

The action of Thompson’s group on a CAT(0) boundary

Daniel Farley

Abstract. For a given locally finite CAT(0) cubical complex X with base vertex $*$, we define the *profile* of a given geodesic ray c issuing from $*$ to be the collection of all hyperplanes (in the sense of Sageev) crossed by c . We give necessary conditions for a collection of hyperplanes to form the profile of a geodesic ray, and conjecture that these conditions are also sufficient.

We show that profiles in diagram and picture complexes can be expressed naturally as infinite pictures (or diagrams), and use this fact to describe the fixed points at infinity of the actions by Thompson’s groups F , T , and V on their respective CAT(0) cubical complexes. In particular, the actions of T and V have no global fixed points. We obtain a partial description of the fixed set of F ; it consists, at least, of an arc c of Tits length $\pi/2$, and any other fixed points of F must have one particular profile, which we describe. We conjecture that all of the fixed points of F lie on the arc c .

Our results are motivated by the problem of determining whether F is amenable.

Mathematics Subject Classification (2000). 20F65, 20F69.

Keywords. Amenability, CAT(0) cubical complex, Thompson’s group, diagram group, space at infinity.

1. Introduction

Thompson’s group F is the group of piecewise linear homeomorphisms $h: [0, 1] \rightarrow [0, 1]$ satisfying:

- (1) the finitely many points at which h is non-differentiable are all dyadic rational numbers, and
- (2) if h is differentiable at x_0 , then $h'(x_0) \in \{2^i \mid i \in \mathbb{Z}\}$.

Thompson also described two other groups, T and V , which are (respectively) the groups of piecewise linear homeomorphisms h of the circle $[0, 1]/(0 = 1)$ and the right-continuous bijections h of $[0, 1)$; in both cases the functions h are required to satisfy (1) and (2). The survey by Cannon, Floyd, and Parry [5] is a useful introduction to all of these groups.

We are interested in the following question:

Question 1.1 ([10]). Is Thompson's group F amenable?

To explain the original interest in Question 1.1, we will need a few definitions. A group G is *elementary amenable* if it is in the smallest class of groups that is closed under extensions and direct limits, and contains finite and abelian groups. A group G is *amenable* if there is a measure $\mu: \mathcal{P}(G) \rightarrow [0, 1]$ ($\mathcal{P}(G)$ is the power set of G) such that: i) μ is finitely additive; ii) μ is left invariant; and iii) $\mu(G) = 1$. We let EG , AG , and NF denote the classes of elementary amenable groups, amenable groups, and groups with no free non-abelian subgroups (respectively). Von Neumann showed that $EG \subseteq AG \subseteq NF$. The problem of determining whether these inclusions are proper was posed by Day [7].

Brin and Squier [4] showed that $F \in NF$. It is proved in [5] that $F \notin EG$. Thus, the existence of F implies that at least one of the inclusions $EG \subseteq AG \subseteq NF$ is proper for finitely presented groups: a positive answer to Question 1.1 shows that $F \in AG - EG$, and a negative answer shows that $F \in NF - AG$. Since at least 1980, when Geoghegan posed Question 1.1, Thompson's group was expected by many to be an example of a finitely presented non-amenable group with no free subgroups.

We know today that $EG \subsetneq AG \subsetneq NF$, and that the inclusions are proper for both finitely generated and finitely presented groups. Grigorchuk found examples of finitely generated groups in $AG - EG$ [11] and, in 1998, finitely presented examples as well [12]. In 1980, Ol'shanskii [15] found finitely generated groups in $NF - AG$. He and Sapir constructed finitely presented groups in $NF - AG$ in 2002 [16].

Although the original reason to consider Question 1.1 is thus obsolete, the problem of determining whether F is amenable is still of great interest, and motivates much of the current work about F .

Here we attempt to resolve Question 1.1 negatively using CAT(0) geometry. Two earlier results are of vital importance in this. First, Adams and Ballmann [1] showed that an amenable group G which acts by isometries on a locally compact CAT(0) space X must either leave a finite-dimensional flat invariant or fix a point at infinity. Second, [8], [9] showed that Thompson's groups F , T , and V act properly, discretely, and by isometries on proper CAT(0) cubical complexes X_F , X_T , and X_V . If F is amenable, it must therefore either leave a flat invariant or fix a point at infinity in X_F . Elementary properties of F (for instance, the fact that $\bigoplus_{n=1}^{\infty} \mathbb{Z} \subseteq F$ [5]) imply that F cannot act properly and freely on a finite-dimensional flat, so we would have a proof that F is non-amenable if we showed that F has no global fixed points in ∂X_F . (Note that the groups T and V are known to be non-amenable, since both are known to contain non-abelian free subgroups.)

Most of the effort in this paper goes into describing the spaces at infinity of the locally finite complexes X_F , X_T , and X_V . (In fact, our methods apply to all diagram groups and picture groups [9].) At this point, some background on X_F , X_T , and X_V is in order; we restrict our remarks to X_F for the sake of simplicity. The constructions in [8], [9] come from the theory of diagram groups, which is due to

Guba and Sapir [13]. Each vertex of X_F is labelled by a semigroup diagram, which is essentially a picture demonstrating how to derive an equality $w_1 = w_2$ between words w_1, w_2 over a semigroup presentation. The group F is itself a diagram group, so every element in F can be represented by a semigroup diagram as well. The action of F on X_F is given by a natural operation on diagrams: if $x \in F$ and $v \in X_F^0$, then $x \cdot v$ is obtained by stacking the pictures x and v , and then “reducing dipoles”. (We refer the reader to Section 3 for more specifics, or to [13] for a complete introduction.)

Our description of ∂X_F is of the same character. We represent regions of ∂X_F as infinite diagrams, which we call *profiles*. The action of F on profiles is determined, as before, by stacking diagrams. As a result, we can largely reduce the problem of finding fixed points in ∂X_F to a much easier algebra problem, which can be handled by a case analysis. Our main theorem is as follows:

Theorem 1.2. *Thompson's group F fixes an arc in the boundary ∂X_F of Tits length $\pi/2$. Any other fixed points on the boundary of F lie in the profile Δ_∞ . Thompson's groups T and V act without global fixed points on their respective boundaries.*

This unfortunately leaves Question 1.1 open.

The problem of finding any remaining fixed points of the action by F appears to be rather delicate. In Section 7, we give some evidence for and against the existence of additional fixed points in ∂X_F .

The paper is organized as follows. In Section 2, we collect various facts about CAT(0) geometry which will be useful in later sections. In Section 3, we briefly sketch the definitions of diagram groups and the cubical complexes, called diagram complexes, on which they act. In Section 4, we describe regions in the space at infinity of diagram complexes using infinite diagrams, and describe the action of a diagram group on this space at infinity. Section 5 contains the main part of the argument, where it is proved that Thompson's group F fixes the profiles $\Delta_L, \Delta_R, \Delta_{L-R}$, and Δ_∞ , each of which can be described by an infinite tree. The groups T and V fix only the profile Δ_∞ . In Section 6, we show that the profiles Δ_{L-R}, Δ_L , and Δ_R represent (respectively) the interior of, the “left” endpoint of, and the “right” endpoint of an arc of length $\pi/2$ in the Tits metric. We show moreover that all points on this arc are fixed by F . Finally, in Section 7 we show that the region at infinity which we call Δ_∞ contains no fixed points of T or V , even though it is fixed as a set. As a result, one has a proof that T and V fix no point at infinity. We also discuss the problem of determining whether F fixes any points in Δ_∞ .

I would like to thank Kim Ruane for suggesting a version of Conjecture 2.9 to me. I also thank the Max Planck Institute for Mathematics, where I wrote most of the paper, for the excellent working conditions it provided. I thank the referee for making helpful comments, and especially for telling me about the work [14] of Dan Guralnik, which led me to alter Conjecture 2.9. Guralnik independently develops a more general theory of the *Roller boundary* of a CAT(0) cubical complex. I believe

that the description of the boundaries ∂X_F , ∂X_T , and ∂X_V using profiles is really a description of their Roller boundaries in different language and notation.

2. Background on CAT(0)spaces

2.1. Basic definitions. We begin by recalling several basic facts about CAT(0) spaces, all of which are taken directly from [2].

A metric space X is *geodesic* if, for any $x_1, x_2 \in X$, there is an isometric embedding $c: [0, d(x_1, x_2)] \rightarrow X$, called a *geodesic*, such that $c(0) = x_1$ and $c(d(x_1, x_2)) = x_2$. We frequently confuse a geodesic with its image. A *geodesic triangle* $\Delta(x, y, z)$ consists of three points $x, y, z \in X$ and choices of geodesics $[x, y]$, $[y, z]$, $[x, z]$ connecting them. Given such a triangle, it is always possible to find points $\bar{x}, \bar{y}, \bar{z}$ in two-dimensional Euclidean space \mathbb{E}^2 such that $d_X(x, y) = d_{\mathbb{E}^2}(\bar{x}, \bar{y})$, $d_X(y, z) = d_{\mathbb{E}^2}(\bar{y}, \bar{z})$, and $d_X(x, z) = d_{\mathbb{E}^2}(\bar{x}, \bar{z})$. The triangle $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ in \mathbb{E}^2 determined by \bar{x}, \bar{y} , and \bar{z} is called a *comparison triangle* for Δ . There is a map $h: \Delta \rightarrow \bar{\Delta}$ which sends sides of Δ isometrically to the corresponding sides of $\bar{\Delta}$. We say that the triangle Δ satisfies the CAT(0) *inequality* if $d_X(a, b) \leq d_{\mathbb{E}^2}(h(a), h(b))$ whenever $a, b \in \Delta$. A geodesic metric space X is CAT(0) if all geodesic triangles in X satisfy the CAT(0) inequality. CAT(0) spaces are contractible, and *uniquely geodesic*, i.e., given any two points x_1, x_2 in a CAT(0) space X , there is a unique geodesic connecting x_1 to x_2 .

If X is an arbitrary metric space, and $c: [0, a] \rightarrow X, c': [0, a'] \rightarrow X$ are geodesic segments satisfying $c(0) = c'(0)$, then we define the *Alexandrov angle* $\angle(c, c')$ as follows:

$$\angle(c, c') := \lim_{\epsilon \rightarrow 0} \sup_{0 < t, t' < \epsilon} \bar{\angle}_{c(0)}(c(t), c'(t')).$$

Here $\bar{\angle}_{c(0)}(c(t), c'(t'))$ is the angle at $\overline{c(0)}$ in the comparison triangle $\bar{\Delta}$ for $\Delta(c(0), c(t), c'(t'))$. Given three points x, y, z in a CAT(0) space X , we let $\angle_y(x, z)$ denote the Alexandrov angle between the (unique) geodesics $[y, x]$ and $[y, z]$.

The CAT(0) inequality can also be expressed in terms of the Alexandrov angle. If Δ is a geodesic triangle in the metric space X , then Δ satisfies the CAT(0) inequality if and only if each Alexandrov angle in Δ measures less than the corresponding angle in the comparison triangle $\bar{\Delta}$. We say that a geodesic metric space X is CAT(0) if every geodesic triangle in X satisfies this version of the CAT(0) inequality. Bridson and Haefliger [2] show that this definition of CAT(0) spaces is equivalent to the earlier one.

A complete CAT(0) space X has a natural space at infinity ∂X , which we now define. Two geodesic rays $c, c': [0, \infty) \rightarrow X$ are said to be *asymptotic* if there exists a constant K such that $d(c(t), c'(t)) \leq K$ for all $t \geq 0$. The set ∂X of *boundary points* of X (or *points at infinity*) is the set of equivalence classes of geodesic rays,

where two geodesic rays are equivalent if and only if they are asymptotic. In practice, we will always use a basepointed version of this construction. Fix a point $x \in X$. We define ∂X to be the set of geodesic rays $c: [0, \infty) \rightarrow X$ issuing from x , i.e., satisfying $c(0) = x$. These two definitions of ∂X are equivalent in a complete CAT(0) space by the following proposition:

Proposition 2.1 ([2]). *If X is a complete CAT(0) space and $c: [0, \infty) \rightarrow X$ is a geodesic ray issuing from x , then for every point $x' \in X$ there is a unique geodesic ray c' which issues from x' and is asymptotic to c .*

If a group G acts by isometries on the CAT(0) space X , then it is clear that there is an induced action on ∂X , if we regard the latter as the collection of equivalence classes of geodesic rays in X . If we use the basepointed version of the construction, then the action $*$ can be described as follows: Let $c \in \partial X$; i.e., $c: [0, \infty) \rightarrow X$ is a geodesic ray and $c(0) = x$. For an isometry $g \in G$, $g * c$ is the unique geodesic ray issuing from x and asymptotic to the left-translate $g \cdot c$ of c .

2.2. Convexity in CAT(0) spaces. A subset C of a CAT(0) space X is *convex* if, given any two points $x_1, x_2 \in C$, the (unique) geodesic segment $[x_1, x_2]$ is contained in C . A function $f: X \rightarrow \mathbb{R}$ on a geodesic metric space is *convex* if, for any geodesic $c: I \rightarrow X$, the composition $f \circ c$ is convex in the ordinary sense, i.e., if, for any $t, t' \in I$ and $s \in [0, 1]$,

$$(f \circ c)((1 - s)t + st') \leq (1 - s)(f \circ c)(t) + s(f \circ c)(t').$$

Bridson and Haefliger [2] show that there is a natural projection $\pi_C: X \rightarrow C$ defined whenever X is a complete CAT(0) space and C is a closed convex subspace. We collect some basic properties of π_C here.

Proposition 2.2. *Let C be a closed, convex subspace of a complete CAT(0) space X .*

(1) *Let $d_C: X \rightarrow \mathbb{R}$ be defined by the rule*

$$d_C(x) = \inf_{y \in C} d(x, y).$$

The function d_C is convex ([2]).

(2) *For any $x \in X$, there is a unique point $\pi_C(x) \in C$ such that $d(x, \pi_C(x)) = d_C(x)$. The function $\pi_C: X \rightarrow C$ does not increase distances ([2]).*

(3) *If $x \in X - C$ and $y \in C - \{\pi(x)\}$, then $\angle_{\pi(x)}(x, y) \geq \pi/2$ ([2]).*

(4) *Fix $x \in X - C$. If $y \in C$ satisfies $\angle_y(x, y') \geq \pi/2$ for all $y' \in C - \{y\}$, then $y = \pi_C(x)$.*

Proof. (4) Assume that y satisfies the above condition and $y \neq \pi(x)$. Consider the comparison triangle $\bar{\Delta}(\bar{x}, \bar{y}, \bar{\pi(x)})$ for the geodesic triangle $\Delta(x, y, \pi(x))$ in X . Since the comparison triangle is non-degenerate by our assumptions, at least one of the comparison angles $\bar{\angle}_{\bar{y}}(\bar{x}, \bar{\pi(x)})$, $\bar{\angle}_{\bar{\pi(x)}}(\bar{x}, \bar{y})$ measures less than $\pi/2$. Our assumptions and the CAT(0) inequality imply that $\bar{\angle}_{\bar{\pi(x)}}(\bar{x}, \bar{y}) < \pi/2$. By the CAT(0) inequality, $\angle_{\pi(x)}(x, y) < \pi/2$ as well. This violates (3). \square

We say that a geodesic ray $c: [0, \infty) \rightarrow X$ crosses a closed, convex subset C if $(\text{Im } c) \cap C$ is a non-empty, compact interval and $(\text{Im } c) - C$ is disconnected.

