corresponding case above.

Product Axiom. Suppose that $\Sigma' = \mathrm{st}_{\Sigma'}(\Sigma)$.

(1) In this case the Product Axiom follows from the Product Axiom in $\mathcal{X}'$ or $\mathcal{X}''$. Otherwise we assume that Σ' satisfies $(*)$ and $\Sigma' \neq \Delta \cup \Theta$.
(2) If $\Sigma = \Sigma_i$ and $\Sigma \neq \Sigma_j$ with $\{i,j\} = \{1,2\}$, then $\Sigma' = \mathrm{st}_{\Sigma'}(\Sigma)$ also satisfies $\Sigma' = \Sigma'_i$, and so we are in the previous case.
(3) The standard copy $\widetilde{Z}_{\mathrm{st}_{\Sigma'}(\Sigma)_1}$ used in case (3) above is equal to the image of $\widetilde{\iota}_{\Sigma_1,\mathrm{st}_{\Sigma'}(\Sigma)_1}$, and hence the corresponding statement is still true after the gluing.
(4) In this case we defined the map $\widetilde{\iota}_{\Sigma,\Sigma'}$ using $\widetilde{\iota}_{\mathrm{lk}_{\Sigma'}(\Sigma),\Sigma'}$, and the graph $\mathrm{lk}_{\Sigma'}(\Sigma)$ is covered by the previous cases.

Orthogonal Axiom. Let $\Lambda = \mathrm{lk}_{\Sigma'}(\Sigma)$.

(1) If $\Sigma' \subseteq \Gamma'$ or $\Sigma' \subseteq \Delta \cup \Theta$ then the axiom is satisfied, since it is satisfied in $\mathcal{X}'$ and $\mathcal{X}''$. Otherwise we assume that Σ' satisfies $(*)$ and $\Sigma' \neq \Delta \cup \Theta$.
(2) If $\Sigma = \Sigma_i$ and $\Sigma \neq \Sigma_j$ with $\{i,j\} = \{1,2\}$, then also $\Lambda = \Lambda_i$. Suppose that $\Lambda \neq \Lambda_j$. In this case both maps $\widetilde{\iota}_{\Sigma,\Sigma'}$ and $\widetilde{\iota}_{\Lambda,\Sigma'}$ factorise through $\widetilde{\iota}_{\mathrm{st}_{\Sigma'}(\Sigma),\Sigma'}$, and thus it is enough to verify the axiom within $\widetilde{X}_{\mathrm{st}_{\Sigma'}(\Sigma)}$. But $\mathrm{st}_{\Sigma'}(\Sigma) = \mathrm{st}_{\Sigma'}(\Sigma)_i$ and so we are done by the previous case.

   Now suppose that $\Lambda \subseteq \partial\Gamma'$. In this case the axiom follows trivially from the construction of the map $\widetilde{\iota}_{\Lambda,\Sigma'}$.
(3) In this case we have $\Lambda \subseteq \partial\Gamma'$ and so we are done as above.
(4) When $\Sigma \subseteq \partial\Gamma'$ we are again done by construction.

Intersection Axiom. Let us now verify that $\mathcal{X}$ satisfies the Intersection Axiom. Take $\Sigma, \Sigma', \Omega \in \mathrm{L}^\varphi$ such that $\Sigma \subseteq \Omega$ and $\Sigma' \subseteq \Omega$, and let $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ be standard copies of, respectively, $\widetilde{X}_\Sigma$ and $\widetilde{X}_{\Sigma'}$ in $\widetilde{X}_\Omega$ with non-empty intersection. We need to show that the intersection is the image of a standard copy of $\Sigma \cap \Sigma'$ in each.

As in the first two cases, the details depend on the inclusions $\Sigma, \Sigma' \subseteq \Omega$. The cases will thus be labeled by pairs of integers (n,m), the first determining in which step the map $\widetilde{\iota}_{\Sigma,\Omega}$ was constructed, and the second playing the same role for $\widetilde{\iota}_{\Sigma',\Omega}$. By symmetry we only need to consider $n \leqslant m$.

(1,1) In this case $\Omega = \Omega_i$, and so the axiom follows from the Intersection Axiom in $\mathcal{X}'$. In what follows we can assume that $\Omega \neq \Delta \cup \Theta$ satisfies $(*)$. Hence $\widetilde{X}_\Omega$ is obtained by gluing $\widetilde{Z}_{\Omega_1}$ and $\widetilde{R}_{\Omega_2}$ along a subcomplex of $\widetilde{Z}_\Delta$ which is a standard copy of $\widetilde{X}_{\Delta\cap\Omega}$; let us denote it by $\widetilde{Z}_{\Delta\cap\Omega}$.
(2,2) This splits into two cases. If $\Sigma = \Sigma_i$ and $\Sigma' = \Sigma'_i$ then both maps $\widetilde{\iota}$ factor through $\widetilde{\iota}_{\Omega_i,\Omega}$, and so the problem is reduced to checking the axiom for the triple $\Sigma, \Sigma', \Omega_i$, for which it holds.

   In the other case we have, without loss of generality, $\Sigma = \Sigma_1$ and $\Sigma' = \Sigma'_2$. By construction, the given standard copies $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ must lie within the standard copies $\widetilde{Z}_{\Omega_1}$ and $\widetilde{R}_{\Omega_2}$ respectively, and hence intersect within $\widetilde{Z}_{\Delta\cap\Omega}$.

We use the Intersection Axiom of $\mathcal{X}'$ for $\widetilde{Z}_{\Delta\cap\Omega}$ and $\widetilde{Y}_\Sigma$ inside $\widetilde{Z}_{\Omega_1}$ and see that the two copies intersect in a copy of $\widetilde{X}_{\Delta\cap\Sigma}$, which is also the image of a standard copy of $\widetilde{X}_{\Delta\cap\Sigma}$ in $\widetilde{Z}_{\Delta\cap\Omega}$.

