Groups of type FP via graphical small cancellation

Thomas Brown   Ian J. Leary

April 9, 2020

Abstract

We construct an uncountable family of groups of type FP. In contrast to every previous construction of non-finitely presented groups of type FP we do not use Morse theory on cubical complexes; instead we use Gromov's graphical small cancellation theory.

1   Introduction

The first examples of non-finitely presented groups of type FP were constructed in the 1990's by Bestvina and Brady, using Morse theory on CAT(0) cubical complexes [2]. Brady also used similar techniques to construct finitely presented subgroups of hyperbolic groups that are not themselves hyperbolic because they are not FP_3 [5]. With the benefit of hindsight, examples due to Stallings and Bieri of groups that are FP_n but that are not FP_{n+1} can be viewed as special cases of the Bestvina-Brady construction [25] and [3, pp.37–40]. In [10], Bux and Gonzalez pointed out the close connection between the Bestvina-Brady construction and Brown's criterion for finiteness properties [6]. The computations of finiteness properties that Brown made in [6] using his new criterion can also be rephrased in terms of Morse theory. Since then, Morse theory on polyhedral complexes has been a vital tool in computing the finiteness properties of many families of groups—for some notable recent examples see [11, 9]. The Bestvina-Brady argument has been extended in a number of ways; in particular the second named author has constructed continuously many isomorphism types of groups of type FP [18]. Nevertheless, it is remarkable that until now, every construction of a non-finitely presented group of type FP has relied on the same Morse-theoretic techniques that were used by Bestvina-Brady.

There are of course a number of ways to make 'new groups of type FP from old': notable examples include the Davis trick, which has been used to produce non-finitely presented Poincaré duality groups [12, 13, 18], the method used by Skipper-Witzel-Zaremsky to construct simple groups with given homological finiteness properties [26] and the proof that every countable group embeds in a group of type FP_2 [19]. Nevertheless, these constructions require a pre-existing family of groups of type FP, and so they rely ultimately on the Morse theoretic techniques of Bestvina-Brady.

Here, we give a new construction for non-finitely presented groups of type FP, that relies on Gromov's graphical small cancellation theory [15, sec. 2] instead of the techniques used by Bestvina-Brady. Gromov introduced graphical small cancellation as a method to embed certain families of graphs, in particular an expanding family, as subgraphs of a

Cayley graph. Our main idea is to use graphical small cancellation to construct families of surjective group homomorphisms with acyclic kernels.

The remainder of this introduction consists of a sketch of the construction and the statements of our main results. For the definitions and background results that we assume concerning homological finiteness properties and graphical small cancellation see Section 2.

The Bestvina-Brady construction takes as its input a finite flag simplicial complex L, which should be acyclic but not contractible in order to produce a group BB_L that is FP but not finitely presented. For constructing continuously many generalized Bestvina-Brady groups $G_L(S)$ as in [18], one also requires L to be aspherical. Our construction takes as input a finite polygonal cell complex K with properties as listed below, together with an infinite set Z of non-zero integers.

Definition 1. *A suitable complex K is a finite polygonal cell complex K with the following properties.*

1. *The 1-skeleton K^1 of K is a simplicial graph;*

2. *The attaching map for each polygon P of K is an embedding of the boundary circle ∂P into K^1;*

3. *The girth g of the graph K^1 is at least 13;*

4. *The perimeter l_P of each polygon P satisfies $l_P > 2g$;*

5. *The polygons of K satisfy a $C'(1/6)$ condition: for each pair $P \neq Q$ of polygons of K, each component of the intersection $\partial P \cap \partial Q$ of their boundaries contains strictly fewer than $\min\{l_P/6, l_Q/6\}$ edges;*

6. *K is acyclic.*

The existence of a suitable complex K will be established below. For now we suppose that we are given a 2-complex K as above and an infinite set $Z \subseteq \mathbb{Z} - \{0\}$. For each subset $S \subseteq Z$, we define a group $H(S)$ via a graphical presentation. The generators of $H(S)$ are the directed edges of K, where the two orientations of the same edge are mutually inverse group elements. For each $n \in Z - S$ and each polygon P of K, we take the ordinary relation which is the 'degree n subdivision of ∂P'. For each $n \in S$ we instead take the single graphical relation that can be described as the 'degree n subdivision of the graph K^1'. The degree n subdivision of a graphical relation is defined in Section 2, but see also the example in figure 1. For P a polygon of K, we define $H_P(S)$ to be the subgroup of $H(S)$ generated by the edges of P.

Theorem 2. *For each $S \subseteq Z$ the graphical presentation for $H(S)$ given above satisfies the graphical small cancellation condition $C'(1/6)$. For each $S \subsetneq T \subseteq Z$ the natural bijection between generating sets extends to a surjective group homomorphism $H(S) \to H(T)$, whose kernel $K_{S,T}$ is a non-trivial acyclic group. For each polygon P the intersection $K_{S,T} \cap H_P(S)$ is the trivial group and so the natural map restricts to an isomorphism $H_P(S) \cong H_P(T)$.*
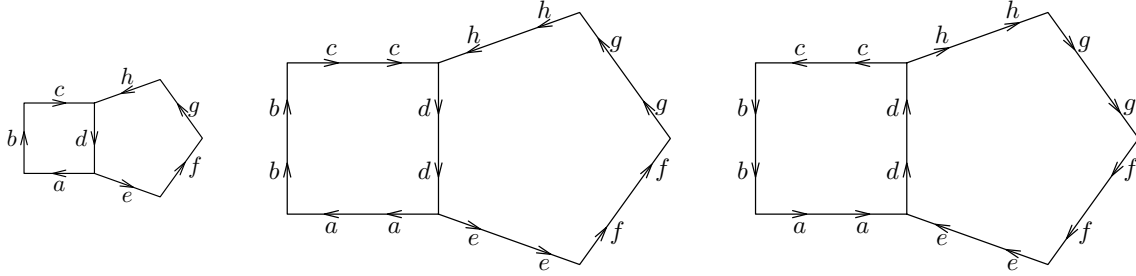
Figure 1: A graphical relation and its degree 2 and -2 subdivisions

