PERGAMON

Topology 42 (2003) 1065–1082

TOPOLOGY

www.elsevier.com/locate/top

Finiteness and CAT(0) properties of diagram groups

Daniel S. Farley

The Pennsylvania State University, 218 McAllister Building, University Park, PA 16802, USA

Received 15 October 2000; accepted 16 April 2002

Abstract

Any diagram group over a finite semigroup presentation acts properly, freely, and cellularly by isometrices on a proper CAT(0) cubical complex.

The existence of a proper, cellular action by isometries on a CAT(0) cubical complex has powerful consequences for the acting group G. One gets, for example, a proof that G satisfies the Baum–Connes conjecture.

Any diagram group over a finite presentation of a finite semigroup is of type $\mathscr{F}_\infty$.

MSC: 20F65; 20F67; 57M07; 20F06

Keywords: Thompson's group; Diagram groups; Cubical complex; CAT(0) space; Type $\mathscr{F}_\infty$

1. Introduction

Richard Thompson first defined his group F in 1965, in the course of his work in logic [14]. One can define F by the following presentation:

$$\langle x_0, x_1, x_2, \ldots, x_n, \ldots \mid x_i^{-1} x_j x_i = x_{j+1} \ (i < j) \rangle.$$

It can also be conveniently defined as a certain group of piecewise linear homeomorphisms of the unit interval [5]. Since its introduction, F has appeared in many different mathematical settings—see [8].

Victor Guba discovered that F is a diagram group in the sense recalled in Section 2. Guba and Sapir, motivated in part by this connection, developed a substantial theory of diagram groups in [10,11].

E-mail address: farley@math.psu.edu (D.S. Farley).

PII: S0040-9383(02)00029-0

Diagram groups are defined in terms of semigroup presentations $\mathcal{P}=\langle\Sigma \mid \mathcal{R}\rangle$. If a positive word w in the alphabet Σ is fixed, then a certain set of semigroup diagrams over $\mathcal{P}$, "the set of (w,w)-diagrams", forms a group with respect to a natural operation. This group is called the *diagram group over $\mathcal{P}$ based at w*, and denoted $\mathcal{D}(\mathcal{P},w)$. An isomorphism between Thompson's group F and the diagram group over $\mathcal{P}=\langle x \mid x=x^2\rangle$ based at x is described in the appendix.

The first goal of this paper is to build a contractible, free $\mathcal{D}(\mathcal{P},w)$-complex $\tilde{K}(\mathcal{P},w)$ for each diagram group $\mathcal{D}(\mathcal{P},w)$. This is done in Section 3. The definition of $\tilde{K}(\mathcal{P},w)$ was originally suggested by a construction from [10]. In that paper, Guba and Sapir defined a two-dimensional complex $\mathcal{K}(\mathcal{P})$, which they call Squier's complex. The 2-skeleton of $\tilde{K}(\mathcal{P},w)$ can be described as the universal cover of the connected component of $\mathcal{K}(\mathcal{P})$ containing the base point w. (Vertices in $\mathcal{K}(\mathcal{P})$ are words in the alphabet Σ, and vice versa.) The complex $\tilde{K}(\mathcal{P},w)$ is a natural extension of $\widetilde{\mathcal{K}(\mathcal{P})}$ into higher dimensions.

Theorem 1.1. *If $\mathcal{P}$ is a finite semigroup presentation, then $\tilde{K}(\mathcal{P},w)$ is a proper* CAT(0) *cubical complex with respect to its natural metric and $\mathcal{D}(\mathcal{P},w)$ acts properly, freely, cellularly and by isometrices.*

The existence of such a group action on a CAT(0) cubical complex has many consequences for the acting group, most notably:

Corollary 1.2. *If $\mathcal{P}$ is a finite semigroup presentation then $\mathcal{D}(\mathcal{P},w)$ satisfies the Baum–Connes conjecture.*

The specific references which establish this connection, and more corollaries, are collected at the end of Section 3.

In [10] Guba and Sapir showed that any diagram group over a finite presentation of a finite semigroup must be finitely presented. Since F is the primary example of such, and F is known to be of type $\mathcal{F}_\infty$ [7], it is natural to expect the stronger conclusion that diagram groups over such presentations are always of type $\mathcal{F}_\infty$.

Theorem 1.3. *If $\mathcal{P}$ is a finite presentation of a finite semigroup then $\mathcal{D}(\mathcal{P},w)$ is of type $\mathcal{F}_\infty$.*

This paper is organized in the following way:

The second section contains a brief review of semigroup diagrams and diagram groups. An important partial order, called the prefix partial order, is also defined. Everything in the section, except for the partial order, has already appeared in [10].

The third section contains the construction of the CAT(0) cubical complex $\tilde{K}(\mathcal{P},w)$. The section ends with a subsection containing the corollaries of the construction.

The fourth section contains a proof that every diagram group over a finite presentation of a finite semigroup is of type $\mathcal{F}_\infty$.

The appendix contains a sketch of the proof that Thompson's group F is the diagram group over $\langle x \mid x=x^2\rangle$ based at x. The proof is intended to be understandable to the reader who has read Section 2 and is familiar with the interpretation of F as a group of PL homeomorphisms.

There is some overlap with Melanie Stein's paper [18]. She built contractible cubical complexes for many groups of PL homeomorphisms, among them Thompson's group F and, more generally, the generalized Thompson's groups F_n (from [6]). Guba and Sapir show that F_n is isomorphic to the diagram group over the presentation $\langle x \mid x = x^n \rangle$ based at x. Stein's constructions for these particular groups are actually equivalent to the ones provided in Section 3, but the fact that these cubical complexes are CAT(0) is new.

2. Semigroup diagrams

2.1. Definitions and basic properties of diagrams

Here some general facts about diagrams and diagram groups are recalled. A more detailed account of these subjects is in [10]. Note that this paper carries over the conventions from [10], and in particular the convention that $\mathcal{R}$ contains no relation of the form $u = u$, where $u \in \Sigma^+$.

Consider, for example, the semigroup presentation $\mathcal{P} = \langle a, b, c, | ab = ba, ac = ca, bc = cb \rangle$. The *semigroup diagrams over* $\mathcal{P}$ are defined inductively. A positive word in the generators, for instance acb, may be represented by a labelled oriented arc as in Fig. 1(a), where the edges are oriented from left to right. This is the simplest type of semigroup diagram. If Δ is a diagram and some positive subpath of its bottom path is labelled by a word $u_1 \in \Sigma^+$, and $(u_1 = u_2) \in \mathcal{R}$, then one may attach a positive path labelled by u_2 below the bottom path of Δ to obtain a new diagram Δ'. Two steps of this procedure are illustrated in Figs. 1(b) and (c).

A diagram whose top path is labelled by u and whose bottom path is labelled by v is called a (u, v)-*diagram*. It will be convenient in this connection to let $*$ denote an arbitrary or dummy word. Thus, a diagram whose top label is u is a $(u, *)$-*diagram* and a diagram whose bottom label is u is a $(*, u)$-*diagram*.

The closure of a bounded complementary region of a diagram Δ is a *cell of* Δ. A diagram is *trivial* if it has no cells. An *atomic* diagram is one with at most one cell.

Two diagrams Δ_1 and Δ_2 are *isotopic* if there is an isotopy of the plane carrying Δ_1 to Δ_2 which takes vertices to vertices and edges to edges, and matches the orientations and labels on the edges. If two diagrams are isotopic, then one writes $\Delta_1 \equiv \Delta_2$.

If Δ_1 is a (u, v)-diagram and Δ_2 is a (v, w)-diagram, then the *concatenation* $\Delta_1 \circ \Delta_2$ is obtained by identifying the bottom path of Δ_1 with the top path of Δ_2, using suitable representatives of the isotopy classes of Δ_1 and Δ_2, as in Fig. 2. This operation is well-defined on isotopy classes and associative. Thus, for a fixed word $w \in \Sigma^+$, the (w, w)-diagrams over $\mathcal{P}$ form a monoid; the identity element is the trivial (w, w)-diagram.

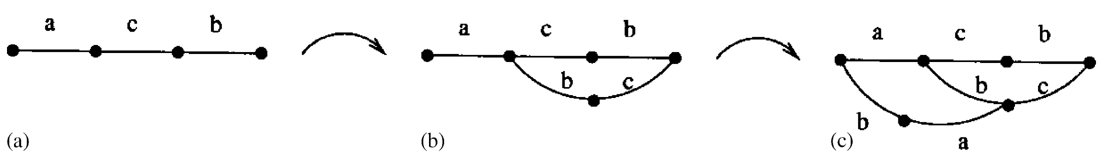

Fig. 1.

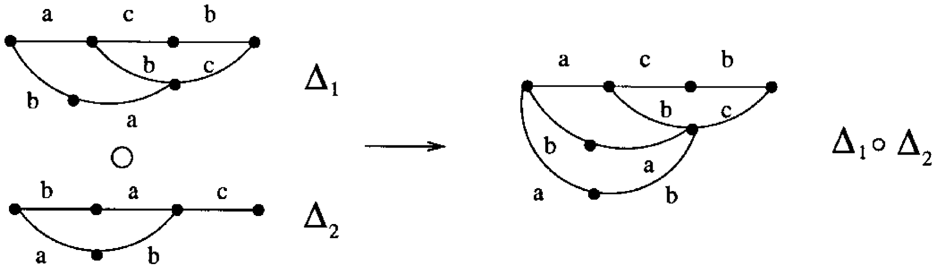

Fig. 2.