We repeat the argument for Σ' and obtain a standard copy of $\widetilde{X}_{\Delta\cap\Sigma'}$ in $\widetilde{Z}_{\Delta\cap\Omega}$. Now this copy intersects the one of $\widetilde{X}_{\Delta\cap\Sigma}$, and hence, applying the Intersection Axiom again, they intersect in a copy of $\widetilde{X}_{\Delta\cap\Sigma\cap\Sigma'}$ in $\widetilde{Z}_{\Delta\cap\Omega}$. But $\Delta\cap\Sigma\cap\Sigma' = \Sigma\cap\Sigma'$, and so we have found the desired standard copy in $\widetilde{Z}_{\Delta\cap\Omega}$. Now the Composition Property (Lemma 5.7) implies that this is also a standard copy in $\widetilde{X}_\Sigma$, $\widetilde{X}_{\Sigma'}$ and $\widetilde{X}_{\Omega_i}$ for any $i \in \{1,2\}$, and thus this is also a standard copy in $\widetilde{X}_\Omega$ by construction.

(2,3) The non-trivial intersection of any standard copy of $\widetilde{X}_\Sigma$ and any standard copy of $\widetilde{X}_{\Sigma'}$ in $\widetilde{X}_\Omega$ is in fact contained in the standard copy of $\widetilde{X}_{\Omega_i}$, since any copy of $\widetilde{X}_\Sigma$ is contained therein. Therefore the intersection is also contained in a standard copy of $\widetilde{X}_{\Sigma'_i}$ by construction of $\widetilde{\iota}_{\Sigma',\Omega}$. We apply the Intersection Axiom in $\widetilde{X}_{\Omega_i}$, and observe that the standard copy of $\widetilde{X}_{\Sigma\cap\Sigma'}$ in $\widetilde{X}_{\Omega_i}$ obtained this way is also a standard copy in $\widetilde{X}_\Omega$ by construction.

(2,4) In this case $\widetilde{Y}_\Sigma$ must in fact be contained in $\widetilde{Q}$, where $\widetilde{Q}$ is either $\widetilde{Z}_{\Omega_1}$ or $\widetilde{R}_{\Omega_2}$.

We have $\widetilde{Y}_\Sigma \cap \widetilde{Y}_{\Sigma'} \subseteq \widetilde{Q}$, and hence we only need to prove that $\widetilde{Y}_{\Sigma'}$ is a standard copy in $\widetilde{Q}$.

Let $\Lambda = \mathrm{lk}_\Omega(\Sigma')$. Note that $\mathrm{lk}(\partial\Gamma') \cap (\Delta\cup\Theta) \subseteq \Lambda$. By definition of $\widetilde{\iota}_{\Sigma',\Omega}$, we have

$$\widetilde{Y}_{\Sigma'} = \widetilde{\iota}_{\Lambda,\Omega}(\{\widetilde{x}\} \times \widetilde{X}_{\Sigma'})$$

for some point $\widetilde{x} \in \widetilde{X}_\Lambda$. Since $\widetilde{Y}_{\Sigma'}$ contains a point in $\widetilde{Q}$, there exists $\widetilde{y} \in \widetilde{X}_{\Sigma'}$ such that $\widetilde{\iota}_{\Lambda,\Omega}(\widetilde{x},\widetilde{y}) \in \widetilde{Q}$.

If $\Lambda \not\subseteq \Delta\cup\Theta$ then this is only possible if $\widetilde{x}$ lies in $\widetilde{Z}_{\Lambda_i}$ or $\widetilde{R}_{\Lambda_2}$ (depending on what $\widetilde{Q}$ is), by the construction of $\widetilde{X}_\Lambda$ and $\widetilde{\iota}_{\Lambda,\Omega}$. But then, again by the construction of $\widetilde{X}_\Lambda$, we have $\widetilde{Y}_{\Sigma'} \subseteq \widetilde{Q}$ being a standard copy as claimed.

We still need to check what happens when $\Lambda \subseteq \Delta\cup\Theta$. Suppose that $\widetilde{Q} = \widetilde{R}_{\Omega_2}$. In this case $\widetilde{Y}_{\Sigma'}$ is a standard copy in $\widetilde{Q}$ by the Composition Property of $\mathcal{X}''$. Lastly, let us suppose that $\widetilde{Q} = \widetilde{Z}_{\Omega_1}$. Since $\mathrm{im}(\widetilde{\iota}_{\Lambda,\Omega}) \subseteq \widetilde{R}_{\Omega_2}$ by construction, the Intersection Axiom in $\mathcal{X}''$ tells us that $\widetilde{Y}_{\Sigma'}$ is a standard copy in $\widetilde{Z}_{\Omega\cap\Delta}$. But then it is also a standard copy in $\widetilde{Q} = \widetilde{Z}_{\Omega_1}$ by the Composition Property in $\mathcal{X}'$.

(3,3) In this case $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ are obtained from $\widetilde{Z}_{\Sigma_1}$, $\widetilde{R}_{\Sigma_2}$, and $\widetilde{Z}_{\Sigma'_1}$, $\widetilde{R}_{\Sigma'_2}$ respectively.

Suppose that $\widetilde{Z}_{\Sigma_1}$ and $\widetilde{Z}_{\Sigma'_1}$ do not intersect. Then $\widetilde{R}_{\Sigma_2}$ and $\widetilde{R}_{\Sigma'_2}$ do intersect, and the Intersection Axiom in $\mathcal{X}''$ tells us that they intersect in a standard copy of $\widetilde{X}_{\Sigma_2\cap\Sigma'_2}$. The graph $\Sigma_2\cap\Sigma'_2$ satisfies $(*)$, and so $(\Sigma_2\cap\Sigma'_2)\cup\Delta = \Delta\cup\Theta$. This implies that the standard copy of $\widetilde{X}_{\Sigma_2\cap\Sigma'_2}$ intersects $\widetilde{Z}_\Delta$ non-trivially (by the Matching Property in $\mathcal{X}''$). But this intersection lies in $\widetilde{Z}_{\Sigma_1}$ and $\widetilde{Z}_{\Sigma'_1}$, and hence they did intersect.