Lemma 2.3. *Let X be a CAT(0) space.*

- (1) *Suppose that C is a closed convex subset of X , $c: [0, \infty) \rightarrow X$ is a geodesic ray which crosses C , and $(\text{Im } c) \cap C = c([t_1, t_2])$. The function $d_C \circ c$ is strictly monotonically increasing on $[t_2, \infty)$, and $(d_C \circ c)(t) \rightarrow \infty$ as $t \rightarrow \infty$.*
- (2) *If $c, c': [0, \infty) \rightarrow X$ are two asymptotic geodesic rays in X , then the function $d(c(_), c'(_)): [0, \infty) \rightarrow [0, \infty)$ is non-increasing.*

Proof. Both parts are standard exercises using basic properties of convex functions. Part (1) follows from the fact that $d_C \circ c: [0, \infty) \rightarrow [0, \infty)$ is convex, $(d_C \circ c)(t_2) = 0$, and $(d_C \circ c)(t') > 0$ for some $t' > t_2$. Part (2) follows from the fact that the function $d(c(_), c'(_)): [0, \infty) \rightarrow [0, \infty)$ is convex and bounded (see [2], p. 261). \square

2.3. CAT(0) cubical complexes. We take the following definition of a cubical complex from [2]:

Definition 2.4 ([2], p. 112). A cubical complex K is the quotient of a disjoint union of cubes $X = \coprod_{\Lambda} I^{n_\lambda}$ by an equivalence relation \sim . The restrictions $p_\lambda: I^{n_\lambda} \rightarrow K$ of the natural projection $p: X \rightarrow K = X/\sim$ 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 $T_\lambda \subseteq I^{n_\lambda}$ onto a face $T_{\lambda'} \subseteq I^{n_{\lambda'}}$ such that $p_\lambda(x) = p_{\lambda'}(x')$ if and only if $x' = h_{\lambda, \lambda'}(x)$.

Let x and y be points in X , and let $l(c)$ denote the length of a path c .

$$d_\ell(x, y) = \inf\{l(c) \mid c(0) = x; c(1) = y\}.$$

The function $d_\ell: X \times X \rightarrow [0, \infty]$ defines a metric on any cubical complex, called the *length metric* [2]. A well-known theorem due to Gromov [2] says that the length metric on a cubical complex X is a CAT(0) metric if and only if X satisfies the link condition. We avoid recounting the precise statement here, but will work exclusively with CAT(0) cubical complexes from now on.

Let X be a complete CAT(0) cubical complex. Following [17], define a relation \sim on edges of X , such that $e_1 \sim e_2$ if and only if e_1 and e_2 are opposite sides of a square (2-cell) in X . We will sometimes call this relation *simple square equivalence*, although it is not an equivalence relation. The transitive, reflexive closure of this relation, also denoted \sim , is called *square equivalence*. It is clear that square equivalence is an equivalence relation.

A *combinatorial hyperplane* in X is an equivalence class of edges under \sim . One obtains a geometric hyperplane H as follows: let M_e be the set of all midpoints of all edges square equivalent to e . If C is an arbitrary cube, define $H \cap C$ to be the convex hull of $M_e \cap C$ in the cube C . This description determines H .

We now collect some basic properties of CAT(0) cubical complexes.

Theorem 2.5 ([17]). *Let X be a CAT(0) cubical complex.*

- (1) *If J is a geometric hyperplane in X , then J does not intersect itself and partitions X into two convex components.*
- (2) *If J_1, \dots, J_k are a collection of geometric hyperplanes in X such that $J_m \cap J_n \neq \emptyset$ for all m, n , then $\bigcap J_i \neq \emptyset$.*
- (3) *If x and y are vertices in X connected by a geodesic edge-path p of length n , then p crosses n distinct hyperplanes J_1, \dots, J_n , and these are precisely the hyperplanes which separate x from y . In particular, any other geodesic edge-path p' connecting x to y must cross precisely the same hyperplanes.*
- (4) *Each geometric hyperplane J is itself a CAT(0) cubical complex.*

The following lemma will be useful in Section 2.4.

Lemma 2.6. *Let X be a locally finite CAT(0) cubical complex.*

- (1) *The closed $\frac{1}{2}$ -neighborhood of a hyperplane H in X factors isometrically as $H \times [0, 1]$.*
- (2) *Let $c: [0, \infty) \rightarrow X$ be a geodesic ray issuing from a vertex $*$ of X . If $c([0, \infty)) \cap C \neq \emptyset$ for some open cube C in X , then c crosses every hyperplane passing through C .*
- (3) *If H is a hyperplane in X , d_H is non-constant on an open cell C in X , and $H \cap C = \emptyset$, then there exists some hyperplane H_1 passing through C such that $H_1 \cap H = \emptyset$.*
- (4) *Let H_1 and H_2 be hyperplanes in the CAT(0) cubical complex X , and, for $i = 1, 2$, let H_i^+, H_i^- be the two open, convex components of $X - H_i$. If the intersections $H_1^+ \cap H_2^-, H_1^+ \cap H_2^+, H_1^- \cap H_2^-,$ and $H_1^- \cap H_2^+$ are all non-empty, then $H_1 \cap H_2$ is also non-empty.*

Proof. (1) This is a consequence of the proof for Theorem 4.10 (p. 611) of [17].

(2) Suppose that $c(t_1) \in C$; let H be any hyperplane passing through C . Thus $c(t_1) \in H \times (0, 1)$; say $c(t_1) \in H \times \{t'\}$. Let s_1, s_2 be arbitrary numbers such that $s_1 < t' < s_2$ and $\frac{1}{2} \in (s_1, s_2)$. The hyperplanes $H \times \{s_i\} = H_{s_i}$ ($i = 1, 2$) separate X into three distinct connected components, one of which is $H \times (s_1, s_2)$. Note that $H \times [s_1, s_2]$ does not contain $*$, so c must therefore cross either H_{s_1} or H_{s_2} . If we assume, without loss of generality, that c crosses H_{s_1} , then it must be that $d_{H_{s_1}}(c(t)) \rightarrow \infty$ monotonically on $[t_1, \infty)$, for appropriate t_1 , by Lemma 2.3. The function $d_{H_{s_1}}$ is bounded on $H \times [s_1, s_2]$, so the geodesic ray c eventually leaves $H \times [s_1, s_2]$, and it cannot cross H_{s_1} a second time, due to the monotonicity of $d_{H_{s_1}} \circ c$ on $[t_1, \infty)$. It follows that c crosses H_{s_2} , and thus also $H_{1/2} = H$.

(3) We prove the contrapositive. Suppose that H is a hyperplane, C is an open cell such that $H \cap C = \emptyset$, and every hyperplane H' passing through C satisfies $H' \cap H \neq \emptyset$. We wish to show that d_H is constant on C .

Identify C with $(0, 1)^n$ and fix a factor of $(0, 1)^n$ (the last one, without loss of generality). There is a hyperplane H' such that $H' \cap C = (0, 1)^{n-1} \times \{\frac{1}{2}\}$. Let $x_1, x_2 \in C$ be two points in C which differ only in the last coordinate, say $x_1 = (c_1, c_2, \dots, c_n)$ and $x_2 = (c_1, c_2, c_3, \dots, \hat{c}_n)$. We regard these as points in $H'_{c_n} := H' \times \{c_n\}$ and $H'_{\hat{c}_n} := H' \times \{\hat{c}_n\}$, respectively. Let $x = (c_1, c_2, \dots, c_{n-1}, \frac{1}{2})$ be in H' , which we identify with $H' \times \{\frac{1}{2}\}$.

We consider the projection $\pi_{H \cap H'}: H' \rightarrow H \cap H'$. Suppose that $\pi_{H \cap H'}(x) \in C'$ where $C' = (0, 1)^m$ is an open cube (of dimension at least 2, since H and H' both pass through C' , and $H \neq H'$). We make the identifications $H' \cap C' = \{(d_1, \dots, d_m) \in C' \mid d_m = 1/2\}$ and $H \cap C' = \{(d_1, \dots, d_m) \in C' \mid d_{m-1} = 1/2\}$.

Suppose that $\pi_{H \cap H'}(x) = (d'_1, d'_2, \dots, d'_{m-2}, 1/2, 1/2)$. The requirement that $\angle_{\pi_{H \cap H'}(x)}(x, y) \geq \pi/2$ for all $y \neq \pi_{H \cap H'}(x)$ in $H \cap H'$ guarantees that $[x, \pi_{H \cap H'}(x)] \cap C' \subseteq \{d'_1\} \times \dots \times \{d'_{m-2}\} \times (0, 1) \times \{1/2\}$ (note: the last coordinate must be $1/2$ since $[x, \pi_{H \cap H'}(x)] \subseteq H'$). But it follows from this that $\angle_{\pi_{H \cap H'}(x)}(x, y) \geq \pi/2$ for all $y \neq \pi_{H \cap H'}(x)$ in H . That is: $\pi_{H \cap H'}(x) = \pi_H(x)$, by Proposition 2.2 (4), where $\pi_{H \cap H'}: H' \rightarrow H \cap H'$ and $\pi_H: X \rightarrow H$ are the projections.

Therefore the geodesic segment $[x, \pi_H(x)] (= [x, \pi_{H \cap H'}(x)])$ is a subset of H' . Now we consider the geodesic segments $[x, \pi_H(x)] \times \{c_n\} \subseteq H' \times \{c_n\}$ and $[x, \pi_H(x)] \times \{\hat{c}_n\} \subseteq H' \times \{\hat{c}_n\}$. These run parallel to $[x, \pi_H(x)]$, and meet H perpendicularly for similar reasons. It follows that $[x_1, \pi_H(x_1)] = [x, \pi_H(x)] \times \{c_n\}$ and $[x_2, \pi_H(x_2)] = [x, \pi_H(x)] \times \{\hat{c}_n\}$.

This implies that $d_H(x_1) = d_H(x) = d_H(x_2)$, which implies that the value of d_H is independent of the last coordinate. Now we can argue coordinate by coordinate to conclude that d_H is constant on C .

(4) Assume that the four intersections in the hypothesis are all non-empty. If we assume also that $H_1 \cap H_2 = \emptyset$, then it follows that $\{H_1^+ \cup H_2^+, H_1^- \cup H_2^-\}$ is an open

cover of X . Now each of the half-spaces H_1^+, H_1^-, H_2^+ , and H_2^- is a convex subspace of a CAT(0) space, and therefore CAT(0) itself. It follows that each is contractible. The same reasoning also applies to the four intersections in the hypothesis: each is CAT(0), and therefore contractible.

It then follows that each of the sets $X^+ = H_1^+ \cup H_2^+$, $X^- = H_1^- \cup H_2^-$ is simply connected, since each is the union of two open contractible sets which intersect in an open contractible set. The intersection $X^+ \cap X^-$ is the union of two disjoint open contractible sets: $H_1^+ \cap H_2^-$ and $H_2^+ \cap H_1^-$. Let c be an arc contained in X^+ connecting $H_1^+ \cap H_2^-$ to $H_2^+ \cap H_1^-$ and meeting each in an open segment.

We apply van Kampen's theorem to the pieces $X^- \cup c$ and X^+ . The first piece $X^- \cup c$ satisfies $\pi_1(X^- \cup c) \cong \mathbb{Z}$, while the second is simply connected. The intersection of these two pieces is the simply connected set $(H_1^+ \cap H_2^-) \cup (H_2^+ \cap H_1^-) \cup c$. It follows that $\pi_1(X^- \cup X^+) = \pi_1(X)$ is isomorphic to \mathbb{Z} . The space X is CAT(0), however, and therefore contractible. We have a contradiction. \square

2.4. Profiles of geodesic rays in CAT(0) cubical complexes. Suppose now that X is locally finite, and let $*$ be a vertex, which will serve as a basepoint. If H is any hyperplane in X , let H^+ , the *positive half-space* determined by H , be the open complementary component of $X - H$ that does not contain $*$; we let H^- be the other open complementary component of $X - H$. If H_1 and H_2 are hyperplanes in X , write $H_1^+ \leq H_2^+$ if $H_2^+ \subseteq H_1^+$. Clearly \leq is a partial order on positive half-spaces. We also regard \leq as a partial order on hyperplanes, writing $H_1 \leq H_2$ if $H_1^+ \leq H_2^+$.

For any geodesic ray c in X issuing from $*$, define $P(c)$, the *profile of c* , to be the collection of all positive half-spaces H^+ such that H is crossed by c .

Proposition 2.7. *If $c : [0, \infty) \rightarrow X$ is a geodesic ray issuing from $*$, then $P(c)$ is non-empty and satisfies:*

- (1) *for any finite subset $\{H_1^+, \dots, H_n^+\} \subseteq P(c)$, $H_1^+ \cap \dots \cap H_n^+ \neq \emptyset$;*
- (2) *the partially ordered set $(P(c), \leq)$ has no maximal elements, and*
- (3) *if $H_1^+ \in P(c)$ and $H_2^+ \leq H_1^+$, then $H_2^+ \in P(c)$.*

Proof. If $H_1^+, H_2^+, \dots, H_n^+$ are contained in $P(c)$, then there exist real numbers $t_1, t_2, \dots, t_n > 0$ such that $c([t_i, \infty)) \subset H_i^+$, by Lemma 2.3. If t is the largest number in $\{t_1, \dots, t_n\}$, then clearly $c([t, \infty)) \subseteq H_1^+ \cap H_2^+ \cap \dots \cap H_n^+$. This proves that property (1) holds for $P(c)$.

It is obvious that property (3) is true of $P(c)$.

Suppose that $H^+ \in P(c)$ is a maximal element, and let $c([t, \infty)) \subseteq H^+$. Since $d_H(c(t')) \rightarrow \infty$ as $t' \rightarrow \infty$ by Lemma 2.3, we can choose t so that $d_H(c(t)) > 1/2$. We consider the collection \mathcal{C} of all open cells C such that $c([t, \infty)) \cap C \neq \emptyset$. Note that if $\hat{C} \cap H \neq \emptyset$, then any point in $c([t, \infty)) \cap \hat{C}$ is at most $1/2$ -distant from H , so that every cell C in \mathcal{C} satisfies $C \cap H = \emptyset$.