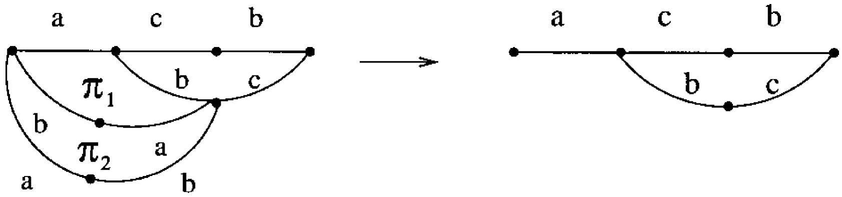

Fig. 3.

A pair of cells π_1 and π_2 in a diagram forms a *dipole* if the bottom path of π_1 is identical to the top path of π_2 and the label of the top path of π_1 is equal to the label of the bottom path of π_2 in the free semigroup Σ^+. To *reduce* a dipole one removes the portion of the diagram lying between the top path of π_1 and the bottom path of π_2, and identifies the top path of π_1 with the bottom path of π_2. This procedure is illustrated in Fig. 3. The inverse process is called *inserting* a dipole. If the diagram Δ_1 may be obtained (up to isotopy) from Δ_2 by repeatedly inserting and removing dipoles, then Δ_1 and Δ_2 are *equivalent modulo dipoles*. One writes $\Delta_1 = \Delta_2$. A diagram without dipoles is *reduced*. Guba and Sapir showed that every diagram Δ is equivalent modulo dipoles to a unique reduced diagram, denoted by $r(\Delta)$.

Concatenation induces a well-defined operation on the equivalence classes of diagrams modulo dipoles. There is, moreover, an inverse for any equivalence class of diagrams modulo dipoles. Fix a word w in Σ^+. Let $\mathscr{D}(\mathscr{P}, w)$ denote the set of equivalence classes of (w, w)-diagrams modulo dipoles.

Theorem 2.1 (Guba and Sapir [10]). *The set $\mathscr{D}(\mathscr{P}, w)$ forms a group under the operation of concatenation.*

The group $\mathscr{D}(\mathscr{P}, w)$ is called the *diagram group over $\mathscr{P}$ with the base w*.

2.2. The prefix partial order

Let Δ_1 and Δ_2 be diagrams. The diagram Δ_1 is a *prefix* of Δ_2 if there is some third diagram Θ such that $\Delta_1 \circ \Theta \equiv \Delta_2$. One writes $\Delta_1 \leqslant \Delta_2$; the relation $\leqslant$ is a partial order on the reduced diagrams over the presentation $\mathscr{P}$, called the *prefix partial order*.

Proposition 2.2. (i) *A reduced diagram Δ has only finitely many prefixes.*

(ii) *If Δ_1 and Δ_2 are reduced diagrams with the same top label then $\{\Delta_1, \Delta_2\}$ has a greatest lower bound in the prefix partial order.*

(iii) *If Δ_1 and Δ_2 are reduced diagrams with an upper bound in the prefix partial order then $\{\Delta_1, \Delta_2\}$ has a least upper bound.*

Proof. (i) This may easily be proved by induction of the number of cells.

(ii) The greatest lower bound of $\{\Delta_1, \Delta_2\}$ may be described as the largest prefix of Δ_1 whose cells are all removed in forming $r(\Delta_2^{-1} \circ \Delta_1)$.

(iii) This is a formal consequence of i) and ii). The proof is an exercise. □

A diagram Ψ is *thin* if the top path of each of its cells is a subpath of the top path of Δ.

Lemma 2.3 (Thin diagram lemma). (i) *Let Ψ be a thin diagram which contains n cells. There is a one-to-one correspondence $\rho_\Psi : \{0,1\}^n \to \{\Psi' \mid \Psi' \leqslant \Psi\}$. If $(a_1,\ldots,a_n),(b_1,\ldots,b_n) \in \{0,1\}^n$, then $\rho_\Psi((a_1,\ldots,a_n)) \leqslant \rho_\Psi((b_1,\ldots,b_n))$ if and only if $a_i \leqslant b_i$, for all $i \in \{1,\ldots,n\}$.*

(ii) *Let $\Psi_1, \Psi_2, \ldots, \Psi_n$ be non-trivial atomic diagrams with a common top label. The set $\{\Psi_1,\ldots,\Psi_n\}$ has an upper bound in the set of all diagrams if and only if $\Psi_1,\ldots,\Psi_n$ represent pairwise disjoint applications of relations to the top label. If there is an upper bound, then there is a least upper bound which is a thin diagram.*

(iii) *Any reduced diagram Δ has a maximal thin suffix Ψ; that is, there is some $\bar{\Delta}$ such that $\bar{\Delta} \circ \Psi \equiv \Delta$ and, if Ψ' is a thin diagram which is such that, for some $\bar{\bar{\Delta}}, \bar{\bar{\Delta}} \circ \Psi' \equiv \Delta$, then $\bar{\Delta} \leqslant \bar{\bar{\Delta}}$.*

(iv) *If Ψ is a thin (v,v)-diagram, for some $v \in \Sigma^+$, then $r(\Psi \circ \Psi)$ is not thin.*

Proof. Suppose that $\pi_1, \pi_2, \ldots, \pi_n$ are the cells of Ψ, numbered from left to right. Let ρ_Ψ send a string $(m_1, m_2, \ldots, m_n)$, $m_i \in \{0,1\}$, to the prefix above the positive path which runs along the top of the cell π_i if $m_i = 0$ and along the bottom if $m_i = 1$. This correspondence has the required properties.

(ii) Let the diagrams $\Psi_1, \Psi_2, \ldots \Psi_n$ represent pairwise disjoint applications of relations to their common top label v. Choose isotopic copies of the Ψ_i so that any two have an identical top path p and $\Psi_i \cap \Psi_j = p$, when $i \neq j$. The least upper bound of $\{\Psi_1,\ldots,\Psi_n\}$ is the union $\Psi_1 \cup \Psi_2 \cup \cdots \cup \Psi_n$.

Conversely, if $\{\Psi_1,\ldots,\Psi_n\}$ has an upper bound Ψ, then $\Psi_1 \cup \cdots \cup \Psi_n$ is a thin diagram and it is clear that this union is the least upper bound.

(iii) Consider the cells π of Δ which are such that the bottom path of π is a subpath of the bottom path of Δ. The union of the boundaries of these cells and the bottom path of Δ is the desired maximal thin suffix Ψ; the portion of Δ which lies above the top path of Ψ is $\bar{\Delta}$.

(iv) Let Ψ be a thin (v,u)-diagram, and let π be the leftmost cell of Ψ. Suppose that the cell π is a (u_1,u_2)-diagram. Let π_1 denote the copy of π in the top part of the concatenation $\Psi \circ \Psi$ and let π_2 denote the copy of π in the bottom. The bottom path of π_1 must intersect the top path of π_2 in an arc containing more than one point. Since $u_1 \neq u_2$ in Σ^+, the cells of π_1 and π_2 also do not form a dipole. These facts together imply that neither cell forms half of a dipole with any other cell in $\Psi \circ \Psi$; therefore both cells remain after reducing all dipoles, and thus $r(\Psi \circ \Psi)$ is not thin. □

Let $I = [\Delta_1, \Delta_2]$ be some non-empty interval in the prefix partial order. The interval I is said to be *elementary* if, for some thin $\Psi, \Delta_1 \circ \Psi \equiv \Delta_2$. Let Δ be a reduced diagram, and let Ψ be some thin diagram such that $\Delta \circ \Psi$ is defined. The *abstract cube determined by* Δ *and* Ψ, denoted by $C(\Delta, \Psi)$, is the set $\{r(\Delta \circ \Psi') \mid \Psi' \leqslant \Psi\}$.

Elementary intervals and abstract cubes are in fact the same thing. If there are no dipoles in the concatenation $\Delta \circ \Psi$, then $C(\Delta, \Psi)$ is equal to $[\Delta, \Delta \circ \Psi]$. If some cells of Δ form dipoles with cells in Ψ, then delete the bottom paths of all such cells in Δ to obtain a new diagram Δ', and interchange the top and bottom paths of the corresponding cells in Ψ to obtain a new thin diagram Ψ'. The cubes $C(\Delta, \Psi)$ and $C(\Delta', \Psi')$ are equal, and the concatenation $\Delta' \circ \Psi'$ is reduced, so $C(\Delta, \Psi) = [\Delta', \Delta' \circ \Psi']$. An inverse procedure shows that, given an elementary interval $[\Delta_1, \Delta_2]$ and $\Delta \in [\Delta_1, \Delta_2]$, there is some Ψ such that $C(\Delta, \Psi) = [\Delta_1, \Delta_2]$. This discussion may be summarized as follows:

Lemma 2.4 (The elementary interval lemma). (i) *Any* $C(\Delta, \Psi)$ *is equal to an interval* $[\bar{\Delta}, \bar{\Delta} \circ \bar{\Psi}]$ *for some* $\bar{\Delta}$ *and* $\bar{\Psi}$, *where* $\bar{\Psi}$ *is thin and* $\bar{\Delta} \circ \bar{\Psi}$ *is reduced. If* $\Delta \circ \Psi$ *is reduced, then* $C(\Delta, \Psi) = [\Delta, \Delta \circ \Psi]$.