Now the Intersection Axiom in $\mathcal{X}'$ tells us that $\widetilde{Z}_{\Sigma_1}$ and $\widetilde{Z}_{\Sigma'_1}$ intersect in a standard copy of $\widetilde{X}_{\Sigma_1\cap\Sigma'_1}$; let us call it $\widetilde{Z}_{\Sigma_1\cap\Sigma'_1}$.

Note that the standard copy $\widetilde{Z}_{\overline{\Sigma}_1}$ which contains $\widetilde{Z}_{\Sigma_1}$ by construction; similarly $\widetilde{Z}_{\overline{\Sigma'}_1}$ contains $\widetilde{Z}_{\Sigma'_1}$. Now $\widetilde{Z}_{\overline{\Sigma}_1}$ and $\widetilde{Z}_{\overline{\Sigma'}_1}$ intersect in the copy $\widetilde{Z}_{\overline{\Sigma}_1 \cap \overline{\Sigma'}_1}$, which contains $\widetilde{Z}_{\Sigma_1 \cap \Sigma'_1}$, since $\widetilde{Z}_{\Sigma_1 \cap \Sigma'_1}$ lies in both $\widetilde{Z}_{\overline{\Sigma}_1}$ and $\widetilde{Z}_{\overline{\Sigma'}_1}$. The copy $\widetilde{Z}_{\overline{\Sigma}_1 \cap \overline{\Sigma'}_1}$ intersects $\widetilde{Z}_\Delta$, and so the Matching Property implies that $\widetilde{Z}_{\Sigma_1 \cap \Sigma'_1}$ intersects $\widetilde{Z}_\Delta$ as well. The Intersection Axiom implies that this intersection is a copy of $\widetilde{X}_{\Sigma_1 \cap \Sigma'_1 \cap \Delta}$. This in turn implies that $\widetilde{R}_{\Sigma_2}$ and $\widetilde{R}_{\Sigma'_2}$ intersect, and we have already shown above that in this case they intersect in a standard copy of $\widetilde{X}_{\Sigma_2 \cap \Sigma'_2}$. The union of this copy with $\widetilde{Z}_{\Sigma_1 \cap \Sigma'_1}$ is by construction a standard copy of $\widetilde{X}_{\Sigma \cap \Sigma'}$ in $\widetilde{X}_\Omega$, and again by construction it is the image of a standard copy of $\widetilde{X}_{\Sigma \cap \Sigma'}$ in $\widetilde{X}_\Sigma$ and $\widetilde{X}_{\Sigma'}$.

(3,4) The standard copy $\widetilde{Y}_\Sigma$ is obtained from $\widetilde{Z}_{\Sigma_1}$ and $\widetilde{R}_{\Sigma_2}$. If $\widetilde{Y}_{\Sigma'}$ intersects $\widetilde{Z}_{\Sigma_1}$, then we apply case (2,4) to the triple $\Sigma_1, \Sigma', \Omega$, and see that $\widetilde{Y}_{\Sigma'}$ intersects $\widetilde{Z}_{\Sigma_1}$ in a copy of $\widetilde{X}_{\Sigma_1 \cap \Sigma'} = \widetilde{X}_{\Sigma \cap \Sigma'}$.

If $\widetilde{Y}_{\Sigma'}$ intersects $\widetilde{R}_{\Sigma_2}$, then we apply case (2,4) to the triple $\Sigma_2, \Sigma', \Omega$, and see that $\widetilde{Y}_{\Sigma'}$ intersects $\widetilde{R}_{\Sigma_2}$ in a copy of $\widetilde{X}_{\Sigma_1 \cap \Sigma'} = \widetilde{X}_{\Sigma \cap \Sigma'}$.

If $\widetilde{Y}_{\Sigma'}$ intersects both $\widetilde{Z}_{\Sigma_1}$ and $\widetilde{R}_{\Sigma_2}$, then the two copies we obtained intersect. But they are copies of the same complex, and hence they coincide.

(4,4) In case (2,4) we have shown that if $\widetilde{Y}_{\Sigma'}$ intersects $\widetilde{Z}_{\Omega_1}$, then it lies within as a standard copy; the analogous statement holds for $\widetilde{R}_{\Omega_2}$. Now the standard copies $\widetilde{Y}_\Sigma$ and $\widetilde{Y}_{\Sigma'}$ intersect, and hence they both lie in $\widetilde{Q}$ as standard copies, where $\widetilde{Q}$ is $\widetilde{Z}_{\Omega_1}$ or $\widetilde{R}_{\Omega_2}$. But now we just need to apply the Intersection Axiom in $\mathcal{X}'$ or $\mathcal{X}''$.

System Intersection Axiom. Take a subsystem $\mathbb{P} \subseteq \mathrm{L}^\varphi$ closed under taking unions. If all elements of $\mathbb{P}$ lie in Γ' or in $\Delta \cup \Theta$, then we are done (from the System Intersection Axiom of $\mathcal{X}'$ or $\mathcal{X}''$). So let us suppose this is not the case, that is suppose that there exists $\Sigma \in \mathbb{P}$ satisfying $(*)$ and $\Sigma \neq \Delta \cup \Theta$. Hence $\bigcup \mathbb{P} \neq \Delta \cup \Theta$ satisfies $(*)$.