Let $C \in \mathcal{C}$. If d_H is non-constant on C , then, by Lemma 2.6(3), there is a hyperplane H' passing through C such that $H' \cap H = \emptyset$. We claim that this implies $H^+ < (H')^+$. If $H^+ \not< (H')^+$, i.e., if $(H')^+ \not\subseteq H^+$, then $(H')^+ \cap H^- \neq \emptyset$. We also know that $* \in H^- \cap (H')^-$, $H^+ \cap (H')^- \neq \emptyset$ (since $C \subseteq H^+$ and $C \cap (H')^- \neq \emptyset$), and $H^+ \cap (H')^+ \neq \emptyset$ (since the geodesic ray must cross H' by Lemma 2.6(2)). Now it follows from Lemma 2.6(4) that $H' \cap H \neq \emptyset$, which is a contradiction. This proves the claim. Now $H^+ < (H')^+$ and c crosses H' by Lemma 2.6(2), which contradicts the maximality of H^+ .

It follows that d_H is constant on all of the cells C in \mathcal{C} . This contradicts the fact that $d_H \circ c$ is strictly monotonically increasing on $[t, \infty)$. It follows that property (2) holds.

It is obvious that $P(c)$ is non-empty. □

From now on, we call a collection of positive half-spaces \mathcal{H} a *profile* if it is non-empty and satisfies properties (1)–(3) in Proposition 2.7. In an earlier version of this paper, I asked if every profile \mathcal{H} in this sense is realized by a geodesic ray, i.e., if there is some geodesic ray c issuing from $*$ such that $\text{Im } c \cap H^+ \neq \emptyset$ if and only if $H^+ \in \mathcal{H}$. The referee pointed out that Guralnik answered this question negatively in [14]. We include his example here.

Example 2.8 ([14], Example 6.11). Consider the space

$$X = \bigcup_{n \in \mathbb{N}} [n - 1, n] \times [n^2, \infty),$$

given the standard cubulation. Let $(0, 0)$ be the basepoint. We note that there are three profiles: one consists of all half-spaces H_i^+ ($i \in \mathbb{N}$) having the form $H_i^+ = \{(x, y) \in X \mid y > i - 1/2\}$; another consists of all half-spaces \hat{H}_i^+ ($i \in \mathbb{N}$) having the form $\hat{H}_i^+ = \{(x, y) \in X \mid x > i - 1/2\}$; a third is the collection of all half-spaces H_i^+ and \hat{H}_i^+ . It is not difficult to check that only the first of these profiles can be realized by a geodesic ray.

It is natural to ask if every profile can be realized in a restricted class of cubical complexes, such as those admitting cocompact actions. Guralnik [14] considers some hypotheses under which one can guarantee a positive answer after his Example 6.11. His Question 6.1 suggests that it is not known whether all profiles are realized, even if one assumes that X has a cocompact action.

The following conjecture was originally suggested to me by Kim Ruane.

Conjecture 2.9. *Let X be a locally finite CAT(0) cubical complex with base vertex $*$. The collection of geodesic rays c having a given, fixed profile \mathcal{H} forms a subset of ∂X of diameter less than or equal to $\pi/2$, where the distance in question is the angle metric (see [2]).*

We note that Guralnik (in [14]) has shown that Conjecture 2.9 is true, provided we replace the bound $\pi/2$ with π .

Conjecture 2.9 implies, in particular, that each profile represents a contractible subset of ∂X , since ∂X is a CAT(1) space with respect to the angular metric, and sets of diameter less than π are contractible in CAT(1) spaces (Proposition 1.4(4) in [2]). If Conjecture 2.9 is true, then the description of ∂X in terms of profiles may therefore give a useful homotopical view of the space at infinity, especially if profiles have a convenient description. We will give such a description of profiles in diagram complexes later in Section 4.

3. Diagram groups

3.1. Basic definitions. If A is a set (alphabet), then the *free semigroup on A* , denoted A^+ , is the collection of all positive, non-empty words in A , with the operation of concatenation. Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation. Thus, Σ is an alphabet and $\mathcal{R} \subseteq \Sigma^+ \times \Sigma^+$ is a collection of equalities between elements of Σ^+ . We will follow the convention of [8] and impose the additional assumption that no relation of the form (w, w) occurs in \mathcal{R} .

We now define pictures over \mathcal{P} . Begin with a *frame* $\partial([0, 1]^2)$, a finite, possibly empty, collection \mathcal{T} of *transistors*, each homeomorphic to $[0, 1]^2$, and a finite, non-empty collection \mathcal{W} of *wires*, each homeomorphic to $[0, 1]$. The frame and transistors all have well-defined top, bottom, left, and right sides, which are the open sides parallel to the coordinate axes, and do not include the corners. A wire has well-defined initial and terminal points (i.e., 0 and 1, respectively). A *picture over \mathcal{P}* , denoted Δ , is a quotient of

$$\partial([0, 1]^2) \coprod \left(\coprod_{T \in \mathcal{T}} T \right) \coprod \left(\coprod_{w \in \mathcal{W}} w \right)$$

(for a choice of sets \mathcal{T} and \mathcal{W}) by an equivalence relation \sim , together with a labelling function $\ell: \mathcal{W} \rightarrow \Sigma$, satisfying:

- (1) The initial point of any given wire is attached either to the bottom of a transistor, or to the top of the frame. The terminal point of any given wire is attached either to the bottom of the frame, or to the top of some transistor. If w is a wire and T is a transistor, then $w \cap T \subseteq \Delta$ is either empty or a singleton set.
- (2) Let T_1 and T_2 be transistors. Write $T_1 < T_2$ if there is some wire w such that the initial point of w is attached to the bottom of T_1 and the terminal point of w is attached to the top of T_2 . Let $<$ also denote the transitive closure of the above relation. The relation $<$ is required to be a strict partial order.
- (3) The equivalence classes of \sim are either singleton sets or consist of exactly two points, exactly one of which is an endpoint of a wire. In other words, the only

identifications in

$$\partial([0, 1]^2) \coprod \left(\coprod_{T \in \mathcal{T}} T \right) \coprod \left(\coprod_{w \in \mathcal{W}} w \right)$$

are generated by the attaching maps of the wires, and no two wires have points in common. The endpoints of the wires are called *contacts*.

- (4) Suppose the top of the transistor T meets the wires $w_{i_1}, w_{i_2}, \dots, w_{i_m}$, reading from the left side of T to the right. Suppose that the bottom of the transistor T meets the wires w_{j_1}, \dots, w_{j_n} , again reading from left to right. The *top label* of T , denoted L_T , is

$$\ell(w_{i_1})\ell(w_{i_2}) \dots \ell(w_{i_m});$$

the *bottom label* of T , denoted L_B , is

$$\ell(w_{j_1})\ell(w_{j_2}) \dots \ell(w_{j_n}).$$

We require that $(L_T, L_B) \in \mathcal{R}$ or $(L_B, L_T) \in \mathcal{R}$.

We can define the top and bottom labels of the frame just as we did for a transistor T . If the top label of the frame is w_1 and the bottom label is w_2 , then Δ is a (w_1, w_2) -picture over \mathcal{P} . We say that Δ is a $(w, *)$ -picture if the top label of Δ is w , and the bottom label is arbitrary.

Two pictures Δ_1 and Δ_2 are *isomorphic*, $\Delta_1 \equiv \Delta_2$, if there is a homeomorphism between them which matches labels and preserves the top-bottom- and left-right-orientations on the frame and transistors.

Given a (u, v) -picture Δ_1 and a (v, w) -picture Δ_2 , one can define the *concatenation* $\Delta_1 \circ \Delta_2$, which is the (u, w) -picture obtained by identifying the bottom of the frame for Δ_1 with the top of the frame for Δ_2 by a homeomorphism which matches the endpoints of the wires, and then removing the line segment corresponding to the bottom of Δ_1 in the quotient, while keeping the wires passing through this line segment intact.

Figure 1 illustrates the operation of concatenation in a particular case. All of the semigroup pictures in the figure are pictures over the presentation $\langle x \mid x = x^2 \rangle$. For this reason, we leave off the labels of the wires, since the label of each one is x . Note that on the top left is an (x, x^3) -picture, and on the bottom left is an (x^3, x) -picture.

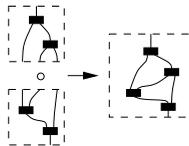


Figure 1. On the left, we have two pictures Δ_1 and Δ_2 (reading from top to bottom); on the right we have the concatenation $\Delta_1 \circ \Delta_2$.

If we denote these pictures Δ_1 and Δ_2 , respectively, then the (x, x) -picture on the right is $\Delta_1 \circ \Delta_2$.

Note that Figure 1 also illustrates our conventions for drawing pictures in the plane. If Δ is a semigroup picture, then a function $\rho: \Delta \rightarrow \mathbb{R}^2$ is a *projection* of Δ if:

- (1) The image of each transistor is a rectangle whose sides are parallel to the coordinate axes. The map ρ takes the top, left, right, and bottom of any given transistor to the corresponding sides in the image.
- (2) The image of the frame is an empty rectangle, and the map ρ is again orientation-preserving, in the sense of (1). The image of ρ is contained inside the image of the frame.
- (3) The image of each wire meets any given horizontal line at most once, and
- (4) The function ρ is an embedding, except possibly at finitely many double points. The inverse image of any double point x is a set of two points on distinct wires w_1 and w_2 . We assume that the images of w_1 and w_2 are transverse at x .

It is rather clear that all of the defining features of a semigroup picture can be recovered from any of its suitably labelled projections. From now on, we will usually confuse a picture with any of its projections without further comment.

Two transistors $T_1 < T_2$ form a *dipole* if the top label of T_1 is identical (as a word in Σ^+) to the bottom label of T_2 , and the bottom contacts of T_1 are paired off by wires in order with the top contacts of T_2 . To *remove a dipole*, delete the transistors T_1 and T_2 and all wires connecting them, and then glue together in order the wires that formed top contacts of T_1 with those that formed bottom contacts of T_2 . The inverse operation is called *inserting a dipole*. Two pictures are *equal modulo dipoles*, $\Delta_1 = \Delta_2$, if one can be obtained from the other by repeatedly inserting and removing dipoles. A picture is called *reduced* if it contains no dipoles. Any equivalence class modulo dipoles contains a unique reduced picture [9], [13].

In Figure 2 we have two $(acbd, abab)$ -pictures over the presentation $\mathcal{P} = \langle a, b, c, d \mid ab = cd, cb = bc, ab = ba \rangle$. In the left picture, we have circled

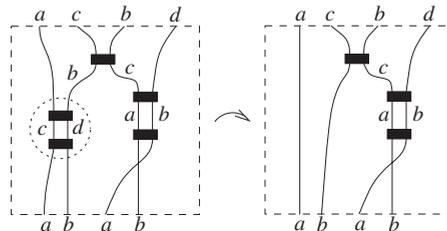


Figure 2. The circled transistors in the left half of the figure form a dipole; on the right, we have the result of removing this dipole.

two transistors which form a dipole. If we remove this dipole, we arrive at the picture on the right. Notice that the two right-most transistors in the right half of the figure do not form a dipole: the top label of the top transistor is cd , but the bottom label of the bottom transistor is ba .

For a fixed word $w \in \Sigma^+$, the set of all (w, w) -pictures over \mathcal{P} , modulo dipoles, forms a group $\mathcal{D}_b(\mathcal{P}, w)$ under the operation of concatenation. We will follow [13] and call $\mathcal{D}_b(\mathcal{P}, w)$ the *braided diagram group* over \mathcal{P} , based at w . (Warning: the word “braided” is rather unfortunate. In fact, as the above definition shows, we do not care about any possible braiding of the wires, since equivalence between pictures does not depend on any embedding into an ambient space. Moreover, there is now a growing literature (see for example [3]) on a braided version of Thompson’s group V , which is something quite different from the older group V which we consider here. Nevertheless, there seems to be no better term.) A picture Δ is *planar* if there is a projection $\rho: \Delta \rightarrow \mathbb{R}^2$ which is also an embedding. The set of all planar (w, w) -pictures over \mathcal{P} , modulo dipoles, forms a group $\mathcal{D}(\mathcal{P}, w)$, which we will call the *diagram group* over \mathcal{P} , based at w . Annular pictures can be defined as follows. Suppose that Δ is a picture, and let Δ' be the space obtained from Δ by removing the sides of the frame. We say that Δ is *annular* if there is an orientation-preserving immersion of Δ' into $A = \{(x, y) \in \mathbb{R}^2 \mid 1 \leq x^2 + y^2 \leq 4\}$ such that: i) the top of the frame for Δ' is wrapped around the circle $x^2 + y^2 = 1$ once in the counterclockwise direction. The initial and terminal points of the top are both mapped to $(1, 0) \in A$; ii) the bottom of the frame for Δ' is wrapped around the circle $x^2 + y^2 = 4$ once in the counterclockwise direction. The initial and terminal points of the bottom are both mapped to $(2, 0) \in A$; iii) the only double points of ρ are $(1, 0)$ and $(2, 0)$; ρ is an embedding otherwise. The set of all annular pictures over \mathcal{P} is a group $\mathcal{D}_a(\mathcal{P}, w)$, called the *annular diagram group* over \mathcal{P} , based at w .

Three groups are of special interest to us. Let $\mathcal{P} = \langle x \mid x = x^2 \rangle$. The groups $\mathcal{D}(\mathcal{P}, x)$, $\mathcal{D}_a(\mathcal{P}, x)$, and $\mathcal{D}_b(\mathcal{P}, x)$ are, respectively, Thompson’s groups F , T , and V . The original observation that $\mathcal{D}(\mathcal{P}, x) \cong F$ was due to Victor Guba; Guba and Sapir (in [13]) sketched the theory of annular and braided diagram groups expressly for the purpose of bringing their techniques to bear on the study of T and V . Section 6 of [9] describes an isomorphism between the groups F , T , and V , and the corresponding diagram groups.

3.2. Diagram complexes. If G is a diagram group of the standard, annular, or braided variety, then a theorem of [9] (see also [8]) says that G acts properly by isometries on a CAT(0) cubical complex. We briefly describe the construction of the cubical complex in this subsection.

Fix a braided diagram group $\mathcal{D}_b(\mathcal{P}, w)$. We define a complex $\tilde{\mathcal{K}}_b(\mathcal{P}, w)$, called the *diagram complex* for $\mathcal{D}_b(\mathcal{P}, w)$, as follows. A vertex $v \in \tilde{\mathcal{K}}_b(\mathcal{P}, w)^0$ is an equivalence class \sim of reduced braided $(w, *)$ -pictures, where $\Delta_1 \sim \Delta_2$ if and only

if there is some braided permutation picture Ψ , such that $\Delta_1 \circ \Psi = \Delta_2$. Here a *permutation picture* is one with no transistors. It is convenient to depict a vertex as a $(w, *)$ -picture in which all wires which would ordinarily be connected to the bottom of the frame have been cut, as in Figure 3.