(ii) *Given any* $\Delta_1 \in [\Delta, \Delta \circ \Psi]$, *where* $\Delta \circ \Psi$ *is reduced and* Ψ *is thin, there is* Ψ_1 *satisfying* $C(\Delta_1, \Psi_1) = [\Delta, \Delta \circ \Psi]$.

3. Diagram groups and cubical complexes

3.1. Abstract cubical complexes and their realizations

There is a well known and widely used theory of abstract simplicial complexes (see [17]). An abstract simplicial complex J is just a set of vertices V and some set S of finite non-empty sets of vertices which covers V and is closed under taking non-empty subsets. Using the abstract data, one can produce a space $|J|$; given an automorphism of J, one can produce a homeomorphism of $|J|$. The corresponding theory of abstract cubical complexes is sketched in this subsection.

The *standard abstract* n-*cube* is the set $\{0,1\}^n$; the *standard abstract* 0-*cube* $\{0,1\}^0$ is equal to $\{0\}$. A *face of the standard abstract* n-*cube* is a product $A_1 \times \cdots \times A_n$, where each A_i is a non-empty subset of $\{0,1\}$. The standard abstract 0-cube is considered to be a face of itself.

Definition 3.1. An abstract cubical complex $K = (\mathscr{V}, \mathscr{C})$ consists of a non-empty set $\mathscr{V}$, called the vertex set, and a set $\mathscr{C}$ of non-empty subsets of $\mathscr{V}$, called cubes, satisfying:

(i) $\mathscr{C}$ is a cover of $\mathscr{V}$.
(ii) For any $C_1, C_2 \in \mathscr{C}, C_1 \cap C_2 \in \mathscr{C}$ or $C_1 \cap C_2 = \emptyset$.
(iii) For any $C \in \mathscr{C}$ there is a bijection $\phi_C : C \to \{0,1\}^n$, for some n, satisfying:
(∗) if $C_1 \subseteq C$, then $C_1 \in \mathscr{C}$ if and only if $\phi_C(C_1)$ is a face of $\{0,1\}^n$.

The *first derived* K' of an abstract cubical complex $K = (\mathscr{V}, \mathscr{C})$ is the abstract simplicial complex whose set of vertices is $\mathscr{C}$ and whose set of simplices consists of all finite non-empty ascending chains of cubes of K. The *realization* $|K|$ of an abstract cubical complex $K = (\mathscr{V}, \mathscr{C})$ is the realization of the abstract simplicial complex K' in the sense of [17]. An *automorphism* of an abstract cubical complex $K = (\mathscr{V}, \mathscr{C})$ is a bijection $\rho : \mathscr{V} \to \mathscr{V}$ which induces a bijection of $\mathscr{C}$.

Definition 3.2 (Bridson and Haefliger [3]). Let I denote the unit interval. A cubical complex L is the quotient of a disjoint union of cubes $X = \amalg_{\Lambda} I^{n_\lambda}$ by an equivalence relation $\sim$. The restrictions $p_\lambda : I^{n_\lambda} \to L$ of the natural projection $p : X \to L = X/\sim$ are required to satisfy:

1. for every $\lambda \in \Lambda$ the map p_λ is injective;
2. if $p_\lambda(I^{n_\lambda}) \cap p_{\lambda'}(I^{n_{\lambda'}}) \neq \emptyset$ then there is an isometry $h_{\lambda,\lambda'}$ from a face $T_\lambda \subset I^{n_\lambda}$ onto a face $T_{\lambda'} \subset I^{n_{\lambda'}}$ such that $p_\lambda(x) = p_{\lambda'}(x')$ if and only if $x' = h_{\lambda,\lambda'}(x)$.

There is a natural metric d_L on any cubical complex L, called the *intrinsic metric*, which is defined in terms of the standard metrics on the individual cubes. The precise definition is in [3]. For this work, it is enough to note that an isometry of (L, d_L) is induced by any self-homeomorphism of L which permutes the cubes isometrically, and that (L, d_L) is a CAT(0) space (see [3]) under certain conditions, to be specified below.

If $C \in \mathscr{C}$, for some abstract cubical complex $K = (\mathscr{V}, \mathscr{C})$, then $(C, P(C) \cap \mathscr{C})$, where $P(C)$ denotes the power set of C, is an abstract cubical complex; its realization is denoted $|C|$.

The proof of the following proposition is straightforward.

Proposition 3.3. *If $K = (\mathscr{V}, \mathscr{C})$ is an abstract cubical complex, then $|K|$ is a cubical complex; the cubes of $|K|$ are the sets $|C|$ where $C \in \mathscr{C}$. An automorphism of K induces an isometry of $|K|$, with respect to its intrinsic metric.*

The link of a vertex v in a cubical complex K is a simplicial complex which may be described topologically as the boundary of a small ε-neighborhood of v in K; the simplices of the link are the intersections of this boundary with the cubes containing v. A fuller description of the link can be found in [3]. The *abstract link* of v, denoted $\mathrm{lk}_{\mathrm{ab}}(v)$, is the underlying abstract simplicial complex.

Lemma 3.4. *Let K be a cubical complex and suppose $v \in K^0$. The abstract link $\mathrm{lk}_{\mathrm{ab}}(v)$ may be described (up to isomorphism) as follows. The vertices of $\mathrm{lk}_{\mathrm{ab}}(v)$ are all vertices of K which are adjacent to v in the* 1*-skeleton of K; a set S of vertices of $\mathrm{lk}_{\mathrm{ab}}(v)$ forms a simplex if there is some closed cube of K containing $S \cup \{v\}$.*

An abstract simplicial complex $J = (V, S)$ is a *flag complex* if every finite subset of V that is pairwise joined by edges is a simplex. The following fundamental theorem is essentially proved in [3]:

Theorem 3.5 (Bridson and Haefliger [3]; Gromov [9]). *If K is a locally finite, simply connected cubical complex and, for any vertex $v \in K^0$, $\mathrm{lk}_{\mathrm{ab}}(v)$ is a flag complex, then K, with its intrinsic metric, is a proper* CAT(0) *space.*

Proof. Theorem 5.2 and Remark 5.3 on p. 206 of [3] together say that an M_κ-polyhedral complex K in which $\varepsilon(x) > 0$ for any $x \in K$ has curvature $\leqslant \kappa$ if and only if it satisfies the link condition. Here $\varepsilon(x)$ is the distance from x to $\overline{\mathrm{st}(x)} - \mathrm{st}(x)$, where $\mathrm{st}(x)$ is the union of the interiors of the cells that contain x. This function is always strictly positive when K is a cubical complex [3, p. 112].

A theorem of Gromov [9], proved on p. 211 of [3], shows that a locally finite cubical complex has curvature $\leqslant 0$ if and only if the (abstract) link of every vertex is a flag complex.

The Cartan–Hadamard Theorem, from p. 193 of [3], says that a simply connected complete metric space of curvature $\leqslant 0$ is CAT(0). □

3.2. The abstract cubical complex $\tilde{K}_{\mathrm{ab}}(\mathscr{P}, w)$

Fix a semigroup presentation $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ and a word $w \in \Sigma^+$ for the rest of this section. Let $\mathscr{V}(\mathscr{P}, w)$ be the set of all reduced diagrams having w as their top label; let $\mathscr{C}(\mathscr{P}, w) = \{C(\Delta, \Psi): \Delta \in \mathscr{V}(\mathscr{P}, w), \Psi$ is thin, $\Delta \circ \Psi$ is defined $\}$. Let $\tilde{K}_{\mathrm{ab}}(\mathscr{P}, w) = (\mathscr{V}(\mathscr{P}, w), \mathscr{C}(\mathscr{P}, w))$.

Theorem 3.6. *$\tilde{K}_{\mathrm{ab}}(\mathscr{P}, w)$ is an abstract cubical complex.*

Proof. (i) It is clear that $\mathscr{C}(\mathscr{P}, w)$ is a cover of $\mathscr{V}(\mathscr{P}, w)$.

(ii) Let $C_1 = [\Delta_1, \Delta_1 \circ \Psi_1]$ and $C_2 = [\Delta_2, \Delta_2 \circ \Psi_2]$ be members of $\mathscr{C}(\mathscr{P}, w)$. Assume $C_1 \cap C_2 \neq \emptyset$. If $\Delta \in C_1 \cap C_2$, then Δ is an upper bound for $\{\Delta_1, \Delta_2\}$ and thus there is a least upper bound by 2.2. A check shows that $C_1 \cap C_2 = [\mathrm{lub}\{\Delta_1 \Delta_2\}, \mathrm{glb}\{\Delta_1 \circ \Psi_1, \Delta_2 \circ \Psi_2\}]$. The latter set is a cube, since subintervals of elementary intervals are elementary.