Define

$$\begin{aligned} \mathbb{P}' &= \{\Sigma \in \mathbb{P} \mid \Theta \subseteq \Sigma\} \\ \overline{\mathbb{P}}' &= \{\overline{\Sigma} = \Sigma \cup \partial\Gamma' \mid \Sigma \in \mathbb{P}'\} \end{aligned}$$

and $\overline{\mathbb{P}} = \mathbb{P} \cup \overline{\mathbb{P}}'$. Observe that $\overline{\mathbb{P}}' \subseteq \mathrm{L}^\varphi$, since all graphs in $\mathbb{P}'$ satisfy $(*)$, and so for any $\Sigma' \in \mathbb{P}'$ we have $\overline{\Sigma'} \in \mathcal{S} \subseteq \mathrm{L}^\varphi$.

Claim 18. $\overline{\mathbb{P}}$ is closed under taking unions.

Proof. Take $\Sigma, \Sigma' \in \overline{\mathbb{P}}$. If both lie in $\mathbb{P}$ then we are done. Let us first suppose that $\Sigma \in \mathbb{P}$ and $\Sigma' \in \overline{\mathbb{P}}'$. Then $\Sigma' = \overline{\Sigma''}$ for some $\Sigma'' \in \mathbb{P}'$. Now

$$\Sigma \cup \Sigma' = \Sigma \cup (\Sigma'' \cup \partial\Gamma') = (\Sigma \cup \Sigma'') \cup \partial\Gamma' \in \overline{\mathbb{P}}'$$

since $\Theta \subseteq \Sigma \cup \Sigma'' \in \mathbb{P}$. If both $\Sigma, \Sigma' \in \overline{\mathbb{P}}'$ then an analogous argument shows that $\Sigma \cup \Sigma' \in \overline{\mathbb{P}}'$. □

Now observe that $\overline{P}'$ is a subsystem of $\overline{\mathbb{P}}$, which is closed under taking unions. Hence the same is true for systems $(\overline{P})_{\Gamma'}$ and $(\overline{P}')_{\Gamma'}$.

We are now going to construct standard copies for all elements in $\mathbb{P}$, such that they all intersect non-trivially.

Observe that $\overline{P}' \subseteq \mathcal{S}$, and so for each element $\overline{\Sigma} \in (\overline{P}')_{\Gamma'}$ we are given a standard copy $\widetilde{Z}_{\overline{\Sigma}}$ in $\widetilde{X}_{\Gamma'}$ which contains $\widetilde{Z}_\Delta$.

Let $\Sigma \in \overline{P}$. Then $\Sigma \cup \bigcap \overline{P}' \in \overline{P}'$, and hence $\Sigma_1 \cup \bigcap (\overline{P}')_{\Gamma'} \in (\overline{P}')_{\Gamma'}$, and so we are able to apply Lemma 5.5 to the collection of standard copies we just discussed, and extend it by adding copies $\widetilde{Z}_{\Sigma_1}$ of $\widetilde{X}_{\Sigma_1}$ with $\Sigma \in \overline{P}$, such that all these copies intersect non-trivially. Moreover, for every $\Sigma \in \mathbb{P}'$, the copy $\widetilde{Z}_{\Sigma_1}$ will intersect $\widetilde{Z}_\Delta$ (thanks to the Matching Property in $\widetilde{Z}_{\overline{\Sigma}_1}$).

By this point we have constructed standard copies $\widetilde{Z}_{\Sigma_1}$ for each $\Sigma \in \mathbb{P}$ which all intersect non-trivially. We will now extend these copies to copies of $\widetilde{X}_\Sigma$.

Let $\Sigma \in \mathbb{P}'$. If $\Sigma \not\subseteq \Delta \cup \Theta$, then we can extend $\widetilde{Z}_{\Sigma_1}$ to a standard copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_\Gamma$ by construction.

When $\Sigma \subseteq \Delta \cup \Theta$ (but still $\Sigma \in \mathbb{P}'$) we need to show that there exists a standard copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_\Gamma$ which contains $\widetilde{Z}_{\Sigma_1}$. By construction it is enough to find such a copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_{\Delta\cup\Theta}$, bearing in mind that $\widetilde{Z}_{\Sigma_1}$ lies in $\widetilde{Z}_\Delta = \widetilde{R}_\Delta$.

By $(*)$ we know that $\Delta \cap \mathrm{lk}(\partial\Gamma') \subseteq \Sigma$.

The map $\widetilde{\iota}_{\mathrm{lk}_{\Delta\cup\Theta}(\partial\Gamma'),\Delta\cup\Theta}$ is onto by the Product Axiom, and so there exists a standard copy of $\widetilde{X}_{\mathrm{lk}_{\Delta\cup\Theta}(\partial\Gamma')}$ which intersects $\widetilde{Z}_{\Sigma_1}$. But we have $\mathrm{lk}_{\Delta\cup\Theta}(\partial\Gamma') \subseteq \Sigma$, and thus there exists a copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_{\Delta\cup\Theta}$ (since $\widetilde{X}_{\Delta\cup\Theta}$ is a product) which contains the given copy of $\widetilde{X}_{\mathrm{lk}_{\Delta\cup\Theta}(\partial\Gamma')}$, and thus is as required by the Intersection Axiom.

We have finished extending the copies for all $\Sigma \in \mathbb{P}'$. For all $\Sigma \in \mathbb{P}$ with $\Sigma = \Sigma_1$ we do not even need to extend.

It is still possible that there exists $\Sigma \in \mathbb{P} \smallsetminus \mathbb{P}'$ such that $\Sigma \not\subseteq \Gamma'$. Such a Σ must satisfy $\Sigma \subseteq \overline{\Theta}$ (by $(*)$) and $\Sigma \cap \Theta \notin \{\emptyset, \Theta\}$. But then for all $\Sigma' \in \mathbb{P}$ we have $\Theta \subseteq \Sigma'$ or $\Sigma' \subseteq \overline{\Theta}$, as otherwise $\Sigma \cup \Sigma'$ would violate $(*)$. So in this situation $\mathbb{P} \smallsetminus \mathbb{P}' \subseteq \mathrm{L}^\varphi_{\Delta\cup\Theta}$.