None of the groups $H(S)$ are of type FP, however they form the main building block in our construction of such groups.

The isomorphism type of the polygon subgroup H_P depends only on Z and on the perimeter of the polygon P. If P has perimeter $l = l_P$, then H_P has the following presentation:

$$H_P = \langle a_1, \ldots, a_l \ : \ a_1^n a_2^n \cdots a_l^n = 1 \ n \in Z \rangle.$$

The final ingredient that we need is an embedding of H_P into a group G_P, where G_P is of type F. The existence of such an embedding puts some constraints on the set Z. Using Sapir's version of the Higman embedding theorem [24], one can show that there is such an embedding if and only if Z is recursively enumerable. However, for our purposes we only require an example of such a Z, and we give explicit constructions for G_P in the cases $Z = \mathbb{Z} - \{0\}$ and $Z = \{k^n : n \geq 0\}$ for any integer k with $|k| > 1$.

Given such an embedding, we define a group $G(S)$ as the fundamental group of a star-shaped graph of groups. The underlying graph has one central vertex and arms of length one indexed by the polygons of K. The central vertex group is $H(S)$, with $H_P = H_P(S)$ on the edge indexed by P and G_P on the outer vertex indexed by P.

Proposition 3. *The group $G(\emptyset)$ is of type F. For each $S \subsetneq T \subseteq Z$ there is a surjective group homomorphism $G(S) \to G(T)$ with non-trivial acyclic kernel.*

Corollary 4. *There are continuously many isomorphism types of the groups $G(S)$ for varying S. Each group $G(S)$ is of type FP. $G(S)$ is finitely presented if and only if S is finite, and $G(S)$ embeds as a subgroup of a finitely presented group if and only if S is recursively enumerable. Provided that each G_P has geometric dimension two, $G(\emptyset)$ also has geometric dimension two and each $G(S)$ has cohomological dimension two.*

Corollary 4 should be compared with theorems 1.2 and 1.3 from [18], in the special case when the flag complex L used there is acyclic. However, we claim that our argument is simpler than those in [2, 18], even if one only wants to establish the existence of some non-finitely presented group of type FP.

Although the proofs are rather different, there is an overlap between the groups $G(S)$ obtained from the new construction and the generalized Bestvina-Brady groups of [18]. In the case when $Z = \mathbb{Z} - \{0\}$, each H_P is isomorphic to a Bestvina-Brady group and we may take for G_P the corresponding right-angled Artin group. For these choices, the group $G(S)$ is naturally isomorphic to the generalized Bestvina-Brady group $G_L(S \cup \{0\})$ of [18], where L is the flag triangulation of K obtained by viewing each polygon as a

cone on its boundary. In particular with these choices $G(\mathbb{Z} - \{0\})$ is isomorphic to the Bestvina-Brady group BB_L.

In the next section, we give some background material concerning finiteness properties and graphical small cancellation. In Section 3 we use graphical small cancellation methods to prove Theorem 2. In Section 4 we use standard methods from graphs of groups to prove Proposition 3 and deduce Corollary 4. In Section 5 we discuss embeddings of the polygon groups H_P into 2-dimensional groups G_P of type F, and then in Section 6 we establish the existence of a 2-complex K with the required properties. This ordering of the material reflects the history of our work: in particular we had a rough version of the main theorem long before we had established the existence of a suitable complex K. Finally, Section 7 discusses some questions that remain open.

This work was done while the first named author was working on his PhD under the supervision of the second named author. Further properties of the groups $G(S)$ and other methods for constructing 2-complexes K with the required properties will appear in the first named author's PhD thesis [8]. The main motivation for this work was the observation that the relations in the presentations for the groups $G_L(S)$ given in [18, Defn. 1.1] consist of a large family of 'long' relations together with a finite number of 'short' relations. Another motivation (which predates the work in [18]) was a conversation between the second named author and Martin Bridson, in which Martin Bridson pointed out that there ought to be other constructions of non-finitely presented groups of type FP apart from that of Bestvina-Brady. The authors gratefully acknowledge this inspiration. The authors also gratefully acknowledge helpful comments on this work by Tim Riley.

2   Background

An *Eilenberg-Mac Lane space* for a group G is a connected CW-complex whose fundamental group is isomorphic to G and whose universal cover is contractible. Any two such spaces are based homotopy equivalent.