(iii) Let $C \in \mathscr{C}(\mathscr{P}, w)$. The abstract cube $C = C(\Delta, \Psi)$ where $\Delta \circ \Psi$ is reduced, by 2.4. Suppose that Ψ contains n cells. Define a partial order on the standard abstract n-cube as follows: $(a_1, \ldots, a_n) \leqslant (b_1, \ldots, b_n)$ if $a_i \leqslant b_i$ for all $i \in \{1, \ldots, n\}$. The intervals of the partially ordered set $\{0,1\}^n$ are precisely the abstract faces. By 2.3, there is an isomorphism of partially ordered sets ρ from the set of prefixes of Ψ to $\{0,1\}^n$. For a prefix Ψ' of Ψ, let $\phi_C(\Delta \circ \Psi') = \rho(\Psi')$.

The map ϕ_C is an isomorphism between the partially ordered sets $\mathscr{C}(\mathscr{P}, w)$ and $\{0,1\}^n$, which therefore carries elementary intervals to faces of $\{0,1\}^n$; thus 3.1 is satisfied. □

3.3. The cubical complex $\tilde{K}(\mathscr{P}, w)$

Using the results of Theorems 3.3 and 3.6, one gets a natural cubical complex which depends only on $\mathscr{P}$ and w.

Definition 3.7. $\tilde{K}(\mathscr{P}, w)$ is the realization of $\tilde{K}_{\mathrm{ab}}(\mathscr{P}, w)$. That is, $\tilde{K}(\mathscr{P}, w) = |\tilde{K}_{\mathrm{ab}}(\mathscr{P}, w)|$.

If $S \subseteq \mathscr{V}(\mathscr{P}, w)$, then $\tilde{K}_S(\mathscr{P}, w)$ is, by definition, the largest subcomplex of $\tilde{K}(\mathscr{P}, w)$ having S as its 0-skeleton.

3.3.1. The group action on $\tilde{K}(\mathscr{P}, w)$

Proposition 3.8. *The diagram group $\mathscr{D}(\mathscr{P}, w)$ acts freely, cellularly, and by isometries on $\tilde{K}(\mathscr{P}, w)$. Two vertices of $\tilde{K}(\mathscr{P}, w)$ are in the same orbit of the action if and only if both have the same bottom label.*

Proof. If $\Delta \in \mathscr{D}(\mathscr{P}, w)$ and $\Delta_1 \in \mathscr{V}(\mathscr{P}, w)$, then $\Delta \cdot \Delta_1 := r(\Delta \circ \Delta_1)$. It is easy to check that this is the definition of a group action on $\mathscr{V}(\mathscr{P}, w)$.

Let $C \in \mathscr{C}(\mathscr{P}, w)$. The abstract cube C is equal to $C(\Delta, \Psi)$ for some $\bar{\Delta}$ and $\bar{\Psi}$. Let $\Delta \in \mathscr{D}(\mathscr{P}, w)$.

$$\Delta \cdot C(\bar{\Delta}, \bar{\Psi}) = \{r(\Delta \circ r(\bar{\Delta} \circ \bar{\Psi}')) : \bar{\Psi}' \leqslant \bar{\Psi}\} = \{r(r(\Delta \circ \bar{\Delta}) \circ \bar{\Psi}') : \bar{\Psi}' \leqslant \bar{\Psi}\} = C(r(\Delta \circ \bar{\Delta}), \bar{\Psi})$$

Thus $D(\mathscr{P}, w)$ acts by automorphisms on $\tilde{K}_{ab}(\mathscr{P}, w)$, and so $\mathscr{D}(\mathscr{P}, w)$ acts cellularly by isometries on $\tilde{K}(\mathscr{P}, w)$ (3.3). If $\Delta \in \mathscr{D}(\mathscr{P}, w)$ fixes a point in $\tilde{K}(\mathscr{P}, w)$ then Δ leaves some $C(\bar{\Delta}, \bar{\Psi})$ invariant. Thus $\Delta \circ \bar{\Delta} = \bar{\Delta} \circ \bar{\Psi}'$ for some $\bar{\Psi}' \leqslant \bar{\Psi}$, and $\Delta \circ \bar{\Delta} \circ \bar{\Psi}' = \bar{\Delta} \circ \bar{\Psi}''$ for some $\bar{\Psi}'' \leqslant \bar{\Psi}$. It follows that $\bar{\Psi}' \circ \bar{\Psi}' = \bar{\Psi}''$, which, by 2.3, is impossible unless $\bar{\Psi}'$ is trivial. It follows that $\Delta = 1$, so $\mathscr{D}(\mathscr{P}, w)$ acts freely.

The second statement is an easy observation. Indeed, if $\Delta_1, \Delta_2 \in \mathscr{V}(\mathscr{P}, w)$ and $\Delta \cdot \Delta_1 = \Delta_2$, then the bottom labels of Δ_1 and Δ_2 must be identical since the bottom label of $\Delta \circ \Delta_1$ is the same as the bottom label of Δ_1, and the bottom (or top) label of a diagram is not changed by reducing dipoles. Conversely, if Δ_1 and Δ_2 are members of $\mathscr{V}(\mathscr{P}, w)$ with the same bottom label, then $r(\Delta_2 \circ \Delta_1^{-1}) \in \mathscr{D}(\mathscr{P}, w)$ takes Δ_1 to Δ_2. □

There is a useful way to produce $\mathscr{D}(\mathscr{P}, w)$-invariant subcomplexes of $\tilde{K}(\mathscr{P}, w)$. If W is a subset of [w], the equivalence class of w modulo $\mathscr{P}$, set $\bar{W} = \{\Delta \in \mathscr{V}(\mathscr{P}, w)$: the bottom label of Δ is in $W\}$.

Proposition 3.9. *If $W \subseteq [w]$, then $\tilde{K}_{\bar{W}}(\mathscr{P}, w)$ is a $\mathscr{D}(\mathscr{P}, w)$-invariant subcomplex of $\tilde{K}(\mathscr{P}, w)$. If W is finite and $\mathscr{P}$ is a finite semigroup presentation, then $\tilde{K}_{\bar{W}}(\mathscr{P}, w)$ is also $\mathscr{D}(\mathscr{P}, w)$-finite.*

Proof. The first statement is an easy consequence of Proposition 3.8. If $\mathscr{P}$ is a finite semigroup presentation, then $\tilde{K}(\mathscr{P}, w)$ is locally finite; this will follow (without circularity) from the argument in 3.3.3. Now, for each word v in W, choose a vertex of $\tilde{K}(\mathscr{P}, w)$ having v as its bottom label, and call the resulting (finite) set of vertices V. The union of all the closed cubes of $\tilde{K}(\mathscr{P}, w)$ which contain a vertex from V is a compact fundamental domain for the action of $\mathscr{D}(\mathscr{P}, w)$ on $\tilde{K}_{\bar{W}}(\mathscr{P}, w)$. □

3.3.2. The contractibility of $\tilde{K}(\mathscr{P}, w)$

Theorem 3.10. *$\tilde{K}(\mathscr{P}, w)$ is contractible with respect to the weak topology.*

Proof. Let $f : S^n \to \tilde{K}(\mathscr{P}, w)$ be a continuous map. By Whitehead's theorem, it is enough to show that f is null homotopic. Since $f(S^n)$ is compact, $f(S^n)$ is contained in the union of finitely many closed cubes. Let S be the (finite) set of vertices which is contained in this union. Let $S' = \{\Delta \in \mathscr{V}(\mathscr{P}, w) \mid \Delta \leqslant \bar{\Delta}$, for some $\bar{\Delta} \in S\}$. Note that S' is still finite, by 2.2.

The complex $\tilde{K}_{S'}(\mathscr{P}, w)$ is collapsible. For let Δ be a maximal element of S'. By 2.3, Δ has a maximal thin suffix Ψ. Suppose $\Delta' \leqslant \Delta$ is such that $\Delta \equiv \Delta' \circ \Psi$. One needs to show, first, that $\tilde{K}_S(\mathscr{P}, w) = \tilde{K}_{S-\{\Delta\}}(\mathscr{P}, w) \cup |C(\Delta', \Psi)|$. So let $|C(\bar{\Delta}, \bar{\Psi})|$ be a cube in $\tilde{K}_{S'}(\mathscr{P}, w)$ and assume that $\bar{\Delta} \circ \bar{\Psi}$ is reduced. If $\bar{\Delta} \circ \bar{\Psi} \not\equiv \Delta$, then $|C(\bar{\Delta}, \bar{\Psi})| \subseteq \tilde{K}_{S'-\{\Delta\}}(\mathscr{P}, w)$. If $\bar{\Delta} \circ \bar{\Psi} \equiv \Delta \equiv \Delta' \circ \Psi$, then $\Delta' \leqslant \bar{\Delta}$, so $|C(\bar{\Delta}, \bar{\Psi})| \subseteq |(C(\Delta', \Psi)|$. The space $\tilde{K}_{S'-\{\Delta\}}(\mathscr{P}, w) \cap |C(\Delta', \Psi)|$ is the union of all codimension one faces of $|C(\Delta', \Psi)|$ which do not contain Δ. Topologically, then, this intersection is a ball of dimension one less than the dimension of $|C(\Delta', \Psi)|$. It follows that $\tilde{K}_{S'}(\mathscr{P}, w)$ collapses onto $\tilde{K}_{S'-\{\Delta\}}(\mathscr{P}, w)$. It follows by an induction on the cardinality of S' that $\tilde{K}_{S'}(\mathscr{P}, w)$ is collapsible. □

In case $\tilde{K}(\mathscr{P},w)$ is locally finite, the weak topology and the metric topology on $\tilde{K}(\mathscr{P},w)$ coincide; since $\tilde{K}(\mathscr{P},w)$ is locally finite when $\mathscr{P}$ is a finite presentation (see the lemma at the beginning of the next subsubsection), $\tilde{K}(\mathscr{P},w)$ is contractible in either topology when $\mathscr{P}$ is finite.