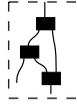


Figure 3. This is a vertex in the cubical complex $\tilde{K}_b(\mathcal{P}, x)$, where $\mathcal{P} = \langle x \mid x = x^2 \rangle$.

An n -dimensional cube in $\tilde{K}_b(\mathcal{P}, w)$ is denoted by a reduced braided $(w, *)$ -picture Δ in which all of the bottom wires have been cut (as above), and n of the maximal transistors of Δ have been drawn as white. The picture Figure 4 a) denotes a 2-cube, for example.

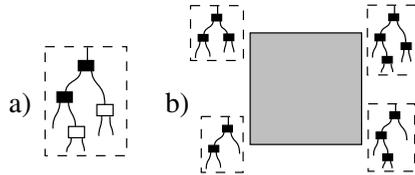


Figure 4. a) This notation describes a cube in the complex $\tilde{K}_b(\mathcal{P}, x)$. b) This is the labelled cube in $\tilde{K}_b(\mathcal{P}, x)$ denoted by the picture in a).

If we arbitrarily number the white transistors $1, 2, \dots, n$, then there is a natural way to label the vertices of an n -cube $[0, 1]^n$, corresponding to this numbering of Δ . Namely, if $(a_1, \dots, a_n) \in \{0, 1\}^n$ label (a_1, \dots, a_n) by the picture $\Delta_{(a_1, \dots, a_n)}$, where the i th transistor is left off if $a_i = 0$ and the i th transistor is filled in if $a_i = 1$. For instance, Figure 4 b) shows how to label the corners of $[0, 1]^2$ if Δ is as in Figure 4 a) and the white transistors are numbered from left to right.

If we let Δ vary over all possible isomorphism classes of cube representatives (where isomorphisms send white transistors to white transistors), and, for each Δ , choose as above a labelling of $[0, 1]^n$ for the appropriate n , then $\tilde{K}_b(\mathcal{P}, w)$ is the quotient of the resulting labelled cubes by the equivalence relation which identifies the cubes along faces with the same labels. It is proved in [9] that $\tilde{K}_b(\mathcal{P}, w)$ is a proper CAT(0) cubical complex if \mathcal{P} is a finite presentation, and that $\mathcal{D}_b(\mathcal{P}, w)$ acts properly and cellularly on $\tilde{K}_b(\mathcal{P}, w)$. The action is usually not cocompact, and, in particular, is not for any of the groups F , T , and V .

We note that entirely similar statements are true for ordinary and annular diagram groups. It is only necessary to replace pictures with planar pictures and annular pictures (respectively) in the above discussion to get the descriptions of $\tilde{K}(\mathcal{P}, w)$ and $\tilde{K}_a(\mathcal{P}, w)$, respectively.

Lastly, we recall a useful partial order on vertices. If $[\Delta_1]$ and $[\Delta_2]$ are vertices in a diagram complex, we write $[\Delta_1] \leq [\Delta_2]$ if there exists some picture θ such that $\Delta_1 \circ \theta \equiv \Delta_2$. Note that this means $\Delta_1 \circ \theta$ and Δ_2 are isomorphic before reducing dipoles. It is not difficult to see that \leq is a well-defined partial order.

Suppose $\mathcal{T}' \subseteq \mathcal{T}_\Delta$, where \mathcal{T}_Δ is the collection of transistors in a picture Δ . We say that \mathcal{T}' is an *initial subset* of \mathcal{T}_Δ if whenever $T_1 < T_2$ and $T_2 \in \mathcal{T}'$, then $T_1 \in \mathcal{T}'$ also. We reproduce a lemma from [9].

Lemma 3.1 ([9]). *Let Δ be a vertex, and let \mathcal{T}_Δ be its set of transistors. There is a one-to-one correspondence ψ between initial subsets of \mathcal{T}_Δ and vertices Δ_1 satisfying $\Delta_1 \leq \Delta$. The function ψ is order-preserving and has an order-preserving inverse, i.e., the initial subsets $\mathcal{T}', \mathcal{T}''$ satisfy $\mathcal{T}' \subseteq \mathcal{T}''$ if and only if $\psi(\mathcal{T}') \leq \psi(\mathcal{T}'')$.*

The map ψ in the above lemma is easy to define: if \mathcal{T}' is an initial subset of transistors, then $\psi(\mathcal{T}')$ is obtained by removing all transistors in $\mathcal{T}_\Delta - \mathcal{T}'$, along with all of their bottom wires. The result is easily seen to be a vertex. The argument that the map ψ is injective can be extended to prove that the automorphism group of a diagram is trivial, at least combinatorially speaking. That is, if $\phi: \Delta \rightarrow \Delta$ is an isomorphism, then ϕ leaves the frame, each transistor, and each wire invariant, and restricts to a self-homeomorphism of each of these. It therefore follows, for instance, that in a concatenation $\Delta_1 \circ \Delta_2$ of pictures, one can speak of the transistors that were contributed by Δ_i for $i = 1, 2$, and this is a well-defined notion even after reducing dipoles. We shall need this observation in future sections, and use it without further comment.

4. Geodesic profiles in diagram complexes

We now describe profiles in diagram complexes. Our main goals here are, first, to describe a profile as an infinite picture of a certain kind, and then to describe the action on profiles in terms of picture multiplication.

Throughout this section, we use only the complex $\tilde{K}_b(\mathcal{P}, w)$, but the discussion carries over to $\tilde{K}(\mathcal{P}, w)$ and $\tilde{K}_a(\mathcal{P}, w)$ in an obvious way.

4.1. Description of profiles. The first step is to describe hyperplanes in $\tilde{K}_b(\mathcal{P}, w)$. Recall that a combinatorial hyperplane is an equivalence class of 1-cells under the relation \sim of square equivalence. The square equivalence relation is generated by simple square equivalence (also denoted \sim), where two 1-cells e_1, e_2 are simple square equivalent if they are opposite faces of a 2-cell (square).

We describe combinatorial hyperplanes in $\tilde{K}_b(\mathcal{P}, w)$ with help from an example. First, fix a 1-cell in $\tilde{K}_b(\mathcal{P}, w)$, such as the one in Figure 5 a), which we will call Δ .

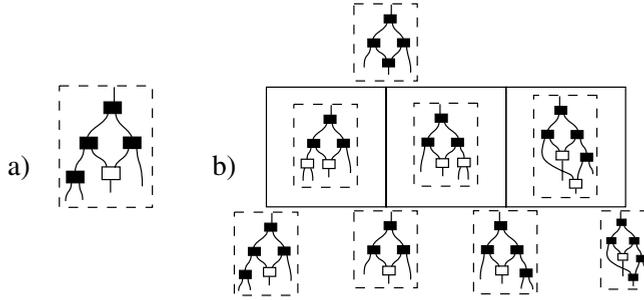


Figure 5. a) An edge Δ in $\tilde{K}_b(\mathcal{P}, x)$, and b) a collection of edges that are square equivalent to Δ .

We consider a small number of 1-cells that are square equivalent to Δ (there are infinitely many such 1-cells for this Δ). These are the vertical edges in Figure 5 b). We denote these edges $\Delta_1, \Delta_2 = \Delta, \Delta_3$, and Δ_4 , reading from left to right. Note the interpretation of simple square equivalence in terms of diagrams: for $i = 1, 2, 3$, $\Delta_i \sim \Delta_{i+1}$ since Δ_i can be obtained from Δ_{i+1} by removing a maximal shaded transistor and all of its bottom wires from Δ_{i+1} , or the reverse, i.e., Δ_{i+1} can be obtained in the same way from Δ_i . This observation is general, and holds true in all of the complexes $\tilde{K}(\mathcal{P}, w)$, $\tilde{K}_a(\mathcal{P}, w)$, and $\tilde{K}_b(\mathcal{P}, w)$, for all \mathcal{P} and w , and indeed follows easily from the definition of the 2-cells in a diagram complex. We record this in a lemma.

Lemma 4.1. *Let Δ', Δ'' be 1-cells in $\tilde{K}_b(\mathcal{P}, w)$. The following statements are equivalent:*

- (1) Δ' and Δ'' are simple square equivalent;
- (2) *There is some maximal shaded transistor T in Δ' such that Δ'' is the result of removing T and all of its bottom wires from Δ' (or the reverse statement is true, with Δ' and Δ'' reversing roles).*

Fix a 1-cell $\Delta \subseteq \tilde{K}_b(\mathcal{P}, w)$. Let H_Δ denote the combinatorial hyperplane corresponding to Δ . Let \mathcal{T}_Δ denote the collection of transistors of Δ . Let T denote the (unique) white transistor in \mathcal{T}_Δ . Consider the collection $M_\Delta = \{T' \mid T' \text{ is a transistor in } \Delta, T' \leq T\}$ (the inequality sign refers to the partial order on transistors). We can associate to this collection of transistors a vertex $\min(H_\Delta)$, called the *minimal vertex of H_Δ* . Simply remove all transistors in $\mathcal{T}_\Delta - M_\Delta$ along with their bottom wires, and then shade the white transistor. The result is necessarily a vertex by Lemma 3.1. It is clear that $\min(H_\Delta)$ depends only on the hyperplane H_Δ .

For example, if $\Delta = \Delta_1, \Delta_2, \Delta_3$, or Δ_4 from Figure 5 b), then $\min(H_\Delta)$ is the vertex at the top of Δ_2 . We note one property of minimal vertices: a vertex is

minimal if and only if it contains a unique maximal transistor. If a vertex Δ has a unique maximal transistor, then the hyperplane H_Δ corresponding to Δ is the square equivalence class of the edge obtained by painting the maximal transistor white.

We are interested in $\min(H_\Delta)$ because of the following lemma. In all that follows, we let our basepoint $*$ be the unique vertex in $\tilde{K}_b(\mathcal{P}, w)$ having no transistors.

Lemma 4.2. *Let H be a hyperplane in $\tilde{K}_b(\mathcal{P}, w)$. If $\bar{\Delta}$ is an arbitrary vertex in $\tilde{K}_b(\mathcal{P}, w)$, then $\bar{\Delta}$ and the basepoint $*$ lie in different components of $\tilde{K}_b(\mathcal{P}, w) - H$ if and only if $\min(H) \leq \bar{\Delta}$.*

Proof. Let $\bar{\Delta}$ be a vertex in $\tilde{K}_b(\mathcal{P}, w)$; let $\mathcal{T}_{\bar{\Delta}}$ denote the collection of transistors in $\bar{\Delta}$. Choose a function $\alpha: \mathcal{T}_{\bar{\Delta}} \rightarrow \{1, \dots, |\mathcal{T}_{\bar{\Delta}}|\}$ satisfying:

- (1) α is one-to-one;
- (2) if $T_1 < T_2$, then $\alpha(T_1) < \alpha(T_2)$.

We associate a sequence of vertices $* = \bar{\Delta}_0, \bar{\Delta}_1, \dots, \bar{\Delta}_{|\mathcal{T}_{\bar{\Delta}}|} = \bar{\Delta}$, where $\bar{\Delta}_i$ is the (unique) vertex determined by $\alpha^{-1}(\{1, 2, \dots, i\})$ under the correspondence in Lemma 3.1. It is not difficult to see that $\bar{\Delta}_i$ is connected to $\bar{\Delta}_{i+1}$ by a unique edge for $i = 0, 1, \dots, |\mathcal{T}(\bar{\Delta})| - 1$. We let p_α denote the edge-path consisting of these edges.

For example, Figure 6 a) shows a picture with a numbering α of its transistors, along with the corresponding edge-path (Figure 6 b)).

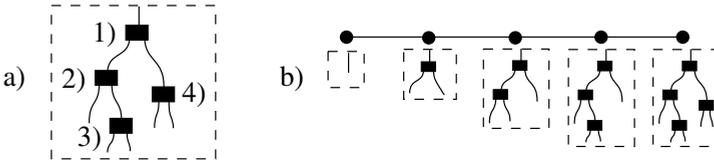


Figure 6. a) A labelling α of the picture $\bar{\Delta}$, and b) the associated edge-path in $\tilde{K}_b(\mathcal{P}, x)$.

We claim that p_α is a geodesic in the 1-skeleton $\tilde{K}_b(\mathcal{P}, w)^1$. Suppose that p is an arbitrary edge-path connecting $*$ to $\bar{\Delta}$; let $* = \bar{\Delta}'_0, \bar{\Delta}'_1, \dots, \bar{\Delta}'_m = \bar{\Delta}$ be the vertices lying along the path p , listed in the order they are visited. It is clear from the definition of edges in $\tilde{K}_b(\mathcal{P}, w)$ that $\bar{\Delta}'_{i+1}$ is obtained from $\bar{\Delta}'_i$ ($0 \leq i \leq m - 1$) by either removing a maximal transistor from the bottom of $\bar{\Delta}'_i$, or adding a new maximal transistor to $\bar{\Delta}'_i$. It immediately follows from this that $\ell(p) \geq |\mathcal{T}_{\bar{\Delta}}|$. This proves the claim.

Suppose that $\min(H_\Delta) \leq \bar{\Delta}$. Lemma 3.1 implies that $\min(H_\Delta)$ corresponds to an initial collection \mathcal{T} of transistors. Suppose that $|\mathcal{T}| = n$. It follows that we can define $\alpha: \mathcal{T}_{\bar{\Delta}} \rightarrow \{1, \dots, n, \dots, |\mathcal{T}_{\bar{\Delta}}|\}$ in such a way that $\alpha|_{\mathcal{T}}: \mathcal{T} \rightarrow \{1, \dots, n\}$ is another labelling function satisfying (1) and (2) above. In this case, $\bar{\Delta}_n = \min(H_\Delta)$ and $\bar{\Delta}_{n-1}$ is the vertex obtained by removing the (unique) maximal transistor in $\bar{\Delta}_n$. It

immediately follows that the edge $[\bar{\Delta}_{n-1}, \bar{\Delta}_n]$ is a member of the combinatorial hyperplane H_Δ . Since a geodesic edge-path p_α from $*$ to $\bar{\Delta}$ crosses H_Δ , Theorem 2.5 (3) implies that $*$ and $\bar{\Delta}$ lie on opposite sides of H_Δ , proving one direction.