A group G is *type F* if G admits an Eilenberg-Mac Lane space with finitely many cells. A space is *acyclic* if it has the same homology as a point. A group G is *type FH* if there is a free G-CW-complex that is acyclic and has only finitely many orbits of cells. A group G is *type FL* if the trivial module $\mathbb{Z}$ for its group algebra $\mathbb{Z}G$ admits a finite resolution by finitely generated free $\mathbb{Z}G$-modules. Finally, a group G is *type FP* if $\mathbb{Z}$ admits a finite resolution by finitely generated projective $\mathbb{Z}G$-modules. From the definitions it is easy to see that
$$F \implies FH \implies FL \implies FP,$$
and that any finitely presented group of type FL is of type F. For further details concerning these properties see [3, 7] and [2, sec. 1]. We require one general result concerning finiteness properties. A group is said to be *acyclic* if its Eilenberg-Mac Lane space is acyclic.

Proposition 5. *Suppose that G is type F and that N is an acyclic normal subgroup of G. The group G/N is type FH, and the cohomological dimension of G/N is bounded above by the geometric dimension of G.*

Proof. Let X be the universal covering space of an Eilenberg-Mac Lane space for G, and consider the quotient X/N. This is an Eilenberg-Mac Lane space for N, equipped with a

free cellular action of G/N. Since N is acyclic, X/N is acyclic. The G/N-orbits of cells in X/N correspond to the G-orbits of cells in X, and so if G is type F then G/N is type FH. If X has dimension n then so does X/N, and the dimension of X/N is an upper bound for the cohomological dimension of G/N. □

A *graph of groups* indexed by a graph Γ consists of groups G_v and G_e for each vertex v and edge e of Γ, together with two injective group homomorphisms $G_e \to G_v$ from the edge group G_e to the vertex groups corresponding to the ends of e. A *graph of based spaces* is defined similarly. If each vertex and edge space in a graph of spaces is an Eilenberg-Mac Lane space and the induced maps on fundamental groups are all injective, then the homotopy colimit (in the category of unbased spaces) of the graph of spaces is also an Eilenberg-Mac Lane space, whose fundamental group is by definition the *fundamental group of the graph of groups*.

Next we describe *graphical small cancellation theory* as in [15, sec. 2] and [23]. The theory subsumes the classical small cancellation theory [21, ch. V], which is the case in which the graph Γ considered below is a disjoint union of cycles. Note that we consider only a simplified version, which corresponds in the classical case to excluding the possibility that a relator is a proper power.

A *labelling* of a graph Γ consists of a set L of labels, together with a fixed-point free involution $\tau : L \to L$, and a labelling function ϕ from the set of directed edges of Γ to L so that the label on the opposite edge to e is $\tau \circ \phi(e)$. A labelling is said to be *reduced* if the graph Γ contains no vertices of valence 0 or 1, and for every vertex v the labelling function ϕ is injective on the outward-pointing neighbours of v. If a labelling is reduced then any word in the elements of L will describe at most one edge-path starting at any vertex v of Γ. A reduced word in L is a finite sequence of elements of L that contains no subword $(l, \tau(l))$. A *piece* is a reduced word in L that defines an edge path starting at more than one vertex of Γ. A piece is said to belong to each component of Γ that contains at least one of these vertices. The length of a piece is the number of letters in the word, or equivalently the length of the corresponding edge paths.

Given a reduced labelled graph Γ, the corresponding *graphical presentation complex* is a 2-dimensional complex obtained as follows. Start with a rose R_L (i.e., a 1-dimensional CW-complex with one vertex) with petals in bijective correspondence with the τ-orbits in L, and fix a bijection from the oriented edges of the rose to L. The labelling of the edges of Γ describes a cellular map $f : \Gamma \to R_L$ which is unique up to homotopy. The graphical presentation complex W is obtained by attaching to R_L the cone on every component of Γ using the map f. The group presented by the graphical presentation is the fundamental group $\pi_1(W)$ of W, and is a quotient of the fundamental group of R_L. The graphical Cayley complex X is the universal cover of W, and the inverse image of R_L in X is the *graphical 1-skeleton X^1* of X. The graphical 1-skeleton X^1 is the Cayley graph (in the usual sense) of $\pi_1(W)$ with respect to the generators given by the edges of R_L. Each lift to X of a copy of the cone on a component of Γ is called a *graphical 2-cell*.

The *girth* of a graph Γ is the minimal length of any cycle in Γ. A graphical presentation satisfies the *small cancellation condition $C'(1/6)$* if the length of any piece is strictly less than $1/6$ of the girth of each component to which the piece belongs. We are now ready to state the main theorem of graphical small cancellation theory.

Theorem 6. *If the graphical presentation Γ, L, τ, ϕ satisfies $C'(1/6)$ then the associated*

graphical Cayley complex X is contractible, and the attaching map for each graphical 2-cell is an injection from a component of Γ to the graphical 1-skeleton of X. Moreover, the Cayley graph X^1 has 'Dehn's property': any closed loop in X^1 contains strictly more than half of some cycle of Γ as a subpath.

We close this section by describing some labellings that we will use.

Any graph Γ admits a *tautological labelling*, in which the set L is just the set of directed edges of Γ, and the function ϕ is the identity map.

Given a labelling of a graph Γ and n a non-zero integer, the *degree n subdivision* of Γ is the labelled graph obtained by subdividing each edge of Γ into $|n|$ parts. If $n > 0$, then the label attached to each of the n new directed edges contained the directed edge e is $\phi(e)$, whereas if $n < 0$ the label attached to each new directed edge contained in e is is $\tau \circ \phi(e)$. It may be helpful to imagine that the graph is rescaled by a factor of $|n|$, so that the edges of the degree n subdivision are 'the same length' as the edges of the original graph. See figure 1 for an example of a labelled graph and two of its subdivisions.