3.3.3. The link of a vertex in $\tilde{K}(\mathscr{P},w)$

Suppose $\Delta \in \tilde{K}(\mathscr{P},w)$ and the bottom label of Δ is v. Define an abstract simplicial complex $J(\mathscr{P},v)$ as follows: the vertex set of $J(\mathscr{P},v)$ consists of all non-trivial atomic diagrams with top label v; a simplex of $J(\mathscr{P},v)$ is a set of vertices having an upper bound in the prefix partial order.

Lemma 3.11. *The abstract simplicial complex $J(\mathscr{P},v)$ has the following properties*:

(i) *its realization is finite dimensional for any semigroup presentation $\mathscr{P}=\langle\Sigma \mid \mathscr{R}\rangle$ and $v\in\Sigma^+$*;

(ii) *it is finite when the presentation $\mathscr{P}$ is finite*;

(iii) *it is a flag complex.*

Proof. (i) Any thin diagram having v as its top label can have at most $\ell(v)$ cells. It directly follows that no simplex in $J(\mathscr{P},v)$ can have more than $\ell(v)$ elements. This implies that $\dim(|J(\mathscr{P},v)|) \leqslant \ell(v)-1$.

(ii) Under these hypotheses, there are only finitely many relations and there can be only finitely many ways to apply each one to v; it follows that $J(\mathscr{P},v)$ has only finitely many vertices.

(iii) Let $\{\Psi_1,\ldots,\Psi_n\}$ be vertices of $J(\mathscr{P},v)$ which, pairwise, span a simplex of $J(\mathscr{P},v)$. Thus $\Psi_1,\ldots,\Psi_n$ represent pairwise disjoint applications of relations to v. The set $\{\Psi_1,\ldots,\Psi_n\}$ has a least upper bound Ψ which is a thin diagram (2.3, part (ii)); therefore $\{\Psi_1,\ldots,\Psi_n\}$ is a simplex of $J(\mathscr{P},v)$. □

Now define $\phi_\Delta:\{\Psi \mid \Psi$ is a non-trivial atomic diagram and v is its top label$\} \to \mathscr{A}(\Delta)=\{\Delta' \mid \Delta'$ is adjacent to $\Delta\}$ by the equation: $\phi_\Delta(\Psi)=r(\Delta\circ\Psi)$.

Proposition 3.12. *Let $v\in\Sigma^+$ and $\Delta\in\mathscr{V}(\mathscr{P},w)$ which has v as its bottom label. The function ϕ_Δ induces an isomorphism between the abstract simplicial complexes $J(\mathscr{P},v)$ and* $\mathrm{lk}_{\mathrm{ab}}(\Delta)$.

Proof. A vertex in $\mathrm{lk}_{\mathrm{ab}}(\Delta)$ is a diagram $\Delta'\in\mathscr{V}(\mathscr{P},w)$ which is adjacent to Δ. Such a Δ' is necessarily $r(\Delta\circ\Psi)$ for some non-trivial atomic diagram Ψ with top label v. It follows that Φ_Δ is a surjection. The injectivity of Φ_Δ follows from the left cancellation property for diagrams [10].

A simplex of $J(\mathscr{P},v)$ is a set $\{\Psi_1,\ldots,\Psi_n\}$ of atomic diagrams having a least upper bound Ψ in the prefix partial order; such a least upper bound must be thin. Now $\Phi_\Delta(\{\Psi_1,\ldots,\Psi_n\})=\{r(\Delta\circ\Psi_1),\ldots,r(\Delta\circ\Psi_n)\}$. This is a collection of vertices adjacent to Δ, all of which are contained in the elementary interval $C(\Delta,\Psi)$. This implies $\Phi_\Delta(\{\Psi_1,\ldots,\Psi_n\})$ is a simplex in $\mathrm{lk}_{\mathrm{ab}}(\Delta)$, by Proposition 3.4.

Now suppose that S is a simplex in $\mathrm{lk}_{\mathrm{ab}}(\Delta)$. The set S, by Proposition 3.4, is $\mathscr{A}(\Delta)\cap C(\Delta,\Psi)$ for some thin diagram Ψ. Let $\{\Psi_1,\ldots,\Psi_n\}$ be the set of atomic diagrams which are prefixes of Ψ. It is clear that $\{\Psi_1,\ldots,\Psi_n\}$ is a simplex of $J(\mathscr{P},v)$ and that $\Phi_\Delta(\{\Psi_1,\ldots,\Psi_n\})$ is S. □

3.4. Actions on CAT(0) cubical complexes and corollaries

The following is one of the main theorems of this paper.

Theorem 3.13. *If* $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ *is a finite semigroup presentation and* $w \in \Sigma^+$, *then* $\tilde{K}(\mathscr{P}, w)$ *is a proper* CAT(0) *cubical complex on which* $\mathscr{D}(\mathscr{P}, w)$ *acts properly, cellularly and freely by isometries. The action is cocompact if and only if the equivalence class mod* $\mathscr{P}$ *of the word* w *is finite.*

Proof. By Theorem 3.10, $\tilde{K}(\mathscr{P}, w)$ is contractible; it follows from Lemma 3.11 and Proposition 3.12 that $\tilde{K}(\mathscr{P}, w)$ is locally finite and the link of each vertex is a flag complex. Now it follows from Theorem 3.5 that $\tilde{K}(\mathscr{P}, w)$ is a proper CAT(0) cubical complex.

By Theorem 3.3 and Proposition 3.8, the action of $\mathscr{D}(\mathscr{P}, w)$ on $\tilde{K}(\mathscr{P}, w)$ is free, cellular, and by isometries. Since the action of $\mathscr{D}(\mathscr{P}, w)$ is free and cellular, $\mathscr{D}(\mathscr{P}, w)$ acts by covering transformations. The action of $\mathscr{D}(\mathscr{P}, w)$ must be proper since $\mathscr{D}(\mathscr{P}, w)$ acts isometrically by covering transformations on a proper metric space. (Note that the statement that $\mathscr{D}(\mathscr{P}, w)$ acts properly is redundant.)

The last statement is a consequence of Proposition 3.9. □

In [16], Niblo and Reeves proved that any group acting properly, cellularly, and cocompactly on a CAT(0) cubical complex is biautomatic. As an immediate consequence of the previous theorem and their work, one has:

Corollary 3.14. *If* $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ *is a finite semigroup presentation,* $w \in \Sigma^+$, *and* $[w]$ *is finite, then* $\mathscr{D}(\mathscr{P}, w)$ *is biautomatic.*

A group has the *Haagerup property* if it acts properly, affinely, and isometrically on a Hilbert space. In [15], Niblo and Reeves showed that a group acting properly and cellularly by isometries on a CAT(0) cubical complex has the Haagerup property. The work of Higson and Kasparov [12] shows that a group with the Haagerup property must satisfy the Baum–Connes conjecture. The Baum–Connes conjecture is very strong; in particular, it implies the rational Novikov conjecture in topology, as well as the Kaplansky–Kadison conjecture in functional analysis. For more information, the reader is referred to the survey article by Julg [13]. Bridson pointed out the connection between [12,16] in an unpublished paper [4]; the statement that a group acting properly and cellularly by isometries on a CAT(0) cubical complex must satisfy the Baum–Connes conjecture has since appeared in [13].

The following corollary is an immediate consequence of Theorem 3.13 and the preceding remarks:

Corollary 3.15. *If* $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ *is a finite semigroup presentation and* $w \in \Sigma^+$, *then* $\mathscr{D}(\mathscr{P}, w)$ *satisfies the Baum–Connes conjecture.*

Thus, in particular, Thompson's group F satisfies the Baum–Connes conjecture.

It is a well-known fact that any group acting freely by isometries on a complete CAT(0) space must be torsion-free [3]. This fact and Theorem 3.13 imply:

Corollary 3.16. $\mathscr{D}(\mathscr{P}, w)$ *is torsion-free, for any semigroup presentation* $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ *and word* $w \in \Sigma^+$.

Proof. The discussion before the corollary immediately implies that $\mathscr{D}(\mathscr{P}, w)$ is torsion-free when $\mathscr{P}$ is a finite semigroup presentation.