Conversely, suppose that $\bar{\Delta}$ and $*$ are separated by the hyperplane H_Δ . Theorem 2.5 (3) says that a geodesic edge-path p crosses precisely the hyperplanes separating the initial vertex of p from the terminal vertex of p . It follows that, for some $k \in \{0, \dots, |\mathcal{T}_{\bar{\Delta}}| - 1\}$, the edge $[\bar{\Delta}_k, \bar{\Delta}_{k+1}]$ represents the hyperplane H_Δ . Under the correspondence in Lemma 3.1, $\bar{\Delta}_k$ corresponds to a collection $\mathcal{T}_{\bar{\Delta}_{k+1}} - \{T\}$ of transistors in $\bar{\Delta}_{k+1}$, where $T \in \mathcal{T}_{\bar{\Delta}_{k+1}}$. The edge $[\bar{\Delta}_k, \bar{\Delta}_{k+1}]$ can be described in terms of pictures as follows: draw $\bar{\Delta}_{k+1}$, but leave the transistor T unshaded. According to the definition, we obtain $\min(H_\Delta)$ by shading T , and then taking the picture corresponding to $\{T' \in \mathcal{T}_{\bar{\Delta}_{k+1}} \mid T' \leq T\}$. It is thus clear that $\min(H_\Delta) \leq \bar{\Delta}_{k+1} \leq \bar{\Delta}$. \square

Proposition 4.3. *Let H_1^+ and H_2^+ be two positive half-spaces in $\tilde{K}_b(\mathcal{P}, w)$. Let Δ_1 and Δ_2 be their minimal vertices.*

- (1) $H_2^+ \leq H_1^+$ if and only if $\Delta_2 \leq \Delta_1$.
- (2) $H_1^+ \cap H_2^+ \neq \emptyset$ if and only if $\{\Delta_1, \Delta_2\}$ has an upper bound in $\tilde{K}_b(\mathcal{P}, w)^0$.

Proof. (1) (\Rightarrow) Suppose $H_2^+ \leq H_1^+$. This means that $H_1^+ \subseteq H_2^+$. Thus, every vertex $\Delta \in H_1^+$ is separated from $*$ by H_2 . In particular, Δ_1 is so separated from $*$. By the previous lemma, $\Delta_2 \leq \Delta_1$.

(\Leftarrow) Suppose $\Delta_2 \leq \Delta_1$. It is sufficient to check the inclusion $H_1^+ \subseteq H_2^+$ on vertices. If Δ is a vertex in H_1^+ , then $\Delta_1 \leq \Delta$. It follows that $\Delta_2 \leq \Delta$, so $\Delta \in H_2^+$ by the previous lemma.

(2) Both directions are immediate consequences of the previous lemma. \square

We now obtain the desired characterization of profiles in terms of pictures.

Theorem 4.4. *Let the basepoint $*$ $\in \tilde{K}_b(\mathcal{P}, w)$ be the unique vertex having no transistors. Let Δ be an infinite $(w, *)$ -picture over the semigroup presentation \mathcal{P} , i.e., a picture in the sense of Section 3, except that the transistor and wire sets are countably infinite. Let us suppose as well that Δ satisfies the following conditions:*

- (1) *For any transistor $T \in \mathcal{T}_\Delta$, the set $(-\infty, T] = \{T' \in \mathcal{T}_\Delta \mid T' \leq T\}$ is finite.*
- (2) *There are no maximal elements in the set \mathcal{T}_Δ of transistors in Δ .*
- (3) *No wire is attached to the bottom of the frame of Δ .*

The picture Δ determines a unique profile, i.e., a non-empty collection of positive half-spaces \mathcal{H}_Δ in $\tilde{K}_b(\mathcal{P}, w)$ satisfying properties (1)–(3) in Proposition 2.7. Conversely, a profile in $\tilde{K}_b(\mathcal{P}, w)$ determines a unique infinite picture Δ satisfying properties (1)–(3) above.

The indicated correspondences are mutually inverse.

Proof. Suppose that Δ is an infinite $(w, *)$ -picture over the semigroup presentation \mathcal{P} satisfying the properties above. The transistors of Δ are in one-to-one correspondence with a collection of hyperplanes in the following way. Let T be a transistor of Δ ; we consider the collection $(-\infty, T]$ of all transistors in Δ which are less than or equal to T in the partial order on transistors. By Lemma 3.1 and our assumption that $(-\infty, T]$ is finite, this collection of transistors corresponds to a unique vertex, and this vertex is the minimal vertex of a unique hyperplane H_T . Note that Lemma 3.1 also implies that the correspondence between transistors and hyperplanes is one-to-one.

We consider the properties of the collection $\mathcal{H}_\Delta = \{H_T \mid T \in \mathcal{T}_\Delta\}$. First, let H_{T_1}, \dots, H_{T_n} be hyperplanes in \mathcal{H}_Δ . Consider the collection $\mathcal{T}' = \{T \in \mathcal{T}_\Delta \mid T \leq T_i \text{ for some } i \in \{1, \dots, n\}\}$. By Lemma 3.1, the collection \mathcal{T}' corresponds to a vertex $\Delta_{\mathcal{T}'}$. Moreover, we have that $(-\infty, T_i] \subseteq \mathcal{T}'$, for $i = 1, \dots, n$, from which it follows that $\min(H_{T_i}) \leq \Delta_{\mathcal{T}'}$, for $i = 1, \dots, n$. This, in turn, implies that $\Delta_{\mathcal{T}'} \in H_{T_1}^+ \cap \dots \cap H_{T_n}^+$, by Lemma 4.2. Thus property (1) from Proposition 2.7 holds.

If $H_T \in \mathcal{H}_\Delta$, then, by the assumption that T is not maximal, there is some $T_1 \in \mathcal{T}_\Delta$ such that $T < T_1$. It follows from this that $(-\infty, T] \subsetneq (-\infty, T_1]$; this implies that the vertices $\min(H_T)$, $\min(H_{T_1})$ under the correspondence from Lemma 3.1 satisfy $\min(H_T) < \min(H_{T_1})$. By Lemma 4.3, $H_T^+ < H_{T_1}^+$, so property (2) from Proposition 2.7 holds.

Checking property (3) from Proposition 2.7 is an easy exercise using the properties of the correspondence in Lemma 3.1.

Conversely, suppose that \mathcal{H} is a non-empty collection of hyperplanes in $\tilde{\mathcal{K}}_b(\mathcal{P}, w)$ satisfying properties (1)–(3) of Proposition 2.7. Choose a finite collection of hyperplanes $H_{T_1}, H_{T_2}, \dots, H_{T_n}$. By property (1) from Proposition 2.7, $H_{T_1}^+ \cap H_{T_2}^+ \cap \dots \cap H_{T_n}^+ \neq \emptyset$. This implies that there is some vertex $\hat{\Delta}$ in the latter intersection, which means, by Lemma 4.2, that $\min(H_{T_i}) \leq \hat{\Delta}$ for $i = 1, \dots, n$. Since the collection $\{\min(H_{T_i}) \mid i = 1, \dots, n\}$ has an upper bound, Lemma 3.2 (2) of [9] implies that it has a least upper bound. Thus, we have shown that any finite collection of minimal vertices for hyperplanes in \mathcal{H} has a least upper bound. (This least upper bound is a “union” of the labels for these vertices, in an appropriate sense. Note that it will not in general be a minimal vertex itself.)

Let $\Delta_1, \Delta_2, \dots, \Delta_n, \dots$ be the sequence consisting of all minimal vertices for hyperplanes in \mathcal{H} . Since any finite collection of these hyperplanes has a least upper bound, we can identify the direct limit of this sequence with an infinite diagram Δ . It is clear that Δ has properties (1) and (3) from the statement of the Theorem. Property (2) follows easily from the fact that the collection \mathcal{H} satisfies (2) from Proposition 2.7.

We leave the proof of the final statement to the reader. □

We will sometimes require a lemma which gives a necessary condition on the open cube C through which a geodesic ray c with profile $P(c)$ can travel. The

condition involves the largest vertex Δ_C in the closure \bar{C} , which always exists, and can be obtained by shading each transistor in the picture representative for C (see for instance Figure 4 in Section 3).

Lemma 4.5. *Let $c : [0, \infty) \rightarrow \tilde{K}_b(\mathcal{P}, w)$ be a geodesic ray issuing from the base vertex $*$. Let Δ be the infinite picture representing the profile of c . Let C be an open cube of $\tilde{K}_b(\mathcal{P}, w)$ satisfying $(\text{Im } c) \cap C \neq \emptyset$.*

If Δ_C is the largest vertex in \bar{C} , then $\Delta_C \leq \Delta$.

Proof. Let Δ_C be the largest vertex of C . Consider the collection of all maximal transistors T_1, \dots, T_n in Δ_C ; let Δ_{T_i} be the unique vertex determined by $(-\infty, T_i]$ under the correspondence from Lemma 3.1. Note that each Δ_{T_i} is the minimal vertex of a hyperplane H_{T_i} , and all of the hyperplanes H_{T_1}, \dots, H_{T_n} are distinct.

Let $\hat{\Delta}_C$ denote the representative for C , which is a picture consisting of shaded and unshaded transistors, as in Figure 4. Some of the transistors T_i are shaded in $\hat{\Delta}_C$; others are unshaded in $\hat{\Delta}_C$.

If T_i is shaded, then $\Delta_{T_i} \leq \Delta_C$, so that H_{T_i} separates Δ_C from $*$. Moreover, H_{T_i} does not pass through C , since the hyperplanes H satisfying $H \cap C \neq \emptyset$ are precisely the collection of all H_{T_j} such that T_j is unshaded in $\hat{\Delta}_C$. It follows that each point x in $(\text{Im } c) \cap C$ can be connected to Δ_C without crossing H_{T_i} ; therefore $x \in H_{T_i}^+$, so $H_{T_i}^+ \in P(c)$.

If T_i is unshaded in $\hat{\Delta}_C$, then $H_{T_i} \cap C \neq \emptyset$. It follows from Lemma 2.6 (2) that $H_{T_i}^+ \in P(c)$.

The correspondence of Theorem 4.4 implies that $\Delta_{T_1}, \dots, \Delta_{T_n} \leq \Delta$. This means that the least upper bound $\tilde{\Delta}$ of $\{\Delta_{T_1}, \dots, \Delta_{T_n}\}$ exists, and satisfies $\tilde{\Delta} \leq \Delta$. But clearly $\tilde{\Delta} = \Delta_C$, since the latter vertex is an upper bound of $\{\Delta_{T_1}, \dots, \Delta_{T_n}\}$, and any proper initial subset of \mathcal{T}_{Δ_C} would fail to contain at least one of the transistors T_1, \dots, T_n . □

4.2. The action on profiles

Proposition 4.6. *There is a well-defined action of $\mathcal{D}_b(\mathcal{P}, w)$ on the set of all profiles. If Δ is a profile and $\Delta_1 \in \mathcal{D}_b(\mathcal{P}, w)$, then $\Delta_1 * \Delta$ can be computed as follows. First, form the concatenation $\Delta_1 \circ \Delta$ and remove all dipoles. Second, remove all maximal transistors from the resulting infinite diagram, until no maximal transistors remain. The result is $\Delta_1 \cdot \Delta$.*

Proof. Let $\mathcal{H}, \mathcal{H}'$ be two collections of positive half-spaces in $\tilde{K}_b(\mathcal{P}, w)$. We write $\mathcal{H} \sim \mathcal{H}'$ and $[\mathcal{H}] = [\mathcal{H}']$ if \mathcal{H} and \mathcal{H}' are cofinal, i.e., if for any $H_1^+ \in \mathcal{H}$, there exists $(H'_1)^+ \in \mathcal{H}'$ such that $H_1^+ \leq (H'_1)^+$ and for any $(H'_2)^+ \in \mathcal{H}'$, there exists $H_2^+ \in \mathcal{H}$ such that $(H'_2)^+ \leq H_2^+$. It is fairly clear that \sim is an equivalence relation on the set of all collections of positive half-spaces in $\tilde{K}_b(\mathcal{P}, w)$.

The group $\mathcal{D}_b(\mathcal{P}, w)$ does not act in an obvious way on the set of equivalence classes, since a group element $\bar{\Delta} \in \mathcal{D}_b(\mathcal{P}, w)$ does not necessarily map a positive half-space to a positive half-space. Indeed, $\bar{\Delta} \cdot H^+$ is a negative half-space if and only if $* \in \bar{\Delta} \cdot H^+$, i.e., if and only if $\bar{\Delta}^{-1} \cdot * \in H^+$. Now $\bar{\Delta}^{-1} \cdot * \in H^+$ if and only if $\min(H^+) \leq \bar{\Delta}^{-1} \cdot *$. There are only finitely many vertices Δ' satisfying $\Delta' \leq \bar{\Delta}^{-1} \cdot *$ (all determined by the correspondence in Lemma 3.1). It follows from this that a given $\bar{\Delta} \in \mathcal{D}_b(\mathcal{P}, w)$ maps at most finitely many positive half-spaces to negative ones.

We obtain an action on profiles in the following way. Identify a profile Δ with the unique equivalence class $[\mathcal{H}]$ such that $\Delta \in [\mathcal{H}]$. For a given $\bar{\Delta} \in \mathcal{D}_b(\mathcal{P}, w)$, choose a collection $\mathcal{H}' \in [\mathcal{H}]$ such that $\bar{\Delta} \cdot H^+$ is a positive half-space, for each $H^+ \in \mathcal{H}'$. It is possible to do this since each $\mathcal{H}' \in [\mathcal{H}]$ is necessarily infinite. We define $\bar{\Delta} * \Delta$ to be $[\bar{\Delta} \cdot \mathcal{H}']$. It is not difficult to see that $[\bar{\Delta} \cdot \mathcal{H}']$ contains a unique profile, and that the definition of $*$ does not depend upon the choice of $\mathcal{H}' \in [\mathcal{H}]$. It follows that $*$ is an action on profiles.

We claim two things: first, that $\Delta * P(c) = P(\Delta * c)$, for any geodesic ray $c: [0, \infty) \rightarrow \tilde{K}_b(\mathcal{P}, w)$ issuing from $*$; second, that the action $*$ from the previous paragraph has the description promised in the statement of the proposition.

Recall the definition of the action $*$: $\mathcal{D}_b(\mathcal{P}, w) \times \partial \tilde{K}_b(\mathcal{P}, w) \rightarrow \partial \tilde{K}_b(\mathcal{P}, w)$ on the space at infinity. If $c \in \partial \mathcal{D}_b(\mathcal{P}, w)$, then, for any $\bar{\Delta} \in \mathcal{D}_b(\mathcal{P}, w)$, $\bar{\Delta} \cdot c$ is simply the translate of c by the usual action of $\bar{\Delta}$ on $\tilde{K}_b(\mathcal{P}, w)$. The ray $\bar{\Delta} * c$ is the unique ray issuing from $*$ and asymptotic to $\bar{\Delta} \cdot c$, the existence of which is guaranteed by Proposition 2.1.