Consider the subsystem $\mathbb{P}'_{\Delta\cup\Theta}$ of $\mathbb{P}_{\Delta\cup\Theta}$. It is closed under taking unions, since $\mathbb{P}'$ is, and for each $\Sigma \in \mathbb{P}$ we have $\Sigma_2 \cup \bigcap \mathbb{P}'_{\Delta\cup\Theta} \in \mathbb{P}'_{\Delta\cup\Theta}$. Now our standard copies of $\widetilde{X}_\Sigma$ for $\Sigma \in \mathbb{P}'$ give us (by the Intersection Axiom) standard copies of $\widetilde{X}_{\Sigma_2}$ in $\widetilde{X}_{\Delta\cup\Theta}$ which intersect in a standard copy of $\widetilde{X}_{(\bigcap \mathbb{P}')_2}$, and hence non-trivially. Lemma 5.5 gives us a collection of standard copies of $\widetilde{X}_{\Sigma_2}$ in $\widetilde{X}_{\Delta\cup\Theta}$ for all $\Sigma \in \mathbb{P}$, which contains the previously discussed collection, and such that all of these standard copies intersect non-trivially. For each $\Sigma \in \mathbb{P} \smallsetminus \mathbb{P}'$ we have $\Sigma = \Sigma_2$, and the standard copy of $\widetilde{X}_\Sigma$ in $\widetilde{X}_{\Delta\cup\Theta}$ becomes a standard copy in $\widetilde{X}_\Gamma$ by construction. □

Bibliography

[AB] Peter Abramenko and Kenneth S. Brown. *Buildings*, volume 248 of *Graduate Texts in Mathematics*. Springer, New York, 2008. Theory and applications.

[Ago] Ian Agol. The virtual Haken conjecture. *Doc. Math.* **18**(2013), 1045–1087. With an appendix by Agol, Daniel Groves, and Jason Manning.

[Ale] J. W. Alexander. Topological invariants of knots and links. *Trans. Amer. Math. Soc.* **30**(1928), 275–306.

[AB] Emina Alibegović and Mladen Bestvina. Limit groups are CAT(0). *J. London Math. Soc. (2)* **74**(2006), 259–272.

[AS] Michael Aschbacher and Gary M. Seitz. Involutions in Chevalley groups over fields of even order. *Nagoya Math. J.* **63**(1976), 1–91.

[Bar] Laurent Bartholdi. Amenability of groups is characterized by Myhill's Theorem. *J. Eur. Math. Soc.* (To appear). With an appendix by Dawid Kielak.

[BKP] Barbara Baumeister, Dawid Kielak, and Emilio Pierro. On the smasmall non-abelian quotient of $\mathrm{Aut}(F_n)$. *arXiv:1705.02885*.

[Bel] Robert W. Bell. Three-dimensional FC Artin groups are CAT(0). *Geom. Dedicata* **113**(2005), 21–53.

[BGP] A. J. Berrick, V. Gebhardt, and L. Paris. Finite index subgroups of mapping class groups. *Proc. Lond. Math. Soc. (3)* **108**(2014), 575–599.

[BFH] M. Bestvina, M. Feighn, and M. Handel. The Tits alternative for $\mathrm{Out}(F_n)$. II. A Kolchin type theorem. *Ann. of Math. (2)* **151**(2000), 517–623.

[BB] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.* **129**(1997), 445–470.

[BNS] Robert Bieri, Walter D. Neumann, and Ralph Strebel. A geometric invariant of discrete groups. *Invent. Math.* **90**(1987), 451–477.

[BP] O.V. Bogopol'skiĭ and D.V. Puga. On embeddings of $\mathrm{Out}(F_n)$, the outer automorphism group of the free group of rank n, into $\mathrm{Out}(F_m)$ for $m > n$. *Algebra and Logic* **41**(2002), 69–73.

[Bow] B. H. Bowditch. Notes on locally CAT(1) spaces. In *Geometric group theory (Columbus, OH, 1992)*, volume 3 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 1–48. de Gruyter, Berlin, 1995.

[BC] Noel Brady and John Crisp. Two-dimensional Artin groups with CAT(0) dimension three. In *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000)*, volume 94, pages 185–214, 2002.

[Bra] Thomas Brady. Artin groups of finite type with three generators. *Michigan Math. J.* **47**(2000), 313–324.

[BM1] Thomas Brady and Jonathan P. McCammond. Three-generator Artin groups of large type are biautomatic. *J. Pure Appl. Algebra* **151**(2000), 1–9.

[BM2] Tom Brady and Jon McCammond. Braids, posets and orthoschemes. *Algebr. Geom. Topol.* **10**(2010), 2277–2314.

[BHRD] J. N. Bray, D. F. Holt, and C. M. Roney-Dougal. *The maximal subgroups of the low-dimensional finite classical groups*, volume 407 of *London Math. Soc. Lecture Note Ser.* Cambridge University Press, Cambridge, 2013.

[BH] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BV1] Martin R. Bridson and Karen Vogtmann. Homomorphisms from automorphism groups of free groups. *Bull. London Math. Soc.* **35**(2003), 785–792.

[BV2] Martin R. Bridson and Karen Vogtmann. Actions of automorphism groups of free groups on homology spheres and acyclic manifolds. *Comment. Math. Helv.* **86**(2011), 73–90.

[BV3] Martin R. Bridson and Karen Vogtmann. Abelian covers of graphs and maps between outer automorphism groups of free groups. *Mathematische Annalen* **353**(2012), 1069–1102.

[Bro1] Kenneth S. Brown. *Buildings*. Springer-Verlag, New York, 1989.