3   Using graphical small cancellation

In this section we prove Theorem 2. First we need to establish that graphical small cancellation can be applied.

Proposition 7. *The given graphical presentation for $H(S)$ satisfies the small cancellation condition $C'(1/6)$.*

Proof. This is a simple check. The shortest cycle in the degree m subdivision of K^1 is length $g|m| \geq 13|m|$, and the shortest cycle in the degree m subdivision of the polygon boundary ∂P is length $l_P|m| > 26|m|$. Now suppose that $m \neq n$ are non-zero integers with $|m| < |n|$. If $mn > 0$, then the longest pieces contained in both a degree m subdivision and a degree n subdivision are of the form $a^m b^m$, where either (a, b) or (b, a) is a pair of consecutive edges in K^1. If on the other hand $mn < 0$, then the longest pieces are of the form a^m.

Between degree m subdivision of either K^1 or a polygon boundary ∂P and itself, the longest pieces are of the form a^{m-1}, of length $|m| - 1$. The longest pieces between the degree m subdivisions of distinct polygon boundaries ∂P and ∂Q are potentially much longer, of length $|m|$ times the length of a piece of $\partial P \cap \partial Q$. But since the polygons of K satisfy the $C'(1/6)$ condition, it follows that they have length strictly less than both $|m|l_P/6$ and $|m|l_Q/6$. $\square$

For $S \subseteq Z$, let $E(S)$ denote the standard graphical Cayley 2-complex for $H(S)$. For the proof of Theorem 2 we will need to compare, for $S \subseteq T \subseteq Z$, two free $H(T)$-complexes: the standard complex $E(T)$ and the quotient $E(S)/K_{S,T}$. Note that this latter space is an Eilenberg-Mac Lane space for $K_{S,T}$ with a free cellular action of $H(T)$. These two free $H(T)$-spaces have the same graphical 1-skeleta, and they share many of the same graphical 2-cells. The difference between them is that for each $n \in T - S$, $E(S)/K_{S,T}$ contains a free orbit of polygonal 2-cells attached along the degree n subdivision of ∂P for each polygon P of K, whereas $E(T)$ contains a free orbit of graphical 2-cells attached

along the degree n subdivision of the graph K^1 itself. There is a 3-dimensional $H(T)$-complex F that contains both $E(S)/K_{S,T}$ and $E(T)$ as subcomplexes. Inside $E(S)/K_{S,T}$ there is, for each $n \in T - S$, a free $H(T)$-orbit of subspaces homeomorphic to K, in which each polygon P of K is attached to the degree n subdivision of ∂P. To make F from $E(S)/K_{S,T}$, attach a free orbit of cones on K to these copies of K. The inclusion of K^1 into K induces an embedding of the cone $C(K^1)$ on K^1 into the cone $C(K)$, and hence an $H(T)$-equivariant inclusion $E(T) \to F$.

Proposition 8. *With notation as above, the inclusion $E(S)/K_{S,T} \to F$ is an $H(T)$-equivariant homology isomorphism and the inclusion $E(T) \to F$ is an $H(T)$-equivariant deformation retraction.*

Proof. The first claim follows because attaching a cone to an acyclic subspace does not change homology. For the second claim, note that for a polygon P, the cone $C(\partial P)$ on ∂P is a deformation retraction of the cone $C(P)$. Putting these deformation retractions together, it follows that the cone $C(K^1)$ on K^1 is a deformation retraction of the cone $C(K)$ on K. Applying this retraction simultaneously over all such cones appearing in F, it follows that $E(T)$ is an $H(T)$-equivariant deformation retraction of F as claimed. $\square$

Proof. (of Theorem 2) Since the generating sets of $H(S)$ and $H(T)$ are identified with each other, the homomorphism $H(S) \to H(T)$ is surjective. To see that its kernel is non-trivial, let $(a_1, \ldots, a_g)$ be a directed loop in the graph K^1 of length equal to the girth of K^1, let n be an element of $T - S$, and consider the element $h := a_1^n a_2^n \cdots a_g^n$ of $H(S)$. This element is contained in $K_{S,T}$ and is not the identity element by the Dehn property, since it can only contain less than $1/6$ of any relator of $H(S)$ coming from degree $m \neq n$, and it is too short to contain more than half of the degree n subdivision of any ∂P.

The fact that $H_P(S) \cap K_{S,T}$ is trivial follows from the Dehn property, because any reduced cycle in K^1 that consists solely of edges of P must be the nth subdivision of the standard cycle ∂P for some n. Hence no non-trivial element of $H_P(S)$ can map to the identity in $H(T)$.

It remains to show that the kernel $K_{S,T}$ is acyclic. By Proposition 8, $E(S)/K_{S,T}$ is an Eilenberg-Mac Lane space for $K_{S,T}$ and has the same homology as F, which is contractible because it has the same (equivariant) homotopy type as $E(T)$. $\square$

4   Graphs of groups

In this section we prove Proposition 3 and Corollary 4.