We now prove the first claim. Let $c \in \partial \tilde{K}_b(\mathcal{P}, w)$ and let $\bar{\Delta} \in \mathcal{D}_b(\mathcal{P}, w)$. We choose some cofinal subset $\mathcal{H} \subseteq P(c)$ such that $\bar{\Delta} \cdot H^+$ is a positive half-space, for any $H^+ \in \mathcal{H}$. It is clear that $\bar{\Delta} \cdot c$ crosses each of the hyperplanes in $\bar{\Delta} \cdot \mathcal{H}$, since c crosses each of the hyperplanes in \mathcal{H} . Since $\bar{\Delta}$ maps positive half-spaces in \mathcal{H} to positive half-spaces, $\bar{\Delta} \cdot c$ intersects each $\bar{\Delta} \cdot H^+ \in \bar{\Delta} \cdot \mathcal{H}$ in an open ray, just as c intersects each $H^+ \in \mathcal{H}$ in an open ray. By Lemma 2.3, $d_{\bar{\Delta} \cdot H}(\bar{\Delta} \cdot c(t)) \rightarrow \infty$ as $t \rightarrow \infty$ for any H such that $H^+ \in \mathcal{H}$. It follows from this, first, that $d_{\bar{\Delta} \cdot H}(\bar{\Delta} * c(t))$ also goes to infinity as $t \rightarrow \infty$, and second, that $(\bar{\Delta} * c)(t) \in \bar{\Delta} \cdot H^+$ for t sufficiently large. This implies that $\bar{\Delta} \cdot \mathcal{H} \subseteq P(\bar{\Delta} * c)$.

Next, we need to show that, for any positive half-space $H^+ \in P(\bar{\Delta} * c)$, there is $\bar{H}^+ \in \bar{\Delta} \cdot \mathcal{H}$ such that $H^+ \leq \bar{H}^+$. Choose a sequence of positive half-spaces $H^+ = H_0^+ < H_1^+ < H_2^+ < \dots$ in $P(\bar{\Delta} * c)$. By Lemma 2.3, we know that $d_{H_i}((\bar{\Delta} * c)(t)) \rightarrow \infty$ as $t \rightarrow \infty$, for any $i \in \{0, 1, \dots, n, \dots\}$. By the definition of $P(\bar{\Delta} * c)$, for any i , $(\bar{\Delta} * c)(t) \in H_i^+$ for t sufficiently large. It now follows from the fact that $\bar{\Delta} \cdot c$ and $\bar{\Delta} * c$ are asymptotic that, for any i , $(\bar{\Delta} \cdot c)(t) \in H_i^+$ for t sufficiently large. This implies that $\bar{\Delta} \cdot c$ crosses at least one of the H_i , and therefore all H_j for j sufficiently large, for otherwise

$$(\bar{\Delta} \cdot c)([0, \infty)) \subseteq \bigcap_{i=0}^{\infty} H_i^+ = \emptyset.$$

We choose some H_j^+ large enough that $\bar{\Delta}^{-1} \cdot H_j^+$ is a positive half-space. Since c crosses $\bar{\Delta}^{-1} \cdot H_j$, it follows that $\bar{\Delta}^{-1} \cdot H_j^+ \in P(c)$. Since \mathcal{H} is cofinal in $P(c)$, there is $\hat{H}^+ \in \mathcal{H}$ such that $\bar{\Delta}^{-1} \cdot H_j^+ \leq \hat{H}^+$. This implies that $H_j^+ \leq \bar{\Delta} \cdot \hat{H}^+ \in \bar{\Delta} \cdot \mathcal{H}$. Now we have shown that $H^+ \leq H_j^+ \in \bar{\Delta} \cdot \mathcal{H}$.

It follows that $P(\bar{\Delta} * \mathcal{H}) = [\bar{\Delta} \cdot \mathcal{H}]$ under the identification of $[\bar{\Delta} \cdot \mathcal{H}]$ with a profile. We have thus shown that $P(\bar{\Delta} * c) = \bar{\Delta} * P(c)$. It immediately follows from this that the action of $\mathcal{D}_b(\mathcal{P}, w)$ on profiles is well defined: if c, c' have the same profile, then so also do $\bar{\Delta} * c$ and $\bar{\Delta} * c'$.

Now we prove the second claim. Let Δ be an infinite picture representing a profile and let $\bar{\Delta} \in \mathcal{D}_b(\mathcal{P}, w)$. We choose a cofinal collection of transistors \mathcal{T}' in Δ (which are identified with positive half-spaces by the correspondence in Theorem 4.4) such that no transistor in \mathcal{T}' forms a dipole in the concatenation $\bar{\Delta} \circ \Delta$. The above description of the action implies that $\bar{\Delta} * \Delta$ is the collection of all positive half-spaces such that $H^+ \leq \bar{\Delta} \cdot H_{T'}^+$, for some $T' \in \mathcal{T}'$. This collection of positive half-spaces may be identified with the collection of transistors T in the reduced concatenation $\bar{\Delta} \circ \Delta$ satisfying $T \leq T'$, for some $T' \in \mathcal{T}'$. By the cofinality of \mathcal{T}' in Δ , these transistors T are precisely those for which there exists an infinite sequence $T = T_0 < T_1 < \dots < T_n < \dots$ where each T_i is a transistor of $\bar{\Delta} \circ \Delta$. The second claim follows. □

5. Fixed profiles under the actions of F , T , and V

The results of the previous section largely reduce the problem of finding fixed points in $\partial \tilde{K}(\mathcal{P}, x)$, $\partial \tilde{K}_a(\mathcal{P}, x)$, and $\partial \tilde{K}_b(\mathcal{P}, x)$ (where $\mathcal{P} = \langle x \mid x = x^2 \rangle$) to the algebraic problem of finding globally fixed profiles. The latter problem is quite easy; we give a complete classification of fixed profiles for F , T , and V in this section.

5.1. Conventions. We fix some conventions for portraying profiles (and pictures) over the semigroup presentation $\langle x \mid x = x^2 \rangle$.

First, we draw every transistor in a picture or profile as a point and omit the frame, so that, for instance, the element $x_0 \in F$ (depicted as an ordinary picture on the left) looks like the right half of Figure 7.

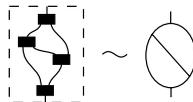


Figure 7. A convention for drawing pictures over the semigroup presentation $\mathcal{P} = \langle x \mid x = x^2 \rangle$. On the left, we have a picture drawn in the usual fashion; on the right is its equivalent.

We also need some conventions that will allow us to portray infinite pictures using a finite amount of space. We draw an empty dot at the end of a wire to indicate that the bottom of the wire does not connect to any transistor. A solid dot at the end of a wire indicates that the wire connects to the top of some transistor. This transistor may be either an (x, x^2) -transistor or an (x^2, x) -transistor. (In particular, one should not assume that the pictures below the solid dots are trees. Indeed, we will consider several cases in which it is possible to connect two solid dots by a path whose interior lies entirely in the omitted portion of the picture.) A wire with no dot at the end may connect to a transistor or not; we make no assumption one way or the other.

Finally, we let T_m denote the full ordered rooted binary tree of depth m . We let \dot{T}_m denote the full ordered rooted binary tree of depth m , where each leaf ends in a solid dot. Thus, \dot{T}_m is the unique ordered rooted binary tree having 2^m dotted leaves, each at distance m from the root. The dot on each leaf (wire) implies that each connects to the top of some transistor. The picture T_m is the same tree, but without the dots on the leaves. Thus, no particular wire in T_m which corresponds to a leaf necessarily leads to the top of a transistor. Notice, however, that in a profile at least one of the leaves beneath a given degree 3 vertex must attach to the top of a transistor, since there are no maximal transistors in a profile (by Theorem 4.4).

5.2. Thompson’s group F . Let Δ be a profile of $\tilde{K}(\mathcal{P}, w)$ which is fixed by all of F . Without loss of generality, we can assume that Δ has one of the forms in Figure 8. (The only other possible cases are $1'-6'$, which are obtained by reflecting 1–6 across a vertical axis. Note that the resulting cases are not mutually exclusive.)



Figure 8. The six cases.

5.2.1. The even-numbered cases. We consider even-numbered cases first. Let x_0 act on any profile Δ covered by Case 2, 4, or 6. The results appear in Figure 9. After cancelling dipoles twice, we arrive at the infinite pictures in the column at the far right of the figure. We claim that these infinite pictures necessarily contain no dipoles, no matter how the wires terminating in black dots are connected to transistors. (It is clear also that these contain no maximal transistors.)

To prove the claim, first note that any dipole in the product $x_0 \cdot \Delta$ must be formed of one transistor in x_0 and another in Δ . Thus, an infinite picture on the right side of Figure 9 contains a dipole only if one of the two pictured vertices of degree 3 (both of which represent transistors from x_0) can form the top half of a dipole.

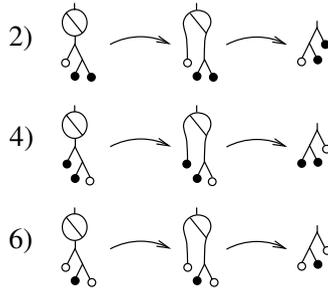


Figure 9. The even-numbered cases.

This is clearly impossible in Cases 2 and 6. In Case 4, it is enough to show that the lower transistor cannot form the top half of a dipole. If we assume that it does, then the original profile Δ would have the form in Figure 10. No profile can have this form, however, since it is impossible to connect the black-dotted wire to a transistor without forming a dipole, and Δ cannot contain dipoles. This proves the claim.



Figure 10. An impossible profile. No matter how we attach the black-dotted wire to the top of a transistor, a dipole will be formed.

Finally, we compare the reduced profiles $x_0 * \Delta$ at the right in Figure 9 with the originals in Figure 8. Since Δ is fixed by all of F , we must have that $x_0 * \Delta = \Delta$. This is impossible, as we easily see. For instance, in Case 2, the left wire dangling from the bottom of the topmost transistor in Δ does not connect to a transistor, but the wire of the same description in $x_0 * \Delta$ does. The other even cases are left as easy exercises.

5.2.2. Cases 3 and 5. If Δ is represented in Case 5, then it must have the form in Figure 11 a), where Δ' is another profile. If we let x_0^{-1} act on Δ , then after removing a dipole and an exposed transistor, we arrive at the profile on the far right of Figure 11 b). A simple induction using the fact that $\Delta = x_0^{-1} * \Delta$ now shows Δ is the (unique) profile of the form depicted in Figure 11 c). It is easy to check that Δ is fixed by all of F ; indeed, it is enough to show that Δ is fixed by the generators x_0 and x_1 . We leave this verification as an exercise.

If Δ is represented in Case 3, then it must have the form in Figure 12 a), where Δ' and Δ'' are profiles. An argument similar to the one used for Case 5 shows that Δ'' has the form depicted in Figure 11 c), and Δ' is the result of reflecting Δ'' across

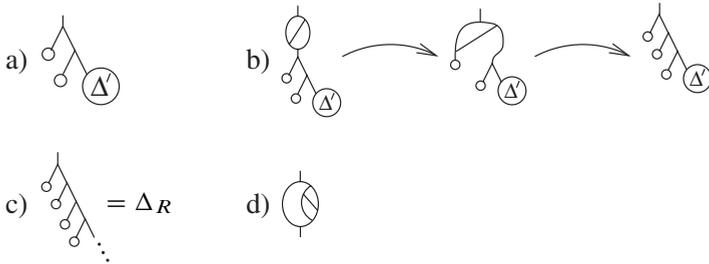


Figure 11. a) The general form of a profile from Case 5. b) The effect of letting x_0^{-1} act. c) The profile Δ_R . d) The element $x_1 \in F$.

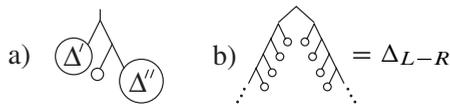


Figure 12. a) The general form of a profile from Case 3. b) The profile Δ_{L-R} .

a vertical axis. The details are left as an exercise. It follows that Δ is the profile depicted in Figure 12 b), which is indeed fixed by both x_0 and x_1 .

5.2.3. Case 1. We now turn to Case 1. Let x_0 act on Δ . There are two subcases to consider: either Δ has the form depicted in Figure 14 a) (and thus the infinite picture at the far right in Figure 13 contains a dipole) or the infinite picture at the far right in Figure 13 is reduced.

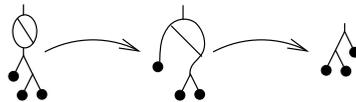


Figure 13. The action of x_0 on a profile from Case 1.

We now rule out the first possibility using the fact that Δ is invariant under the action of F . Let x_1 act on Δ ; after reducing two dipoles we arrive at the infinite picture $\hat{\Delta}$ on the far right of Figure 14 b). The transistors enclosed by the dotted circle were contributed by x_1 , and any dipole in $\hat{\Delta}$ would have to involve one of these three transistors. Now note that, of these, only the transistor labelled $*$ could form half of a dipole; the others could not, even after we cancel any dipole involving $*$. Now we compare Δ and $\hat{\Delta}$. Under any isomorphism between Δ and $\hat{\Delta}$, the transistors labelled 1 and 2 in $\hat{\Delta}$ must correspond (respectively) to the transistors labelled i) and ii) in Δ (as depicted in Figure 14 a). This is not possible, since ii) is a (x^2, x) -transistor and 2 is a (x, x^2) -transistor.

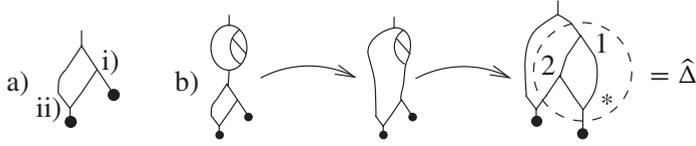


Figure 14. a) If we let x_0 act on this profile, the transistor labelled ii) will form half of a dipole. b) The action of x_1 on the profile from a).

It follows that we can assume that there are no dipoles in the infinite picture at the far right of Figure 13. After comparing this profile with the profile Δ , we can conclude that Δ has the form \hat{T}_2 . We now multiply Δ by $x_0, x_0x_1x_0^{-1}, x_1x_0^{-1}$ and x_0^{-1} . The results are listed in Figure 15 a)–d) (in the same order). Dotted circles enclose the

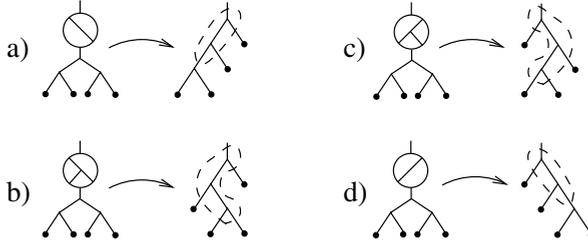


Figure 15. The actions of a) x_0 , b) $x_0x_1x_0^{-1}$, c) $x_1x_0^{-1}$ and d) x_0^{-1} on Δ .

transistors that were contributed by the acting element. If we knew that there were no dipoles in the infinite pictures at the right, we could use the fact that all of them are equal to Δ in order to conclude that Δ has the form \hat{T}_3 . The infinite pictures on the far right of a) and d) in Figure 15 are necessarily reduced. The profiles in b) and c) will be reduced unless Δ has the form in Figure 16 b).