[Bro2] Samuel Brown. A gluing theorem for negatively curved complexes. *J. Lond. Math. Soc. (2)* **93**(2016), 741–762.

[BKM] Inna Bumagin, Olga Kharlampovich, and Alexei Miasnikov. The isomorphism problem for finitely generated fully residually free groups. *J. Pure Appl. Algebra* **208**(2007), 961–977.

[CL] Christopher H. Cashen and Gilbert Levitt. Mapping tori of free group automorphisms, and the Bieri–Neumann–Strebel invariant of graphs of groups. *J. Group Theory* **19**(2016), 191–216.

[CFF] J. C. Cha, S. Friedl, and F. Funke. The Grothendieck group of polytopes and norms. *arXiv:1512.06699* (December 2015).

[CCHO] Jérémie Chalopin, Victor Chepoi, Hiroshi Hirai, and Damian Osajda. Weakly modular graphs and nonpositive curvature. *arXiv:1409.3892.*

[Cha] Ruth Charney. An introduction to right-angled Artin groups. *Geom. Dedicata* **125**(2007), 141–158.

[CCV] Ruth Charney, John Crisp, and Karen Vogtmann. Automorphisms of 2-dimensional right-angled Artin groups. *Geom. Topol.* **11**(2007), 2227–2264.

[CD1] Ruth Charney and Michael Davis. Singular metrics of nonpositive curvature on branched covers of Riemannian manifolds. *Amer. J. Math.* **115**(1993), 929–1009.

[CD2] Ruth Charney and Michael W. Davis. Finite $K(\pi, 1)$s for Artin groups. In *Prospects in topology (Princeton, NJ, 1994)*, volume 138 of *Ann. of Math. Stud.*, pages 110–124. Princeton Univ. Press, Princeton, NJ, 1995.

[CSV] Ruth Charney, Nathaniel Stambaugh, and Karen Vogtmann. Outer space for untwisted automorphisms of right-angled Artin groups. *arXiv:1212.4791.*

[CS] Adam Clay and Lawrence H. Smith. Corrigendum to: "On ordering free groups" [J. Symbolic Comput. 40 (2005) 1285–1290]. *Journal of Symbolic Computation* **44**(2009), 1529 – 1532.

[Cla] M. Clay. ℓ^2-torsion of free-by-cyclic groups. *arXiv:1509.09258* (September 2015).

[CCN$^+$] J. H. Conway, R. T. Curtis, S. P. Norton, R. A. Parker, and R. A. Wilson. *An Atlas of finite groups.* Oxford University Press, Eynsham, 1985.

[Coo] B. N. Cooperstein. Maximal subgroups of $G_2(2^n)$. *J. Algebra* **70**(1981), 23–36.

[Cul] Marc Culler. Finite groups of outer automorphisms of a free group. In *Contributions to group theory*, volume 33 of *Contemp. Math.*, pages 197–207. Amer. Math. Soc., Providence, RI, 1984.

[CV] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.* **84**(1986), 91–119.

[Die] Jean Dieudonné. Les déterminants sur un corps non commutatif. *Bull. Soc. Math. France* **71**(1943), 27–45.

[DM] John D. Dixon and Brian Mortimer. *Permutation groups*, volume 163 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1996.

[DFL] J. Dubois, S. Friedl, and W. Lück. The L^2-Alexander torsion of 3-manifolds. *arXiv:1410.6918* (October 2014).

[Dun1] M. J. Dunwoody. Cutting up graphs. *Combinatorica* **2**(1982), 15–23.

[Dun2] M. J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.* **81**(1985), 449–457.

[Fox] Ralph H. Fox. Free differential calculus. I. Derivation in the free group ring. *Ann. of Math. (2)* **57**(1953), 547–560.

[FH] S. Friedl and S. Harvey. Non-commutative Multivariable Reidemeister Torsion and the Thurston Norm. *Alg. Geom. Top.* **7**(2007), 755–777.

[FL1] Stefan Friedl and Wolfgang Lück. L^2-Euler characteristic and the Thurston norm. *arXiv:1609.07805.*

[FL2] Stefan Friedl and Wolfgang Lück. The L^2-torsion function and the Thurston norm of 3-manifolds. *arXiv:1510.00264.*

[FL3] Stefan Friedl and Wolfgang Lück. Universal L^2-torsion, polytopes and applications to 3-manifolds. *arXiv:1609.07809.*

[FT] Stefan Friedl and Stephan Tillmann. Two-generator one-relator groups and marked polytopes. *arXiv:1501.03489.*

[Fun] F. Funke. The integral polytope group. *arXiv:1605.01217* (2016).

[FK] Florian Funke and Dawid Kielak. Alexander and Thurston norms, and the Bieri–Neumann–Strebel invariants for free-by-cyclic groups. *arXiv:1605.09067.*

[GMSW] Ross Geoghegan, Michael L. Mihalik, Mark Sapir, and Daniel T. Wise. Ascending HNN extensions of finitely generated free groups are Hopfian. *Bull. London Math. Soc.* **33**(2001), 292–298.

[Ger] S. M. Gersten. A presentation for the special automorphism group of a free group. *J. Pure Appl. Algebra* **33**(1984), 269–279.

[Gil] Robert Gilman. Finite quotients of the automorphism group of a free group. *Canad. J. Math.* **29**(1977), 541–551.

[GLS] Daniel Gorenstein, Richard Lyons, and Ronald Solomon. *The classification of the finite simple groups. Number 3. Part I. Chapter A*, volume 40 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1998. Almost simple K-groups.

[Grä] George Grätzer. *Lattice theory: foundation*. Birkhäuser/Springer Basel AG, Basel, 2011.

[Gro1] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Gro2] Edna K. Grossman. On the residual finiteness of certain mapping class groups. *J. London Math. Soc. (2)* **9**(1974/75), 160–164.