Proposition 9. *There is a finite Eilenberg-Mac Lane space for the group $G(\emptyset)$ as in Proposition 3. If each G_P has a finite 2-dimensional Eilenberg-Mac Lane space then so does $G(\emptyset)$.*

Proof. Let X be the presentation 2-complex for $H(\emptyset)$, which is built from standard cells, and has finite 1-skeleton. For each polygon P, let Y_P be the presentation 2-complex for H_P, and let Z_P be a finite Eilenberg-Mac Lane space for G_P. Since the generators and relations for each H_P are a subset of those of $H(\emptyset)$, the inclusion $H_P \to H(\emptyset)$ is induced by an isomorphism between Y_P and a subcomplex of X. Note also that each 2-cell of X is contained in exactly one of these subcomplexes.

If $f_P : Y_P \to Z_P$ is a map that induces the embedding $H_P \to G_p$, then an Eilenberg-Mac Lane space W for $G(\emptyset)$ can be built by taking a copy of the mapping cylinder of $f_P : Y_P \to Z_P$ for each P, and identifying the copy of Y_P with its image in X. If Y_P^1 denotes the finite 1-skeleton of Y_P, then the mapping cylinder of the restriction $f_P|_{Y_P^1}$ is a deformation retract of the mapping cylinder for f_P, and is a finite complex. Since each 2-cell of X belongs to a unique polygon P, these deformation retractions can be combined. This shows that the finite complex W' obtained from X^1 and the mapping cylinders of the maps $f_P|_{Y_P^1}$ by identifying each copy of Y_P^1 with its image in X^1 is a deformation retract of W. If each Z_P is 2-dimensional then so is W'. $\square$

Proposition 10. *For each $S \subseteq T \subseteq Z$, the kernel of the homomorphism $G(S) \to G(T)$ is acyclic.*

Proof. It suffices to show that the kernel of the map $G(S) \to G(T)$ is a countable free product of copies of the acyclic group $K_{S,T}$. To show this we construct an Eilenberg-Mac Lane space for the kernel.

Construct Eilenberg-Mac Lane spaces X_S, X_T for $G(S)$ and $G(T)$ respectively as star-shaped graphs of spaces, with central vertex space an Eilenberg-Mac Lane space for $H(S)$ (resp. $H(T)$), outer vertex spaces Eilenberg-Mac Lane spaces for each G_P and edge spaces Eilenberg-Mac Lane spaces for the groups H_P. A map $f : X_S \to X_T$ realizing the surjection $G(S) \to G(T)$ on fundamental groups can be chosen to preserve the graph of spaces structure. An Eilenberg-Mac Lane space for the kernel can be constructed by pulling back the universal covering of X_T along the map $f : X_S \to X_T$. The universal covering of X_T is a tree of spaces, with each vertex space contractible and each edge space contractible. The pullback along f also has the structure of a tree of spaces. Each vertex space that projects to the central vertex of the star is an Eilenberg-Mac Lane space for $K_{S,T}$ and each other vertex space and each edge space is contractible. Hence the kernel of the map $G(S) \to G(T)$ is a free product of countably many copies of $K_{S,T}$, which is acyclic as claimed. $\square$

Proof. (of Corollary 4) Recall from [18, sec. 15] the set-valued invariant $\mathcal{R}(\mathbf{g}, G) \subseteq \mathbb{Z}$ for a group G and a sequence $\mathbf{g} = (g_1, \ldots, g_l)$ of elements of G, defined by

$$\mathcal{R}(\mathbf{g}, G) = \{n \in \mathbb{Z} \ : \ g_1^n g_2^n \cdots g_l^n = 1\}.$$

In [18, prop. 15.2], three properties of this invariant were established: for a fixed isomorphism type of countable group G, the invariant takes only countably many values as $\mathbf{g}$ varies; if $H \geq G$ then $\mathcal{R}(\mathbf{g}, G) = \mathcal{R}(\mathbf{g}, H)$; if G is finitely presented then $\mathcal{R}(\mathbf{g}, G)$ is recursively enumerable.

Let $a_1, \ldots, a_g$ be a directed loop in K of length equal to the girth of K^1. Viewing this loop as a sequence of elements of $H(S)$, the Dehn property implies that for any $S \subseteq Z$, $a_1^n \cdots a_g^n = 1$ in $H(S)$ if and only if $n \in S$. Hence for any $S \subseteq Z$ one has

$$\mathcal{R}((a_1, \ldots, a_g), G(S)) = \mathcal{R}((a_1, \ldots, a_g), H(S)) = S \cup \{0\}.$$

From this together with the known properties of the invariant $\mathcal{R}$ it follows that there are continuously many isomorphism types of groups $G(S)$, that $G(S)$ is finitely presented if and only if S is finite, and that $G(S)$ can embed in a finitely presented group only when

S is recursively enumerable. For the converse, note that since we know that H_P embeds in a finitely presented group (see the remark at the end of the next section) it follows that Z is recursively enumerable. Now if S and Z are both recursively enumerable $G(S)$ is recursively presented and so by the Higman embedding theorem [16, 21], $G(S)$ does embed in some finitely presented group.

We know already that $G(\emptyset)$ has geometric dimension two and is type F; since the kernel of the map $G(\emptyset) \to G(S)$ is acyclic it follows from Proposition 5 that $G(S)$ is type FH and has cohomological dimension two. $\qquad\square$

5 Embedding polygon subgroups

In this section we construct embeddings of the polygon subgroup H_P, whose isomorphism type depends only on the perimeter l of P and on $Z \subseteq \mathbb{Z} - \{0\}$, into groups G_P of type F. Moreover, each group G_P will have geometric dimension two.