Let $\mathscr{P}$ be an arbitrary semigroup presentation and let Δ be a non-trivial, reduced (w, w)-diagram over $\mathscr{P}$. Consider the semigroup presentation $\mathscr{P}' = \langle \Sigma' \,|\, \mathscr{R}' \rangle$, where Σ' is the alphabet consisting of all letters of Σ which label edges of Δ and $\mathscr{R}'$ is the set of all relations (u, v) for which there is a cell of Δ which is a (u, v)- or (v, u)-diagram. Since a diagram has only finitely many edges and cells, $\mathscr{P}'$ is a finite presentation. There is a natural embedding of $\mathscr{D}(\mathscr{P}', w)$ into $\mathscr{D}(\mathscr{P}, w)$; the general statement of the corollary now follows easily. □

Using different methods, Guba and Sapir have proved this and several stronger group-theoretic statements about diagram groups over semigroup presentations. Much of their work suggests that each group $\mathscr{D}(\mathscr{P}, w)$ acts by *semisimple* isometries on its space $\tilde{K}(\mathscr{P}, w)$. An isometry of a CAT(0) space X is *semisimple* if it has a fixed point or it acts on some geodesic line of X by a translation.

Question 3.17. Does $\mathscr{D}(\mathscr{P}, w)$ act on $\tilde{K}(\mathscr{P}, w)$ by *semisimple* isometries?

An affirmative answer to this question would have several new group-theoretic consequences for diagram groups and, in particular, for F. See [3] for details.

Bridson showed that every isometry of a finite-dimensional cubical complex is semisimple [2], and this may be taken as evidence for an affirmative answer to this question.

4. A class of groups of type $\mathscr{F}_\infty$

A group G is said to be of type $\mathscr{F}_n$ if there is some $K(G, 1)$ complex having a finite n-skeleton; a group G is $\mathscr{F}_\infty$ if there is a $K(G, 1)$ complex with only finitely many cells in each dimension. Thus, for example, a group is finitely presented if and only if it is of type $\mathscr{F}_2$.

In [10], Guba and Sapir proved that if $\mathscr{P} = \langle \Sigma \,|\, \mathscr{R} \rangle$ is a finite presentation of a finite semigroup and $w \in \Sigma^+$, then $\mathscr{D}(\mathscr{P}, w)$ is finitely presented. Brown and Geoghegan showed in [7] that F is of type $\mathscr{F}_\infty$. The rest of this section is devoted to a proof of the following theorem.

Theorem 4.1. *If $\mathscr{P} = \langle \Sigma \,|\, \mathscr{R} \rangle$ is a finite presentation of a finite semigroup and $w \in \Sigma^+$, then $\mathscr{D}(\mathscr{P}, w)$ is of type $\mathscr{F}_\infty$.*

A group G is *infinite dimensional* if there is no finite dimensional $K(G, 1)$ complex. Thompson's group F was the first group known to be infinite dimensional, torsion-free, and of type $\mathscr{F}_\infty$.

Corollary 4.2. *Under the hypotheses from Theorem* 4.1, *if w is equivalent modulo $\mathscr{P}$ to w^n for some $n > 1$, then $\mathscr{D}(\mathscr{P}, w)$ is an infinite dimensional, torsion-free group of type $\mathscr{F}_\infty$.*

Proof. Guba and Sapir show, in [10], that $\mathscr{D}(\mathscr{P}, w_1)$ is isomorphic to $\mathscr{D}(\mathscr{P}, w_2)$ if w_1 and w_2 are equivalent modulo $\mathscr{P}$ and that, for any words u and v in Σ^+, $\mathscr{D}(\mathscr{P}, u) \times \mathscr{D}(\mathscr{P}, v)$ embeds into $\mathscr{D}(\mathscr{P}, uv)$ in a natural way.

It follows that $\Pi_1^n \mathscr{D}(\mathscr{P}, w)$ embeds into $\mathscr{D}(\mathscr{P}, w^n) \cong \mathscr{D}(\mathscr{P}, w)$. Now any group G having G^n $(n > 1)$ as a subgroup must be either trivial or infinite dimensional. Thus, assuming Theorem 4.1 and using Corollary 3.16, it is enough to show that $\mathscr{D}(\mathscr{P}, w)$ is non-trivial.

Let Δ be a reduced (w, w^n)-diagram over $\mathscr{P}$; such a diagram exists since w is equivalent to w^n over $\mathscr{P}$. (It is proved in [10] that a (w_1, w_2)-diagram over $\mathscr{P}$ exists when w_1 is equivalent to w_2 modulo $\mathscr{P}$.) A non-trivial (w^{n+1}, w^{n+1})-diagram may be obtained by identifying the terminal vertex of Δ with the initial vertex of Δ^{-1}. (This diagram is $\Delta + \Delta^{-1}$ in the terminology of [10]). Now, since $\mathscr{D}(\mathscr{P}, w^{n+1})$ embeds in $\mathscr{D}(\mathscr{P}, w^{2n-1}) \cong \mathscr{D}(\mathscr{P}, w)$, $\mathscr{D}(\mathscr{P}, w)$ is non-trivial. □

Recall that if $H \leqslant G$, then H is a *retract* of G if there is a surjection $\phi : G \to H$ which is the identity on H. A *complete* presentation is one for which every relation in $\mathscr{R}$ is an ordered pair of words in Σ^+ and the string rewriting system associated to $\mathscr{P}$ is complete—see [10].

Lemma 4.3 (Guba and Sapir [10]). *If $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ is a finite presentation of a finite semigroup and $w \in \Sigma^+$, then there is some finite complete presentation $\mathscr{P}' = \langle \Sigma' \mid \mathscr{R}' \rangle$ of a finite semigroup, satisfying $\Sigma \subseteq \Sigma'$ and $\mathscr{R} \subseteq \mathscr{R}'$, such that $\mathscr{D}(\mathscr{P}, w)$ is a retract of $\mathscr{D}(\mathscr{P}', w)$.*

Proof. This is proved on p. 61 of [10], in the course of the proof of their Theorem 10.7. □

Any retract of a group of type $\mathscr{F}_n$ is of type $\mathscr{F}_n$, so Theorem 4.1 follows from:

Theorem 4.4. *If $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ is a finite complete presentation of a finite semigroup and $w \in \Sigma^+$, then $\mathscr{D}(\mathscr{P}, w)$ is of type $\mathscr{F}_\infty$.*

The proof of 4.4 occupies the remaining subsections.

4.1. Finite complete presentations and the descending link

For the rest of the section, fix a finite complete semigroup presentation $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ and a word $w \in \Sigma^+$. Write $u \dot{\to} v$ if v may be obtained from u by a (possibly empty) finite sequence of reductions modulo $\mathscr{P}$. The relation $\dot{\to}$ makes Σ^+ into a partially ordered set. Write $u \leqslant v$ if $v \dot{\to} u$.

The proofs of the following lemmas are simple exercises.

Lemma 4.5. *For any word $w \in \Sigma^+$, $(-\infty, w] = \{v \in \Sigma^+ \mid v \leqslant w\}$ is finite.*

Lemma 4.6. *There is a sequence (w_n) in $[w]$ satisfying*:

(i) $\{w_1, \ldots, w_n, \ldots\} = [w]$, *and*

(ii) *For any n, $(-\infty, w_n] \subseteq \{w_1, \ldots, w_n\}$.*

Fix a sequence (w_n) satisfying (i) and (ii). For each n, let $\tilde{K}(\mathscr{P}, w)_n$ denote the complex $\tilde{K}_{\overline{\{w_1,\ldots,w_n\}}}(\mathscr{P}, w)$, (see Section 3.3.1). By Proposition 3.9, each $\tilde{K}(\mathscr{P}, w)_n$ is $\mathscr{D}(\mathscr{P}, w)$-invariant, $\mathscr{D}(\mathscr{P}, w)$-finite, and $\bigcup_n \tilde{K}(\mathscr{P}, w)_n = \tilde{K}(\mathscr{P}, w)$.

One needs to understand the topology of the complex $\tilde{K}(\mathscr{P}, w)_n$ near those vertices having w_n as their bottom label.

Lemma 4.7. *If Δ_1 and Δ_2 are distinct vertices of $\tilde{K}(\mathscr{P},w)_n$, both having w_n as their bottom label, then*:

(i) *There is no cube of $\tilde{K}(\mathscr{P},w)_n$ which contains both Δ_1 and Δ_2.*

(ii) *The link of Δ_1 in $\tilde{K}(\mathscr{P},w)_n$ is isomorphic, as an abstract simplicial complex, to the full subcomplex of $J(\mathscr{P},w_n)$ generated by the set $\{\Psi \mid \Psi$ is a vertex of $J(\mathscr{P},w_n)$ and v, the bottom label of Ψ, satisfies $v \leqslant w_n\}$.*

Proof. (i) This follows from the observation that every cube C in $\tilde{K}(\mathscr{P},w)$ has a vertex whose bottom label is larger than that of any other vertex in C. Note that this is so since $\mathscr{P}$ contains no relation of the form $u=u$.

(ii) The isomorphism is the restriction of ϕ_{Δ_1} to the set of atomic diagrams which represent moves in the rewriting system corresponding to $\mathscr{P}$. The proof that this is an isomorphism is similar to the argument from Proposition 3.12 and depends on property (ii) from Lemma 4.6. □

The link described in part (ii) of the preceding lemma deserves a name.

Definition 4.8. Let $v\in\Sigma^+$. The descending link of v, denoted $\mathrm{lk}_\downarrow(v)$, is defined to be the full subcomplex of $J(\mathscr{P},v)$ generated by $\{\Psi \mid \Psi$ is a vertex of $J(\mathscr{P},v)$ and u, the bottom label of Ψ, satisfies $u\leqslant v\}$ (q.v. [1]).