[GL] F. Grunewald and A. Lubotzky. Linear Representations of the Automorphism Group of a Free Group. *Geometric and Functional Analysis* **18**(June 2006), 1564–1608.

[Gru] I. Gruschko. Über die Basen eines freien Produktes von Gruppen. *Rec. Math. [Mat. Sbornik] N.S.* **8 (50)**(1940), 169–182.

[GL] Vincent Guirardel and Gilbert Levitt. The outer space of a free product. *Proc. Lond. Math. Soc. (3)* **94**(2007), 695–714.

[HKS] Thomas Haettel, Dawid Kielak, and Petra Schwer. The 6-strand braid group is CAT(0). *Geom. Dedicata* **182**(2016), 263–286.

[Hal] Marshall Hall, Jr. Subgroups of finite index in free groups. *Canadian J. Math.* **1**(1949), 187–190.

[HM] Michael Handel and Lee Mosher. Relative free splitting and free factor complexes I: Hyperbolicity. *arXiv:1407.3508*.

[Har1] S. L. Harvey. Monotonicity of degrees of generalized Alexander polynomials of groups and 3-manifolds. *Math. Proc. Cambridge Philos. Soc.* **140**(2006), 431 – 450.

[Har2] Shelly L. Harvey. Higher-order polynomial invariants of 3-manifolds giving lower bounds for the Thurston norm. *Topology* **44**(2005), 895–945.

[HK1] Sebastian Hensel and Dawid Kielak. Nielsen realisation for untwisted automorphisms of right-angled Artin groups. *arXiv:1410.1618v2*.

[HK2] Sebastian Hensel and Dawid Kielak. Nielsen realisation by gluing: Limit groups and free products. *Michigan Math. J.* (To appear).

[HOP] Sebastian Hensel, Damian Osajda, and Piotr Przytycki. Realisation and dismantlability. *Geom. Topol.* **18**(2014), 2079–2126.

[Hig] Graham Higman. The units of group-rings. *Proc. London Math. Soc. (2)* **46**(1940), 231–248.

[Hor] Camille Horbez. The boundary of the outer space of a free product. *arXiv:1408.0543*.

[Ima] Shiro Imamura. On embeddings between outer automorphism groups of right-angled Artin groups. *arXiv:1702.04809*.

[Kap] Irving Kaplansky. *Problems in the theory of rings. Report of a conference on linear algebras, June, 1956, pp. 1-3*. National Academy of Sciences-National Research Council, Washington, Publ. 502, 1957.

[KPS] A. Karrass, A. Pietrowski, and D. Solitar. Finite and infinite cyclic extensions of free groups. *J. Austral. Math. Soc.* **16**(1973), 458–466. Collection of articles dedicated to the memory of Hanna Neumann, IV.

[Ker1] Steven P. Kerckhoff. The Nielsen realization problem. *Bull. Amer. Math. Soc. (N.S.)* **2**(1980), 452–454.

[Ker2] Steven P. Kerckhoff. The Nielsen realization problem. *Ann. of Math. (2)* **117**(1983), 235–265.

[Khr1] D. G. Khramtsov. Finite groups of automorphisms of free groups. *Mat. Zametki* **38**(1985), 386–392, 476.

[Khr2] D. G. Khramtsov. Outer automorphisms of free groups. In *Group-theoretic investigations (Russian)*, pages 95–127. Akad. Nauk SSSR Ural. Otdel., Sverdlovsk, 1990.

[Kie1] Dawid Kielak. Outer automorphism groups of free groups: linear and free representations. *J. London Math. Soc.* **87**(2013), 917–942.

[Kie2] Dawid Kielak. Groups with infinitely many ends are not fraction groups. *Groups Geom. Dyn.* **9**(2015), 317–323.

[Kie3] Dawid Kielak. Low-dimensional free and linear representations of $\mathrm{Out}(F_3)$. *J. Group Theory* **18**(2015), 913–949.

[Kie4] Dawid Kielak. Outer actions of $\mathrm{Out}(F_n)$ on small RAAGs. *arXiv:1607.07031*.

[KP] Dawid Kielak and Emilio Pierro. On the smallest non-trivial quotients of mapping class groups. *arXiv:1705.10223*.

[Kle1] P. B. Kleidman. The maximal subgroups of the Chevalley groups $G_2(q)$ with q odd, the Ree groups ${}^2G_2(q)$, and their automorphism groups. *J. Algebra* **117**(1988), 30–71.

[Kle2] P. B. Kleidman. The maximal subgroups of the Steinberg triality groups ${}^3D_4(q)$ and of their automorphism groups. *J. Algebra* **115**(1988), 182–199.

[KL] Peter Kleidman and Martin Liebeck. *The subgroup structure of the finite classical groups*, volume 129 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1990.

[Koc] D. H. Kochloukova. Some Novikov rings that are von Neumann finite. *Comment. Math. Helv.* **81**(2006), 931–943.

[Krö] Bernhard Krön. Cutting up graphs revisited—a short proof of Stallings' structure theorem. *Groups Complex. Cryptol.* **2**(2010), 213–221.

[LS] Vicente Landazuri and Gary M. Seitz. On the minimal degrees of projective representations of the finite Chevalley groups. *J. Algebra* **32**(1974), 418–443.

[Lau] Michael R. Laurence. A generating set for the automorphism group of a graph group. *J. London Math. Soc. (2)* **52**(1995), 318–334.

[LS] Martin W. Liebeck and Gary M. Seitz. On finite subgroups of exceptional algebraic groups. *J. Reine Angew. Math.* **515**(1999), 25–72.

[LL] P. Linnell and W. Lück. Localization, Whitehead groups, and the Atiyah Conjecture. *arXiv:1602.06906* (2016).