The first case that we deal with is the case $Z = \mathbb{Z} - \{0\}$. Before starting this case we recall that the right-angled Artin group A_L associated to a flag simplicial complex L is the group with generators the vertices of L, subject to the relations that the two vertices incident on each edge commute. The Bestvina-Brady group BB_L is the kernel of the map $A_L \to \mathbb{Z}$ that sends each of the generators to $1 \in \mathbb{Z}$. Provided that L is connected, there is a generating set for BB_L that corresponds to the directed edges of L, where the edge from vertex x to y corresponds to the element xy^{-1} in A_L. A presentation for BB_L in terms of these generators is given in [14]. In the case when $Z = \mathbb{Z} - \{0\}$, the given presentation for H_P is equal to this presentation for $BB_{\partial P}$. Hence we may take for G_P the right-angled Artin group $A_{\partial P}$. However, it may be more helpful to use the natural isomorphism between the right-angled Artin group A_L and the Bestvina-Brady group $BB_{C(L)}$ for the cone on L: if c is the cone vertex, this isomorphism takes the vertex generator x to the generator xc^{-1} corresponding to the edge from x to c.

Proposition 11. *Let K be a suitable 2-complex, let $Z = \mathbb{Z} - \{0\}$ and let L be the flag complex obtained from K by replacing each polygon with the cone on its boundary. With the embedding $H_P \to G_P$ as described above there is an isomorphism, for each $S \subseteq \mathbb{Z} - \{0\}$, from $G(S)$ to the generalized Bestvina-Brady group $G_L(S \cup \{0\})$.*

Proof. The generating set for $H(S)$ consists of the directed edges of K^1, which is a subcomplex of L, the generating set for each H_P is identified with the edges of ∂P, and the generating set for G_P is identified with the edges from vertices of ∂P to the cone vertex c_P. This gives a generating set for the group $G(S)$ consisting of directed edges of L. Since each edge of L is either in the image of K^1 or is incident on some cone vertex, this generating set for $G(S)$ consists of all of the directed edges of L. But the generating set for the presentation for $G_L(S \cup \{0\})$ given in [18, defn. 1.1] is also the directed edges of L. Hence there are natural mutually inverse bijections between the generators of $G(S)$ and the generators of $G_L(S \cup \{0\})$. Using the presentation for $G(S)$ that is implicit in its description as a graph of groups, together with the presentation for $G_L(S \cup \{0\})$ given in [18, Defn. 1.1], these bijections send relators to valid relations, and so they induce group isomorphisms. $\qquad\square$

Next we consider the case $Z = \{k^n : n \geq 0\}$ for some $k \in \mathbb{Z}$ with $|k| > 1$. In this case the group H_P has an injective but non-surjective self homomorphism $\phi = \phi_P$ defined by $\phi(a_i) := a_i^k$ for each of the edge generators a_i. In this case the natural choice for G_P is the ascending HNN-extension

$$G_P = \langle a_1, \ldots, a_l, t : a_i^t = a_i^k,\ a_1 a_2 \cdots a_l = 1 \rangle,$$

in which conjugation by the stable letter $t = t_P$ acts by applying the homomorphism ϕ.

Proposition 12. *The presentation 2-complex for the finite presentation for G_P given above is aspherical.*

Proof. Let $Y = Y_P$ be the presentation 2-complex for the small cancellation presentation for H_P, and let $f : Y \to Y$ be the based cellular map that induces $\phi : H_P \to H_P$ on fundamental groups. The natural choice of Eilenberg-Mac Lane space for G_P is the mapping torus $M = M_f$ of f. Since f sends relators to relators, there is an easy way to put a CW-structure on M, with finite 1-skeleton. Each i-cell of Y contributes one i-cell and one $(i+1)$-cell to M. Hence the cells of M are: one 0-cell, $l + 1$ 1-cells labelled by the generators $a_1, \ldots, a_l, t$, one family of l trapezoidal 2-cells (coming from the 1-cells of Y) whose boundaries are the words $t a_i t^{-1} a_i^{-k}$, an infinite family of 2-cells and an infinite family of 3-cells. For $n \geq 0$, denote by e_n the 2-cell that corresponds to the relator $a_1^{k^n} a_2^{k^n} \cdots a_l^{k^n}$, and for $n \geq 1$ let E_n be the 3-cell coming from the 2-cell e_{n-1}, so that the boundary of E_n consists of e_{n-1}, $-e_n$, and k^{n-1} copies of each of the trapezoidal 2-cells. For $n \geq 0$, let M_n be the subcomplex of M that contains the 1-skeleton, the trapezoidal 2-cells, the 2-cells $e_0, \ldots, e_n$ and the 3-cells $E_1, \ldots, E_n$. There is a deformation retraction of E_n onto $\partial E_n - \mathrm{Int}(e_n)$, and combining this with the identity map on the rest of M_{n-1} defines a deformation retraction of M_n onto M_{n-1}. Applying these retractions successively so that the nth of the retractions happens during the interval $[1/2^n, 1/2^{n-1}]$ gives a deformation retraction of M onto M_0. Since M_0 is the presentation 2-complex described in the statement this implies that M_0 is aspherical. $\square$

Any finitely generated subgroup of a finitely presented group is recursively presented, so if the group H_P embeds into a group G_P of type F, it must be the case that Z is recursively enumerable. Sapir has given an aspherical version of the Higman embedding theorem, stating that any finitely generated group with an aspherical recursive presentation can be embedded into a group with a finite aspherical presentation [24]. Hence one obtains the following statement.