If one takes ε small, then $\tilde{K}(\mathscr{P},w)_n - \mathrm{N}_\varepsilon\overline{(\{w_n\})}$ will strong deformation retract onto $\tilde{K}(\mathscr{P},w)_{n-1}$, since the strong deformation retractions from each vertex in $\overline{\{w_n\}}$ are compatible with one another, by Lemma 4.7(i). The boundary of $N_\varepsilon(\Delta)\subset\tilde{K}(\mathscr{P},w)_n$ is homeomorphic to $|\mathrm{lk}_\downarrow(w_n)|$, for each $\Delta\in\overline{\{w_n\}}$, and each such epsilon neighborhood is topologically a cone on this boundary. One therefore obtains a description of $\tilde{K}(\mathscr{P},w)_n$ in terms of $\tilde{K}(\mathscr{P},w)_{n-1}$: up to homotopy, it is obtained from $\tilde{K}(\mathscr{P},w)_{n-1}$ by attaching countably many cones on $|\mathrm{lk}_\downarrow(w_n)|$ along their bases. This leads to a sufficient condition for $\mathscr{D}(\mathscr{P},w)$ to have type $\mathscr{F}_n$.

Proposition 4.9. *If $\mathscr{P}=\langle\Sigma\mid\mathscr{R}\rangle$ is a finite complete presentation, $w\in\Sigma^+$ and there is N such that $|\mathrm{lk}_\downarrow(w_m)|$ is n-connected for all $m>N$, then $\tilde{K}(\mathscr{P},w)_N$ is n-connected. In particular, $\mathscr{D}(\mathscr{P},w)$ is of type $\mathscr{F}_{n+1}$.*

Proof. Let $n\geqslant 2$. Assume that $|\mathrm{lk}_\downarrow(w_m)|$ is n-connected for all $m>N$. A computation using the Mayer–Vietoris sequence of the tried $(\tilde{K}(\mathscr{P},w)_m,\tilde{K}(\mathscr{P},w)_{m-1},\coprod C_i)$ and the connectivity of $|\mathrm{lk}_\downarrow(w_m)|$ shows that the natural map $H_n(\tilde{K}(\mathscr{P},w)_{m-1})\to H_n(\tilde{K}(\mathscr{P},w)_m)$ is an isomorphism, for all $m>N$. It follows from a direct limit argument that $H_n(\tilde{K}(\mathscr{P},w)_N)\to H_n(\tilde{K}(\mathscr{P},w))$ is an isomorphism, so $H_n(\tilde{K}(\mathscr{P},w)_N)=0$.

In the case $n=0$ one attaches the cones C_i one at a time to $\tilde{K}(\mathscr{P},w)_{m-1}$ and uses the Mayer–Vietoris sequence of reduced homology together with a direct limit argument to conclude that $\tilde{H}_0(\tilde{K}(\mathscr{P},w)_{m-1})\to\tilde{H}_0(\tilde{K}(\mathscr{P},w)_m)$ is an isomorphism for all $m>N$. Another direct limit argument then establishes that $\tilde{H}_0(\tilde{K}(\mathscr{P},w)_N)\to\tilde{H}_0(\tilde{K}(\mathscr{P},w))$ is an isomorphism, so $\tilde{H}_0(\tilde{K}(\mathscr{P},w)_N)=0$.

In the case $n=1$ one argues that $\pi_1(\tilde{K}(\mathscr{P},w)_N)\to\pi_1(\tilde{K}(\mathscr{P},w))$ is an isomorphism. The argument is similar to the one from the previous paragraph, but with van Kampen's theorem taking the place of the Mayer–Vietoris sequence.

The proposition then follows by induction on n and the Hurewicz theorem [17]. □

4.2. The connectivity of the descending link

The rest of the argument requires several new definitions. The *initial word* of the (non-trivial) atomic diagram Ψ, denoted $i(\Psi)$, is the label of the positive path which runs from the initial vertex of Ψ to the initial vertex of its cell; the *terminal word* of Ψ, denoted $t(\Psi)$, is the label of the positive path which runs from the terminal vertex of its cell to its terminal vertex. Note that either word (or both) may be empty. A *left-hand vertex of* $\mathrm{lk}_{\downarrow}(v)$ is one with a reduced (or empty) initial word.

If $v \in \Sigma^+$ and s is a suffix of v, then $r\,\mathrm{lk}_{v,\downarrow}(s)$ denotes the full subcomplex of $\mathrm{lk}_{\downarrow}(v)$ generated by those atomic diagrams which represent applications of relations to the suffix s. The subcomplexes $r\,\mathrm{lk}_{v,\downarrow}(s)$ are linearly ordered by inclusion for a fixed v: if s_1 is a suffix of s_2 then $r\,\mathrm{lk}_{v,\downarrow}(s_1)$ is a subcomplex of $r\,\mathrm{lk}_{v,\downarrow}(s_2)$. The complexes $\mathrm{lk}_{\downarrow}(s)$ and $r\,\mathrm{lk}_{v,\downarrow}(s)$ are also naturally isomorphic.

Recall that the *star* of a vertex v_0 in a simplicial complex is the full subcomplex generated by v_0 and the set of all adjacent vertices. The realization of the star of a vertex in any simplicial complex is contractible. The star of a vertex Ψ in $\mathrm{lk}_{\downarrow}(v)$ will be denoted by $\mathrm{st}(\Psi)$.

Lemma 4.10. *Let $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ be a finite complete semigroup presentation and let $v \in \Sigma^+$.*

(i) $|\mathrm{lk}_{\downarrow}(v)| = \bigcup_{\Psi \in \mathscr{L}} |\mathrm{st}(\Psi)|$, *where $\mathscr{L}$ is the set of all left-hand vertices in $\mathrm{lk}_{\downarrow}(v)$.*

(ii) *Let Ψ_1 and Ψ_2 be distinct left-hand vertices of $\mathrm{lk}_{\downarrow}(v)$. $|\mathrm{st}(\Psi_1)| \cap |\mathrm{st}(\Psi_2)| = |r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_1))|$ or $|r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_2))|$, whichever is smaller in the sense of inclusion.*

Proof. (i) Let $|S|$ be a simplex of $|\mathrm{lk}_{\downarrow}(v)|$. Suppose that the corresponding abstract simplex S is $\{\Psi_1, \ldots, \Psi_n\}$. Assume that $i(\Psi_1)$ is the shortest of all of the initial words $i(\Psi_1), \ldots, i(\Psi_n)$. If $i(\Psi_1)$ is reduced with respect to $\mathscr{P}$, then $|S| \subseteq |\mathrm{st}(\Psi_1)|$, and Ψ_1 is a left-hand vertex. If $i(\Psi_1)$ is not reduced, then one can apply some reduction to $i(\Psi_1)$ to get a new atomic diagram Ψ_{n+1} with top label v, and it is easy to see that one can also arrange for $i(\Psi_{n+1})$ to be reduced. Now $|\{\Psi_1, \ldots, \Psi_n, \Psi_{n+1}\}| \subseteq |\mathrm{st}(\Psi_{n+1})|$, and Ψ_{n+1} is a left-hand vertex.

(ii) The intersection $|\mathrm{st}(\Psi_1)| \cap |\mathrm{st}(\Psi_2)|$ is the full subcomplex of $|\mathrm{lk}_{\downarrow}(v)$ generated by the vertices Ψ which are adjacent to both Ψ_1 and Ψ_2. Such a vertex Ψ will represent the application of a relation to both $t(\Psi_1)$ and $t(\Psi_2)$, since $i(\Psi_1)$ and $i(\Psi_2)$ are reduced by assumption. It follows that $|\mathrm{st}(\Psi_1)| \cap |\mathrm{st}(\Psi_2)| \subseteq |r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_1))| \cap |r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_2))|$.

The reverse inclusion is clear. □

The following proposition and Proposition 4.9 together prove Theorem 4.4. Note that if $\mathscr{P}$ is a finite complete presentation of a finite semigroup, then there are only finitely many reduced words in Σ^+.

Proposition 4.11. *Let $\mathscr{P} = \langle \Sigma \mid \mathscr{R} \rangle$ be a finite presentation of a finite semigroup, let $p \in \mathbf{N}$ be strictly larger than the length of any reduced word in Σ^+, and let $q \in \mathbf{N}$ be the length of the longest word in Σ^+ which is the left side of some relation in $\mathscr{R}$. If $\ell(v) \geqslant p + n(p + q - 1)$, then $|\mathrm{lk}_{\downarrow}(v)|$ is $(n-1)$-connected.*

Proof. The proof is by induction on n. Suppose $n=0$. If $\ell(v) \geqslant p$, then v must not be a reduced word, because of the choice of p. Since v is not reduced, there are non-trivial atomic diagrams having v as their top label, so $|\,\mathrm{lk}_{\downarrow}(v)\,|$ is non-empty.

In the general case, one needs the following fact: if X_1 and X_2 are n-connected subcomplexes of a CW-complex X and $X_1 \cap X_2$ is $(n-1)$-connected, then $X_1 \cup X_2$ is n-connected. This is easily proved by induction using van Kampen's theorem, the Mayer–Vietoris sequence, and the Hurewicz theorem.