[LLS] Peter A. Linnell, Wolfgang Lück, and Thomas Schick. The Ore condition, affiliated operators, and the lamplighter group. In *High-dimensional manifold topology*, pages 315–321. World Sci. Publ., River Edge, NJ, 2003.

[LS] W. Lück and T. Schick. L2-torsion of hyperbolic manifolds of finite volume. *Geom. Func. Anal.* **9**(1999), 518–567.

[Lüc1] Wolfgang Lück. *L^2-invariants: theory and applications to geometry and K-theory*, volume 44 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]*. Springer-Verlag, Berlin, 2002.

[Lüc2] Wolfgang Lück. Twisting L^2-invariants with finite-dimensional representations. *arXiv:1510.00057*.

[Mal1] A. I. Mal'cev. On the embedding of group algebras in division algebras. *Doklady Akad. Nauk SSSR (N.S.)* **60**(1948), 1499–1501.

[Mal2] G. Malle. The maximal subgroups of ${}^2F_4(q^2)$. *J. Algebra* **139**(1991), 52–69.

[McC] Stephen H. McCleary. Free lattice-ordered groups represented as o-2 transitive l-permutation groups. *Trans. Amer. Math. Soc.* **290**(1985), 69–79.

[MR] J.C. McConnell and J.C. Robson. *Noncommutative Noetherian Rings*. Wiley series in pure and applied mathematics. Wiley, 1987.

[McM] Curtis T. McMullen. The Alexander polynomial of a 3-manifold and the Thurston norm on cohomology. *Ann. Sci. École Norm. Sup. (4)* **35**(2002), 153–171.

[MZ] Mattia Mecchia and Bruno Zimmermann. On minimal finite factor groups of outer automorphism groups of free groups. *Atti Semin. Mat. Fis. Univ. Modena Reggio Emilia* **57**(2010), 115–120 (2011).

[Men] Jens L. Mennicke. Finite factor groups of the unimodular group. *Ann. of Math. (2)* **81**(1965), 31–37.

[Min] Igor Mineyev. Submultiplicativity and the Hanna Neumann conjecture. *Ann. Math. (2)* **175**(2012), 393–414.

[Nav] Andrés Navas. On the dynamics of (left) orderable groups. *Annales de l'institut Fourier* **60**(2010), 1685–1740.

[Neu] B. H. Neumann. On ordered division rings. *Trans. Amer. Math. Soc.* **66**(1949), 202–252.

[PR] A. Potapchik and A. Rapinchuk. Low-dimensional linear representations of Aut F_n, $n \geq 3$. *Trans. Amer. Math. Soc.* **352**(2000), 1437–1451.

[Riv] Cristóbal Rivas. Left-orderings on free products of groups. *Journal of Algebra* **350**(2012), 318 – 329.

[Ser1] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[Ser2] Herman Servatius. Automorphisms of graph groups. *J. Algebra* **126**(1989), 34–60.

[Sik] J.-C. Sikorav. Homologie de Novikov associée à une classe de cohomologie réelle de degré un. *Thèse Orsay* (1987).

[Sil] R. J. Silvester. *Introduction to algebraic K-theory*. Chapman & Hall, London, 1981.

[Sta1] John R. Stallings. On Torsion-Free Groups with Infinitely Many Ends. *Annals of Mathematics* **Second Series, 88**(September 1968), 312–334.

[Sta2] John R. Stallings. *Group theory and three-dimensional manifolds.* Yale University Press, New Haven, 1971.

[Sta3] John R. Stallings. Topology of finite graphs. *Invent. Math.* **71**(1983), 551–565.

[Str] Ralph Strebel. *Notes on the Sigma invariants.* arXiv:1204.0214.

[Tam] Dov Tamari. A refined classification of semi-groups leading to generalized polynomial rings with a generalized degree concept. In *Proceedings of the International Congress of Mathematicians, 1954, Vol. 1, Amsterdam*, pages 439–440. North-Holland, 1957.

[Thu] William P. Thurston. A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.* **59**(1986), i–vi and 99–130.

[Toi] Emmanuel Toinet. Conjugacy p-separability of right-angled Artin groups and applications. *Groups Geom. Dyn.* **7**(2013), 751–790.

[TW] Victor Turchin and Thomas Willwacher. Hochschild–Pirashvili homology on suspensions and representations of $\mathrm{Out}(F_n)$. *arXiv:1507.08483.*

[Var] Olga Varghese. Low dimensional linear representations of $\mathrm{SAut}(F_n)$. *arXiv:1509.00187.*

[Wad] Richard D. Wade. Johnson homomorphisms and actions of higher-rank lattices on right-angled Artin groups. *J. Lond. Math. Soc. (2)* **88**(2013), 860–882.

[Wag] Ascher Wagner. The Faithful Linear Representation of Least Degree of Sn and An over a Field of Characteristic 2. *Mathematische Zeitschrift* **151**(1976), 127–138.

[Wal] F. Waldhausen. Algebraic K-Theory of Generalized Free Products, Part 2. *Ann. of Math.* **108**(1978), 205–256.

[Weg] C. Wegner. The Farrell-Jones Conjecture for virtually solvable groups. *arXiv:1308.2432* (2013).

[Wei] S.H. Weintraub. *Representation theory of finite groups: algebra and arithmetic.* Graduate studies in mathematics. American Mathematical Society, 2003.

[Wil] R. A. Wilson. *The finite simple groups*, volume 251 of *Graduate Texts in Mathematics.* Springer-Verlag London, Ltd., London, 2009.

[Zim1] Bruno Zimmermann. Über Homöomorphismen n-dimensionaler Henkelkörper und endliche Erweiterungen von Schottky-Gruppen. *Comment. Math. Helv.* **56**(1981), 474–486.

[Zim2] Bruno P. Zimmermann. On minimal finite quotients of mapping class groups. *Rocky Mountain J. Math.* **42**(2012), 1411–1420.