Proposition 13. *Suppose that P is a polygon of perimeter at least 7, and let H_P be the corresponding polygon group, which depends on $Z \subseteq \mathbb{Z} - \{0\}$ as well as on P. The following statements are equivalent.*

- *The set Z is recursively enumerable;*

- *H_P embeds in a finitely presented group;*

- *H_P embeds in a group admitting a finite 2-dimensional Eilenberg-Mac Lane space.*

6   A suitable 2-complex

We construct a 2-complex K having the required properties in stages. Initially we ignore the requirements of large girth and acyclicity and construct a 2-complex K_1 with perfect fundamental group and the required small cancellation property. By passing to a subcomplex we obtain an acyclic subcomplex K_2. Finally, by subdividing the 1-skeleton of K_2, with a corresponding increase in the number of sides of each polygon, we obtain K.

Fix a prime power q, and fix $d \geq 3$ so that d divides $q + \epsilon$ for some $\epsilon \in \{\pm 1\}$. Now let $G = PGL(2, q)$ be the 2-dimensional projective general linear group over the field with q elements. There is a natural action of G on the projective line, a set of $q + 1$ points. The 1-skeleton K_1^1 is the complete graph with vertex set the projective line. By construction G acts on K_1^1, and the action is triply-transitive on the vertex set K_1^0. Note also that $|G| = q(q^2 - 1)$ is equal to the number of ordered triples of elements of K_1^0.

Let g be an element of G of order d. Our complex will depend on the pair (d, q) and on the conjugacy class of the element g. The centralizer in G of g is cyclic of order $q + \epsilon$, and the normalizer of the subgroup generated by g is dihedral of order $2(q + \epsilon)$. It follows that the conjugacy class of g contains $q(q - \epsilon)$ elements and that it is closed under taking inverses. The element g is a power of an element of order $q + \epsilon$. From this it follows that the permutation action of g on K_1^0 has $1 - \epsilon$ fixed points and $(q + \epsilon)/d$ cycles of length d. Thus each pair of the form x, x^{-1} of elements of the conjugacy class of g gives a recipe for attaching $(q + \epsilon)/2d$ distinct d-sided polygons to K_1^1. The complex K_1 is thus defined by attaching $|G|/2d$ distinct d-gons to K_1^1.

Proposition 14. *Each piece of the intersection of two polygons of K_1 contains at most one edge.*

Proof. Say that an ordered triple (u, v, w) of vertices of K is contained in a polygon P if v is a vertex of P and the edges $\{u, v\}$ and $\{v, w\}$ are both contained in P. Clearly each d-gon contains exactly $2d$ ordered triples. It will suffice to show that no ordered triple is contained in more than one polygon of K_1, because the vertices contained in a piece of length 2 would form such a triple.

Since there are $q(q^2 - 1)$ distinct ordered triples in K_1^0 and $q(q^2 - 1)/2d$ polygons in K_1, the average number of polygons containing any fixed ordered triple is equal to 1. But G acts on K_1, and G acts transitively on the ordered triples. Hence each ordered triple is contained in exactly one polygon of K_1, as required. □

In the case when $d = q+1$, the Euler characteristic of K_1 is equal to 1. These complexes were studied by Aschbacher-Segev, and many of them are rationally acyclic [1]. The only Aschbacher-Segev complexes that are known to be (integrally) acyclic are the trivial case $(d, q) = (3, 2)$ in which case K_1 is a single triangle, and the case $(d, q) = (5, 4)$, which is the well-known example of an acyclic 2-complex referred to as the 2-skeleton of the Poincaré homology 3-sphere. We cannot use these complexes because we need $d \geq 7$ to ensure the $C'(1/6)$ condition.

For $d < q + 1$, the complex K_1 has too many 2-cells to be acyclic, but frequently $H_1(K_1)$ is trivial. Consider the case $(d, q) = (7, 8)$. In this case K_1 has 36 2-cells, and a calculation shows that $H_1(K_1) = 0$ and that $H_2(K_1)$ is free abelian of rank 8. (There are three conjugacy classes of elements of order 7 in $G = PGL(2, 8)$, but the action of

the Galois group of the field of 8 elements by outer automorphisms of G permutes the three classes, so there is only one isomorphism type of complex K_1.) For this K_1 there are many ways to remove eight 2-cells to leave an acyclic 2-complex: randomly removing 2-cells to reduce the rank of H_2 often finds such a complex. The simplest way to produce such a complex that we have found is to fix one of the vertices v_0 of K_1, and to discard the eight 2-cells that are *not* incident on v_0.

This 2-complex K_2 has all of the required properties except that its girth is 3; by subdividing each edge into five, one obtains a new acyclic complex K with girth 15 whose 28 2-cells are 35-gons. Alternatively, one may subdivide the edges not incident on v_0 into five and subdivide the edges incident on v_0 into four to obtain another suitable acyclic complex K with girth 13 whose 28 2-cells are 33-gons.

7 Closing remarks

The groups $G(S)$ for $S \neq \emptyset$ are known to have cohomological dimension 2, but for most of them we have been unable to construct a 2-dimensional Eilenberg-Mac Lane space.