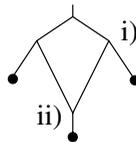


Figure 16. This profile would form dipoles in cases b) and c) from Figure 15.

We now rule out the latter possibility. Let x_1 act on Δ . Any isomorphism between $\hat{\Delta}$ and Δ must match the transistors labelled 1) and 2) in Figure 17 with the transistors labelled i) and ii), respectively. This is impossible, since 2) is an (x, x^2) -transistor and ii) is an (x^2, x) -transistor.

It now follows that Δ has the form of \hat{T}_3 . We multiply Δ by $x_0, x_0x_1x_0^{-1}, x_1x_0^{-1}$,

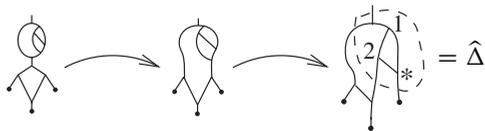


Figure 17. The action of x_1 on the profile from Figure 16.

x_0^{-1} . The results appear in Figure 18. Note that it is no longer possible for the transistors from the acting elements (circled) to form dipoles, so all of the profiles in Figure 18 are reduced. Each of these profiles is equal to Δ , since Δ is invariant under the action of F . It follows that Δ has the form of \hat{T}_4 .

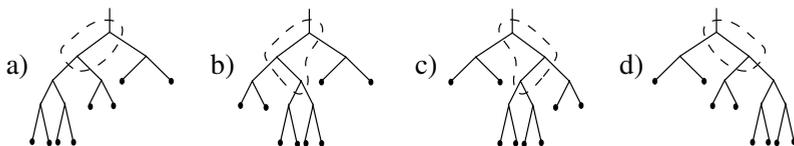


Figure 18. The results of letting a) x_0 , b) $x_0x_1x_0^{-1}$, c) $x_1x_0^{-1}$, and d) x_0^{-1} act on \hat{T}_3 .

We now repeat this argument, letting the same four elements act on Δ . In this way, we conclude by induction that Δ has the form of the full infinite binary T_∞ . We write $\Delta = \Delta_\infty$.

We have now proved Theorem 5.1.

Theorem 5.1. *Thompson’s group F fixes exactly four profiles:*

$$\Delta_L, \Delta_R, \Delta_{L-R}, \text{ and } \Delta_\infty.$$

These profiles come from Cases 5', 5, 3, and 1, respectively.

5.3. Thompson’s groups T and V . This subsection is devoted to a proof of the following theorem:

Theorem 5.2. *Thompson’s groups T and V fix only the profile Δ_∞ .*

Proof. Let $\mathcal{P} = \langle x \mid x = x^2 \rangle$. Suppose that the group $\mathcal{D}_a(\mathcal{P}, x) \cong T$ fixes the profile Δ . It is not difficult to see that Δ has the form in Figure 19 a), without loss of generality. We let $\pi_1 \in \mathcal{D}_a(\mathcal{P}, x)$ act on Δ ; the result is portrayed in Figure 19 c), where the circled transistor is contributed by π_1 . It follows that the picture in c) is reduced. Combining a) and c), which are both equivalent since $\pi_1 * \Delta = \Delta$, we have that Δ has the form T_2 .

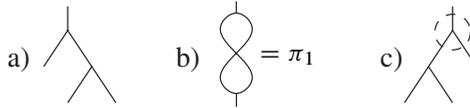


Figure 19. a) The general form of a profile in the cubical complexes for T and V ; b) the acting element π_1 , which represents a half-rotation of the circle; c) the effect of letting π_1 act on the profile from a).

Next, we claim that at least one of the four leaves of T_2 connects to the top of an (x, x^2) -transistor. If not, then consider an (x, x^2) -transistor T in Δ which is distinct from the three (x, x^2) -transistors in T_2 , and minimal among (x, x^2) -transistors with this property. The wire attached to the top of such a transistor could only lead up to the bottom of an (x^2, x) -transistor T' . This implies that T and T' form a dipole, which contradicts the fact that Δ is reduced.

We therefore assume, without loss of generality, that Δ has the form in Figure 20 a). Now we multiply Δ by π_2 , π_2^2 , and π_2^3 , where π_2 is as in Figure 20 b) to get the

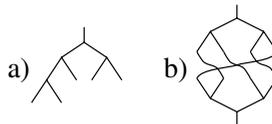


Figure 20. a) Without loss of generality, Δ has this form. b) The element π_2 , which represents a quarter-turn of the circle.

three profiles in Figure 21. Since the profiles from Figure 21 and the profile from Figure 20 a) are all equal to Δ , it follows that Δ has the form T_3 .

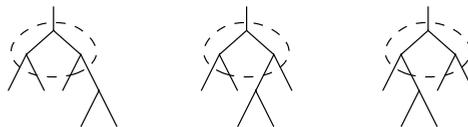


Figure 21. The results of letting the powers of π_2 act on the profile from Figure 20 a).

We then argue, as before, that at least one of the eight leaves at the bottom of T_3 must be attached to the top of an (x, x^2) -transistor. We then multiply Δ by $1, \pi_3, \pi_3^2, \pi_3^3, \dots, \pi_3^7$, where π_3 is a picture representing a one-eighth turn of the circle.

If we compare the eight resulting profiles, we conclude that Δ is equivalent to T_4 . We can continue in a similar way, and eventually conclude that $\Delta = \Delta_\infty$.

This proves the theorem in the case of T , and the proof for V is the same word for word. □

6. The cases of Δ_L , Δ_R , and Δ_{L-R}

Now we consider the geodesic rays $c \in \Delta_L \cup \Delta_R \cup \Delta_{L-R}$. Consider first the profile Δ_{L-R} . By Lemma 4.5, any geodesic ray c in Δ_{L-R} is contained in the subcomplex K of X_F ($= \mathcal{D}(\mathcal{P}, w)$, where $\mathcal{P} = \langle x \mid x = x^2 \rangle$ and $w = x$) pictured in Figure 22.

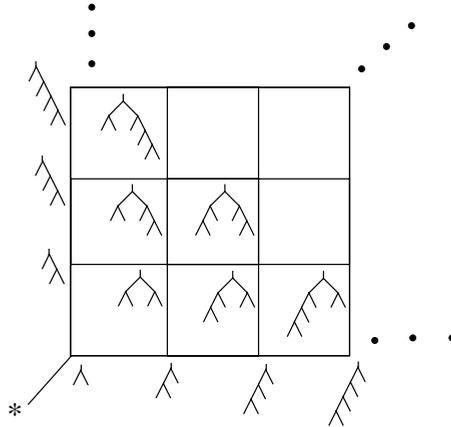


Figure 22. A flat sector in X_F . An integer lattice point (m, n) ($m, n \geq 0$) corresponds to the tree $T_{m,n}$ having a root caret, m carets dangling to the left, and n carets dangling to the right.

Thus K may be naturally identified with $\mathbb{R}^{2,+} \cup I$, where $\mathbb{R}^{2,+} = \{(x, y) \in \mathbb{R}^2 \mid x, y \geq 0\}$, I is the unit interval, and $\mathbb{R}^{2,+} \cap I = \{(0, 0)\}$.

Lemma 6.1. *The inclusion $i : K \rightarrow X_F$ is an isometric embedding.*

Proof. We appeal to Theorem 1 (2) of [6], which says: If X and Y are finite dimensional CAT(0) cubical complexes and $\Phi : X \rightarrow Y$ is a cubical map, then the map Φ is an isometric embedding if and only if, for every vertex $v \in X$, the simplicial map between links $\text{Lk}(v, X) \rightarrow \text{Lk}(\Phi(v), Y)$ induced by Φ is injective with image a full subcomplex of $\text{Lk}(\Phi(v), Y)$. (We refer the reader to [2], p. 102, for a discussion of the link; Crisp and Wiest define cubical maps on p. 443 of [6], and it is clear that the inclusion map is cubical.)

We consider the link of a vertex $T_{m,n}$, where $m, n > 1$, and leave the verifications for the other vertices as an exercise. The link $\text{Lk}(T_{m,n}, K)$ is a square, i.e., the obvious one-dimensional simplicial complex consisting of 4 vertices and 4 edges. This link will be embedded in $\text{Lk}(T_{m,n}, X_F)$ as a full subcomplex if and only if (1) there is no two-dimensional cube C in X_F such that $T_{m,n-1}$ and $T_{m,n+1}$ are both vertices of C , and (2) there is no two-dimensional cube C in X_F such that $T_{m-1,n}$ and $T_{m+1,n}$ are both vertices of C .

We now check (1); the argument for (2) is similar. If there is such a cube C , then C can be represented by a picture as in Figure 4, consisting of k shaded and 2 unshaded transistors. The four corners of C are labelled by pictures having $k, k + 1, k + 1,$ and $k + 2$ transistors. It follows that $k = m + n$, that $T_{m,n-1}$ is the result of leaving off both unshaded transistors, and that $T_{m,n+1}$ is the result of shading both transistors. From this we get a contradiction, since the unshaded transistors of the cube C must both be maximal, and there is no way to remove two transistors that are both maximal in $T_{m,n+1}$ and arrive at $T_{m,n-1}$. \square

Theorem 6.2. *If c is a geodesic ray in X_F issuing from $*$ and $c \notin \Delta_\infty$, then c represents a point at infinity that is fixed by all of F if and only if $c \in \Delta_L \cup \Delta_R \cup \Delta_{L-R}$. The subspace of ∂X_F consisting of $\Delta_L \cup \Delta_R \cup \Delta_{L-R}$ is an arc of Tits length $\pi/2$.*

Proof. If $f : X \rightarrow Y$ is an isometric embedding between CAT(0) spaces, then the induced map $f_\infty : \partial X \rightarrow \partial Y$ is an isometry, where the boundary is endowed with the angular metric ([2], p. 280). By the previous lemma, K is isometrically embedded in X_F ; by Lemma 4.5, a geodesic ray c issuing from $*$ represents $[c] \in \Delta_L \cup \Delta_R \cup \Delta_{L-R}$ if and only if $\text{Im } c \subseteq K$. It follows from this that the image of ∂K under the map $\partial K \rightarrow \partial X_F$ is precisely $\Delta_L \cup \Delta_R \cup \Delta_{L-R}$. The second statement now follows from the fact that ∂K is isometric to $[0, \pi/2]$.

Now suppose that $c \notin \Delta_\infty$ is a geodesic ray in X_F issuing from $*$. If c is fixed by all of F under the action $*$, then the argument of Section 5 shows that $c \in \Delta_L \cup \Delta_R \cup \Delta_{L-R}$.

Conversely, suppose that $c \in \Delta_L \cup \Delta_R \cup \Delta_{L-R}$. It follows from this and Lemma 4.5 that $\text{Im } c \subseteq K$, so every point $x \in \text{Im } c$ is within $1 + 2\sqrt{2}$ of a point in $\{2, 3, \dots\} \times \{2, 3, \dots\} \subseteq \mathbb{R}^{2,+}$. We let $\widehat{T}_{m,n}$ denote the tree in Figure 23, which consists of $T_{m,n}$ and one additional caret:

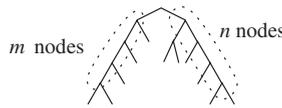


Figure 23. A picture of the tree $\widehat{T}_{m,n}$, which consists of $T_{m,n}$ and one additional caret.

If $m, n \geq 2$, it is routine to check that

$$\begin{aligned} x_0 \cdot T_{m,n} &= T_{m+1,n-1}, \\ x_1 \cdot T_{m,n} &= \widehat{T}_{m,n-1}. \end{aligned}$$

The trees $T_{m,n}$ and $T_{m+1,n-1}$ can be joined by an edge-path of length 2 in X_F . The same goes for $T_{m,n}$ and $\widehat{T}_{m,n-1}$, so $d(x_i \cdot T_{m,n}, T_{m,n}) \leq 2$ ($i = 0, 1$).

Let $t \geq 0$. We have:

$$\begin{aligned} d(c(t), x_i \cdot c(t)) &\leq d(c(t), T_{m,n}) + d(T_{m,n}, x_i \cdot T_{m,n}) + d(x_i \cdot T_{m,n}, x_i \cdot c(t)) \\ &\leq 4 + 4\sqrt{2}. \end{aligned}$$

Since this estimate does not depend upon t , it follows that x_0, x_1 both fix c under the action $*$. This implies that F fixes c , since x_0, x_1 generate F . \square

7. The case of Δ_∞

This section is devoted to an investigation of fixed points in Δ_∞ . Our main result is the following:

Theorem 7.1. *The profile Δ_∞ contains no global fixed point of T or V .*

An immediate consequence of Theorem 7.1 and Theorem 5.2 is:

Corollary 7.2. *Thompson’s groups T and V act without global fixed points on the boundaries-at-infinity of their respective picture complexes.*

The case for Thompson’s group F is more complicated. I do not know whether $\Delta_\infty \subseteq \partial X_F$ contains fixed points of F or not. Example 7.7 and Proposition 7.5 give evidence for and against the existence of fixed points in Δ_∞ , respectively.

Our arguments will use a simple procedure for embedding any CAT(0) cubical complex into Hilbert space.

Proposition 7.3. *If X is a CAT(0) cubical complex, then there is an embedding $\rho: X \rightarrow \ell^2(\mathcal{H})$, where \mathcal{H} is the collection of hyperplanes in X . The map ρ does not increase distances.*

Proof. Let X be a CAT(0) cubical complex with a distinguished base vertex $*$. For any given hyperplane H in X , we identify the closed $1/2$ -neighborhood of H with $H \times [0, 1]$ in such a way that $d(*, H \times \{0\}) < d(*, H \times \{1\})$, and let H_t denote $H \times \{t\}$ for $t \in [0, 1]$.

Let $\rho: X \rightarrow \ell^2(\mathcal{H})$ send x to $\sum_{H \in \mathcal{H}} f_H(x)H$, where

$$f_H(x) = \sup(\{t \in [0, 1] \mid [* , x] \cap H_t \neq \emptyset\} \cup \{0\}).$$

It is clear that ρ is an embedding, that each cube of x is embedded into $\ell^2(\mathcal{H})$ isometrically, and that the restrictions $\rho|_C$ (where C is a cube) agree on overlaps. Note also that each sum $\sum_{H \in \mathcal{H}} f_H(x)H$ is finite, since any pair of vertices in X are separated by at most finitely many hyperplanes. The first statement follows.

Let $[x, y]$ be a geodesic in X . Since the restriction of ρ to each cube is an isometry, ρ preserves the lengths of paths. Therefore,

$$d_X(x, y) = \ell(\rho[x, y]) \geq d_{\ell^2(\mathcal{H})}(\rho(x), \rho(y)). \quad \square$$

Example 7.4. We consider the image of a certain $x \in X_F$ under the map ρ .