Assume that the proposition is true for n. Let v be some word of length greater than or equal to $p+(n+1)(p+q-1)$. Let $\{\Psi_1, \Psi_2, \ldots, \Psi_m\}$ be the set of all left-hand vertices of $\mathrm{lk}_{\downarrow}(v)$, and suppose that $t(\Psi_i)$ is a (not necessarily proper) suffix of $t(\Psi_j)$ if $i<j$. Now $|\,\mathrm{st}(\Psi_1)\,| \cap |\,\mathrm{st}(\Psi_2)\,| = |\,r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_1))\,| \cong |\,\mathrm{lk}_{\downarrow}(t(\Psi_1))\,|$. Using the assumptions about p and q and the fact that Ψ_1 is a left-hand vertex, one concludes that the word $t(\Psi_1)$ is of length at least $p+n(p+q-1)$, so $|\,\mathrm{lk}_{\downarrow}(t(\Psi_1))\,|$ is $(n-1)$-connected. It follows that $|\,\mathrm{st}(\Psi_1)\,| \cup |\,\mathrm{st}(\Psi_2)\,|$ is n-connected, since $|\,\mathrm{st}(\Psi_1)\,|$ and $|\,\mathrm{st}(\Psi_2)\,|$ are contractible. Next consider $(\,|\,\mathrm{st}(\Psi_1)\,| \cup |\,\mathrm{st}(\Psi_2)\,|\,) \cap |\,\mathrm{st}(\Psi_3)\,| = (\,|\,\mathrm{st}(\Psi_1)\,| \cap |\,\mathrm{st}(\Psi_3)\,|\,) \cup (\,|\,\mathrm{st}(\Psi_2)\,| \cap |\,\mathrm{st}(\Psi_3)\,|\,) = |\,r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_1))\,| \cup |\,r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_2))\,| = |\,r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_2))\,|$. The word $t(\Psi_2)$ is of length at least $p+n(p+q-1)$, so $|\,r\,\mathrm{lk}_{v,\downarrow}(t(\Psi_2))\,|$ is $(n-1)$-connected. It again follows that $|\,\mathrm{st}(\Psi_1)\,| \cup |\,\mathrm{st}(\Psi_2)\,| \cup |\,\mathrm{st}(\Psi_3)\,|$ is n-connected since $|\,\mathrm{st}(\Psi_1)\,| \cup |\,\mathrm{st}(\Psi_2)\,|$ is n-connected and $|\,\mathrm{st}(\Psi_3)\,|$ is contractible.

One concludes after finitely many iterations of this argument that $|\,\mathrm{lk}_{\downarrow}(v)\,|$ is n-connected. □

Assume that $\mathscr{P}$ is a finite complete presentation of a finite semigroup, and assume that a sequence satisfying the conditions of Lemma 4.6 has been chosen. Since the length of w_n tends to infinity with n, it follows from Propositions 4.9 and 4.11 that $\mathscr{D}(\mathscr{P}, w)$ is of type $\mathscr{F}_n$, for all n. It follows that $\mathscr{D}(\mathscr{P}, w)$ is of type $\mathscr{F}_\infty$, proving Theorem 4.4.

Acknowledgements

This paper is based on my doctoral dissertation written at Binghamton University (State University of New York). I thank my advisor, Ross Geoghegan, for being a source of constant support and encouragement to me for the past several years. Matt Brin, Victor Cuba, Fernando Guzman, Martin Bridson, and the referee made helpful comments and suggestions for revisions. Thanks are due to Vlad Sergiescu for telling me about the connection between what I'd done and the Baum–Connes Conjecture.

Appendix.

This appendix contains a description of an isomorphism from the diagram group $\mathscr{D}(\mathscr{P}, x)$, where $\mathscr{P} = \langle x \mid x = x^2 \rangle$, to Thompson's group F. Proofs are only sketched, since Guba and Sapir [10] proved that $\mathscr{D}(\mathscr{P}, x) \cong F$ by showing that $\mathscr{D}(\mathscr{P}, x) \cong \langle x_0, x_1, \ldots, x_n, \ldots \mid x_i^{-1} x_j x_i = x_{j+1} (i<j) \rangle$, a well-known presentation for F. Some background from [8] is assumed.

A $(x, *)$-diagram over $\langle x \mid x = x^2 \rangle$ is *positive* if the top label of each cell is x; a $(*, x)$-diagram is *negative* if the top label of each cell is x^2. A positive diagram Δ corresponds to a rooted binary

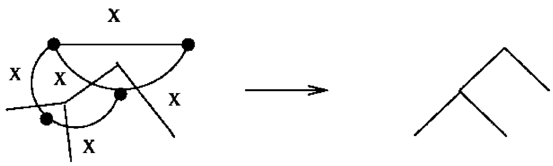

Fig. 4.

tree T_Δ. An example of a positive diagram Δ and its corresponding tree T_Δ are depicted in Fig. 4. The inverse of a negative diagram is positive. If Δ is negative, then $T_\Delta := T_{\Delta^{-1}}$.

Lemma 4.12. *If Δ is a reduced (x,x)-diagram over $\langle x \mid x = x^2 \rangle$, then $\Delta \equiv \Delta_1 \circ \Delta_2$, where Δ_1 is positive and Δ_2 is negative. The diagrams Δ_1 and Δ_2 are, moreover, unique.*

Let Δ be a reduced (x,x)-diagram over $\langle x \mid x = x^2 \rangle$. By the preceding lemma, $\Delta \equiv \Delta_1 \circ \Delta_2$, where Δ_1 is positive and Δ_2 is negative. Let T_{Δ_1} be the tree corresponding to Δ_1 and T_{Δ_2} be the tree corresponding to Δ_2. It is not difficult to see that $(T_{\Delta_1}, T_{\Delta_2})$ is a reduced tree diagram. There is therefore a map $\phi : \mathscr{D}(\mathscr{P}, x) \to F$, where $\phi(\Delta) = (T_{\Delta_1}, T_{\Delta_2})$

This ϕ is an isomorphism. It is easily seen to be bijective, and it is a homeomorphism if one writes functions on the right in the representation of F by PL homeomorphisms. The proof is omitted.

References

[1] M. Bestvina, N. Brady, Morse theory and finiteness properties of groups, Invent. Math. 129 (3) (1997) 445–470.
[2] M.R. Bridson, On the semisimplicity of polyhedral isometries, Proc. Am. Math. Soc. 127 (7) (1999) 2143–2146.
[3] M.R. Bridson, A. Haefliger, Metric Spaces of Non-Positive Curvature, Springer, Berlin, Heidelberg, 1999, xxi+643pp.
[4] M.R. Bridson, Remarks on a talk of Higson, unpublished note.
[5] M.G. Brin, C.C. Squier, Groups of PL homeomorphisms of the real line, Invent. Math. 79 (3) (1985) 485–498.
[6] K.S. Brown, Finiteness properties of groups, J. Pure Appl. Algebra 44 (1–3) (1987) 45–75.
[7] K.S. Brown, R. Geoghegan, An infinite-dimensional torsion-free FP_∞ group, Invent. Math. 77 (2) (1984) 367–381.
[8] J.W. Cannon, W.J. Floyd, W.R. Parry, Introductory notes on Richard Thompson's groups, Enseign. Math. (2) 42 (3–4) (1996) 215–256.
[9] M. Gromov, Hyperbolic groups in: S.M. Gersten (Ed.), Essays in Group Theory, Vol. 8, MSRI Publ., Springer, Berlin, 1987, pp. 75–263.
[10] V.S. Guba, M.V. Sapir, Diagram groups, Mem. Am. Math. Soc. 130 (620) (1997) viii+117.
[11] V.S. Guba, M.V. Sapir, On subgroups of R. Thompson's group F and other diagram groups, Mat. Sb. 190 (8) (1999) 3–60 (in Russian).
[12] N. Higson, G. Kasparov, Operator K-theory for groups which act properly and isometrically on Hilbert space, Electron. Res. Announc. Am. Math. Soc. 3 (1997) 131–142 (electronic).
[13] P. Julg, Travaux de N. Higson et G. Kasparov sur la conjecture de Baum–Connes (French) [The work of N. Higson and G. Kasparov on the Baum–Connes conjecture] Séminaire Bourbaki, Vol. 1997–98; Astérisque 252 (1998) 151–183 (Exp. No. 841, 4).
[14] R. McKenzie, R.J. Thompson, An elementary construction of unsolvable word problems in group theory, in: W.W. Boone, F.B. Cannonito, R.C. Lyndon (Eds.), Word Problems (Conf., Univ. California, Irvine, 1969), Studies in Logic and the Foundations of Mathematics, Vol. 71, North-Holland, Amsterdam, 1973, pp. 457–478.

[15] G.A. Niblo, L.D. Reeves, Groups acting on CAT(0) cube complexes, Geom. Topol. 1 (1997) 7 (electronic).
[16] G.A. Niblo, L.D. Reeves, The geometry of cube complexes and the complexity of their fundamental groups, Topology 37 (3) (1998) 621–633.
[17] E.H. Spanier, Algebraic Topology, McGraw-Hill, New York, Toronto, Ont., London, 1966, xiv+528pp.
[18] M. Stein, Groups of piecewise linear homeomorphisms, Trans. Am. Math. Soc. 332 (2) (1992) 477–514.