In contrast to the Bestvina-Brady construction, graphical small cancellation is purely 2-dimensional, and so our methods cannot be used to construct (for example) groups that are finitely presented but not type F.

We have been unable to find a suitable complex K admitting an action of a non-trivial finite group Q so that the fixed point set K^Q is empty. Such a pair might give an alternative construction for the main examples in [20].

It is easy to see that the groups $H(S)$ fall into uncountably many quasi-isometry classes using Bowditch's argument [4]. It is less clear what happens with the groups $G(S)$, except in the case when $G(S)$ is isomorphic to a generalized Bestvina-Brady group $G_L(S)$, which was covered in [17]. The results in [22] do not seem to apply directly. In the language of [22], $\psi : S \mapsto G(S)$ defines an injective function from the Cantor set of subsets of Z to the space of marked groups generated by the edges of L. To apply the main theorem of [22] it would suffice to show that the image of ψ is closed.

Conditions 1–5 of Definition 1 are essential for our construction to work. But condition 6 may be weakened, giving a weaker conclusion. For example, if K is only assumed to be 1-acyclic (or equivalently to have perfect fundamental group), then each $G(S)$ will be FP_2 and if K is only assumed to be rationally acyclic, then each $G(S)$ will be $FP(\mathbb{Q})$. Of course in these cases one could also weaken the hypotheses on G_P.

References

[1] M. Aschbacher and Y. Segev, *A fixed point theorem for groups acting on finite 2-dimensional acyclic simplicial complexes*, Proc. London Math. Soc. **67** (1993) 329–354.

[2] M. Bestvina and N. Brady, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997) 445–470.

[3] R. Bieri, *Homological dimension of discrete groups*, second edition, Queen Mary College Mathematical Notes, Queen Mary College, Department of Pure Mathematics, London (1981).

[4] B. H. Bowditch, *Continuously many quasi-isometry classes of 2-generator groups*, Comment. Math. Helv. **73** (1998) 232–236.

[5] N. Brady, *Branched coverings of cubical complexes and subgroups of hyperbolic groups*, J. London Math. Soc. (2) **60** (1999) 461–480.

[6] K. S. Brown, *Finiteness properties of groups*, J. Pure Appl. Algebra **44** (1987), 45–75.

[7] K. S. Brown, *Cohomology of Groups* Graduate Texts in Mathematics **87**, Springer Verlag, New York-Berlin (1982).

[8] T. Brown, *Ph. D. Thesis*, University of Southampton (2020).

[9] K.-U. Bux, M. G. Fluch, M. Marschler and S. Witzel, *The braided Thompson's groups are F_∞. With an appendix by Zaremsky*, J. reine angew. Math. **718** (2016) 59–101.

[10] K.-U. Bux and C. Gonzalez, *The Bestvina-Brady construction revisited: geometric computation of Σ-invariants for right-angled Artin groups*, J. London Math. Soc. (2) **60** (1999) 793–801.

[11] K.-U. Bux, R. Köhl and S. Witzel, *Higher finiteness properties of reductive arithmetic groups in positive characteristic: the Rank Theorem*, Ann. Math. **177** (2013) 311–366.

[12] M. W. Davis, *The cohomology of a Coxeter group with group ring coefficients*, Duke Math. J. **91** (1998) 297–314.

[13] M. W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series **32** Princeton Univ. Press, Princeton, NJ (2008).

[14] W. Dicks and I. J. Leary, *Presentations for subgroups of Artin groups*, Proc. Amer. Math. Soc. **127** (1999) 343–348.

[15] M. Gromov, *Random walks in random groups*, Geom. Funct. Anal. **13** (2003) 73–146.

[16] G. Higman, *Subgroups of finitely presented groups*, Proc. Roy. Soc. London Ser. A **262** (1961) 455–475.

[17] R. P. Kropholler, I. J. Leary and I. Soroko, *Uncountably many quasi-isometry classes of groups of type FP*, arXiv:1712.05826.

[18] I. J. Leary, *Uncountably many groups of type FP*, Proc. London Math. Soc. (3) **117** (2018) 246–276.

[19] I. J. Leary, *Subgroups of almost finitely presented groups*, Math. Ann. **372** (2018) 1383–1391.

[20] I. J. Leary and B. E. A. Nucinkis, *Some groups of type VF*, Invent. Math. **151** (2003) 135–165.

[21] R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*, Classics in Mathematics, Springer-Verlag, Berlin (2001).

[22] A. Minasyan, D. Osin and S. Witzel, *Quasi-isometric diversity of marked groups*, arXiv:1911.01137.

[23] Y. Ollivier, *On a small cancellation theorem of Gromov*, Bull. Belg. Math. Soc. Simon Stevin **13** (2006) 75–89.

[24] M. Sapir, *A Higman embedding preserving asphericity*, J. Amer. Math. Soc. **27** (2014) 1–42.

[25] J. Stallings, *A finitely presented group whose 3-dimensional integral homology is not finitely generated*, Amer. J. Math. **85** (1963) 541–543.

[26] R. Skipper, S. Witzel and M. Zaremsky, *Simple groups separated by finiteness properties*, Invent. Math. **215** (2019) 713–740.

Author's addresses:

`i.j.leary@soton.ac.uk`   `t.brown@soton.ac.uk`

School of Mathematical Sciences,
University of Southampton,
Southampton,
SO17 1BJ