Note that each hyperplane occurring in the sum $\rho(x)$ corresponds in a straightforward way to a particular transistor Δ_∞ . We can use this fact to simplify our notation – compare Figure 25 to the right half of Figure 24. In this way, we identify each point

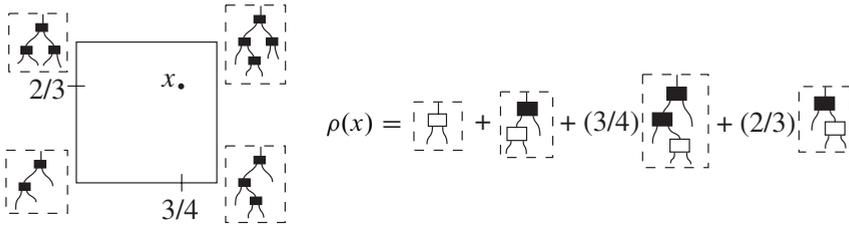


Figure 24. A point x in the cubical complex for Thompson's group F (left) and its image under the map ρ (right).

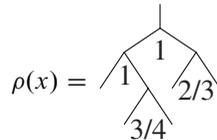


Figure 25. A simpler notation for $\rho(x)$.

of $\text{Im } \rho$ with a picture Δ such that each maximal transistor is labelled by a number $t \in (0, 1]$, and every other transistor is labelled by the number 1. We will usually just omit the label of a transistor T if T is not maximal.

With this convention, an element of F , T , or V acts on a point of $\text{Im } \rho$ by the usual picture multiplication. This action is somewhat tricky to describe if a transistor from the acting element forms a dipole with an exposed transistor labelled by a number $t \neq 1$. In practice, however, we will always be able to avoid considering this situation. If there are no such dipoles, then the action is simple to describe: concatenate and reduce dipoles.

It will be helpful to have a vocabulary for describing subtrees of a given labelled tree T . If T_∞ is the full ordered rooted binary tree of infinite depth, we use binary strings to denote the vertices of degree three in T_∞ . We give each edge in T_∞ a label of 0 or 1; the label is 0 if the edge forms the left half of a caret, and 1 if the edge forms

the right half of a caret. Now label each vertex v in T_∞ by the label of the unique geodesic path from the root to v . For instance, if the geodesic path from the root to v passes through the right half of a caret twice, and then through the left half of a caret, and then finally through the right half of a caret again, then the label of v is 1101. The root has the empty label.

Now if T is an arbitrary labelled subtree (as in Figure 25) of T_∞ and bin is a binary string, we let T_{bin} denote the labelled tree having the vertex bin as its root. For instance, if T is the labelled tree in Figure 25, then T_1 consists of a single caret, labelled by the number 2/3. The tree T_{01} is a single caret labelled by 3/4.

The following partial result will be used in the proof of Theorem 7.1.

Proposition 7.5. *Let $c: [0, \infty) \rightarrow X$ be a geodesic ray, where $X = \tilde{K}(\mathcal{P}, w)$, $\tilde{K}_a(\mathcal{P}, w)$, or $\tilde{K}_b(\mathcal{P}, w)$, $\mathcal{P} = \langle x \mid x = x^2 \rangle$, and $w = x$. Let T be some rooted ordered binary tree in which each transistor is labelled by a 1 (and thus T corresponds to a vertex in $\tilde{K}(\mathcal{P}, w)$, $\tilde{K}_a(\mathcal{P}, w)$, or $\tilde{K}_b(\mathcal{P}, w)$, as the case may be) such that the subtrees T_{10} and T_{11} each contain at least one caret.*

*If $c(t) = T$ for some $t \in [0, \infty)$, then $x_0 * c \neq c$.*

Proof. We can express $\rho(c(t))$ as a tree of the form in Figure 26 a), where each tree $T_i, i = 1, 2, 3$, is a labelled tree, i.e., a picture in which every transistor is an (x, x^2) -transistor, and the trees T_2 and T_3 each have at least one transistor. The effect of the action by x_0 is to transform the tree in Figure 26 a) into the tree in Figure 26 b).

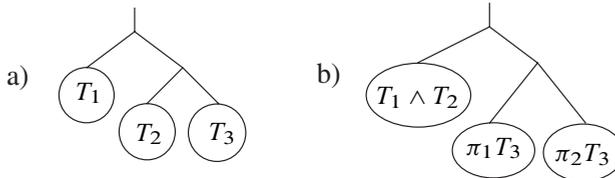


Figure 26. a) The form of $\rho(c(t))$. The trees T_2 and T_3 both contain at least one caret. b) The effect of letting x_0 act on $\rho(c(t))$.

Figure 26 b) also serves as a definition of three different operations on labelled trees. We now make this more explicit. If T is a labelled tree having at least one transistor labelled 1, then the act of removing the topmost transistor of T leaves an ordered pair of trees $(\pi_1 T, \pi_2 T)$. If T_1 and T_2 are labelled trees, then $T_1 \wedge T_2$ is the unique labelled tree satisfying $\pi_1(T_1 \wedge T_2) = T_1$ and $\pi_2(T_1 \wedge T_2) = T_2$.

Note that

$$\|T - x_0 \cdot T\|_2^2 = \|T_1 - T_1 \wedge T_2\|_2^2 + \|T_2 - \pi_1 T_3\|_2^2 + \|T_3 - \pi_2 T_3\|_2^2.$$

Since each caret in each tree is labelled with a 1, each of the terms on the right side of the equation is an integer which counts the number of carets that are in one tree but

not in the other. It follows that $\|T_1 - T_1 \wedge T_2\|_2^2 \geq 2$ and $\|T_3 - \pi_2 T_3\|_2^2 \geq 1$. This implies that

$$d(c(t), x \cdot c(t)) \geq \|T - x_0 \cdot T\|_2 \geq \sqrt{3}.$$

Now we appeal to Lemma 2.3 (2), which implies that if there is $t' < t$ such that $d(c(t'), x_0 \cdot c(t')) < \sqrt{3}$, then c and $x_0 \cdot c$ are not asymptotic, i.e., $c \neq x_0 * c$. We produce a $t' < t$ where $d(c(t'), x_0 \cdot c(t')) = \sqrt{2}$ in Figure 27. \square

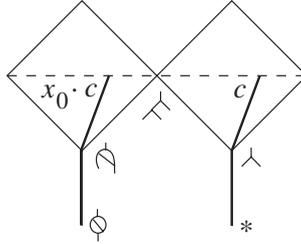


Figure 27. On the right, we have the general picture describing a geodesic ray c issuing from the basepoint $*$. Every such geodesic ray must cross the dotted horizontal line; at the moment it does so, it is precisely $\sqrt{2}$ units distant from its translate $x_0 \cdot c$.

Remark 7.6. (1) It is reasonably clear from the proof that there are variations on this Proposition, in which the hypothesis that T_{10} and T_{11} are non-trivial is replaced by similar assumptions on different subtrees.

(2) Let v be a vertex in the cubical complex for Thompson's group F such that $\rho(v)$ is an ordered labelled rooted binary tree \hat{T} . Note that the coefficient of each caret in \hat{T} is either 1 or 0, since v is a vertex. The above argument shows that

$$d(v, x_0 \cdot v) \geq \sqrt{(1 + (\# \text{ of carets in } \hat{T}_{10})) + (1 + (\# \text{ of carets in } \hat{T}_{11}))}, \quad (*)$$

provided the trees \hat{T}_{10} and \hat{T}_{11} both contain at least one caret.

This suggests a strategy for proving that Δ_∞ contains no fixed points of Thompson's group F . Suppose that c is a geodesic ray having the profile Δ_∞ . It follows from Lemma 4.5 that $\rho(c(t))$ is a labelled ordered rooted binary tree $T(t)$, for any $t \geq 0$. For any given caret $C \in T_\infty$, C occurs in $T(t)$ with the coefficient 1 for t sufficiently large, since c crosses every hyperplane in Δ_∞ . If $c(t)$ passed close to a vertex v for a large value of t , then the inequality (*) shows that $d(v, x_0 \cdot v)$ would be large, so that $d(c(t), x_0 \cdot c(t))$ would also be large, and thus $c \neq x_0 * c$. (Indeed, the proof of Proposition 7.5 shows it suffices to prove that $d(c(t), x_0 \cdot c(t)) > \sqrt{2}$.) Unfortunately, I know of no way to control the distance of $c(t)$ from a vertex, since c will generally travel through cubes of higher and higher dimension, whose diameters go to infinity.

Proof of Theorem 7.1. Let $c \in \Delta_\infty$ be a geodesic ray in either X_T or X_V which represents a fixed point at infinity. We note first that $c(1)$ is the vertex labelled by the finite tree T_1 . If $\pi_1 \in T \subseteq V$ is as in Figure 19 b), it is clear that $\pi_1 \cdot T_1 = T_1$. Now since $(\pi_1 \cdot c)(1) = c(1)$ and $\pi_1 \cdot c, c$ are asymptotic by our assumptions, it follows from Lemma 2.3 (2) that $(\pi_1 \cdot c)(t) = c(t)$ for $t \geq 1$.

Now we refer to Figure 27. The element π_1 flips the diamond on the right along its vertical axis, stabilizing the vertices T_1 and T_2 . Since the geodesic ray travels through this diamond and $(\pi_1 \cdot c)(t) = c(t)$ for $t \geq 1$, it must be that the ray c travels along the straight line connecting T_1 to T_2 , which is vertical in the figure.

Next, we consider the action of $\pi_2 \in T$ (see Figure 20 b)) on c . Note first that π_2 stabilizes the vertex T_2 . It follows from this, Lemma 2.3 (2), the equality $c(1 + \sqrt{2}) = T_2$, and the assumption that π_2 fixes c , that $(\pi_2 \cdot c)(t) = c(t)$ for $t \geq 1 + \sqrt{2}$. Lemma 4.5 implies that $c(1 + \sqrt{2} + \epsilon)$ is in the four dimensional cube C having T_2 as its minimal vertex and T_3 as its maximal vertex. If we identify the cube C with $[0, 1]^4$ in such a way that $(0, 0, 0, 0) = T_2$ and $(1, 1, 1, 1) = T_3$, then π_2 acts by cyclically permuting the coordinates of C . It follows from this that the geodesic ray travels along the diagonal of C from T_2 to T_3 . In particular, c passes through T_3 .

It now follows from Proposition 7.5 that $c \neq x_0 * c$, which is a contradiction. \square

We conclude with an example giving some evidence that there may be fixed points of F in Δ_∞ .

Example 7.7. Figure 28 depicts a labelled tree T (i.e., point in Hilbert space) such that $\|T - x_0 \cdot T\|_2 = \sqrt{7}/2$. The check is left as an exercise. It is not difficult to see that x_0 acts on (most) of the trees along the bottom by a leftward shift. In particular, the tree $T_{1\dots 10}$ (where there are n ones and a single 0) is mapped to $T_{1\dots 10}$ (where there are $n - 1$ ones); the tree $T_{0\dots 01}$ (n zeros) is mapped to $T_{0\dots 01}$ ($n + 1$ zeros). This suggests a principle for building larger trees that are moved only a small distance by x_0 : Begin with the tree $T_{m,m}$ and attach new trees T' to the leaves, making sure that the tree attached at a given leaf is within ϵ of its neighbor to the immediate left, where ϵ will depend on m . The object is to make sure that $\|T - x_0 \cdot T\|_2$ is less than or equal to $\sqrt{2}$, and it is not too difficult to see that this can be done for any m . (A little experimentation also shows that it is useful to label the leftmost and rightmost carets of $T_{m,m}$ with $1/2$.) Moreover, the trees T' that are attached “close” to the root (in a sense that depends on m) can be made arbitrarily large. Thus, we can make a sequence of labelled trees \bar{T}_k , which are each moved less than $\sqrt{2}$ units in the Hilbert metric, and gradually fill up the complete binary tree T_∞ . One can then hope that the corresponding points z_k in X_F are also moved only a small distance by x_0 (as seems likely), so that some subsequence of z_k converges to a point ζ at infinity which is fixed by the action of x_0 . With additional care, it should also be possible to do this

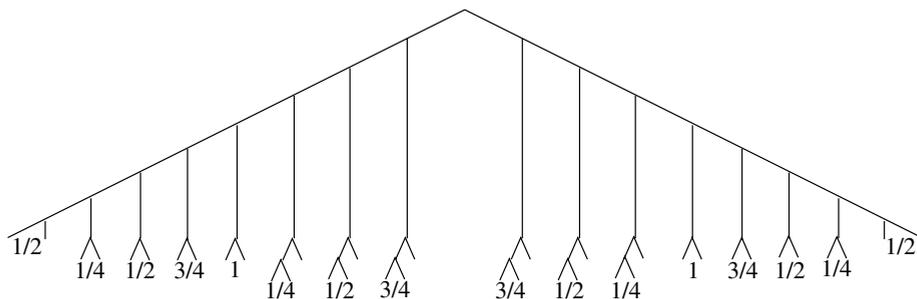


Figure 28. This labelled tree represents a point in X_F which is moved only a small distance by the generators of F . It should be possible to build a sequence of similar, arbitrarily large labelled trees which converge to a fixed point at infinity.

so that each point z_k is likewise moved only a small distance by x_1 , and therefore ζ would be fixed by all of F .

It seems very likely that all of the above can be done. This is not enough, however, because the point ζ may well fail to have the profile Δ_∞ . Indeed, the tree in Figure 28 has the property that most of its norm is contributed by $T_{7,7}$. It appears likely that any tree in the sequence \bar{T}_k will have most of its norm contributed by $T_{n,n}$ (for appropriate n), and this may mean that $\zeta \in \Delta_L \cup \Delta_{L-R} \cup \Delta_R$. I conjecture the following:

Conjecture 7.8. *A point $\zeta \in \partial X_F$ is fixed by all of F if and only if $\zeta \in \Delta_L \cup \Delta_{L-R} \cup \Delta_R$. In particular, the only fixed points for the action of F on ∂X_F lie on an arc of Tits length $\pi/2$.*

References

- [1] S. Adams and W. Ballmann, Amenable isometry groups of Hadamard spaces. *Math. Ann.* **312** (1998), 183–195. [Zbl 0913.53012](#) [MR 1645958](#)
- [2] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer-Verlag, Berlin 1999. [Zbl 0988.53001](#) [MR 1744486](#)
- [3] M. G. Brin, The algebra of strand splitting. I. A braided version of Thompson's group V . *J. Group Theory* **10** (2007), 757–788. [Zbl 05222400](#) [MR 2364825](#)
- [4] M. G. Brin and C. C. Squier, Groups of piecewise linear homeomorphisms of the real line. *Invent. Math.* **79** (1985), 485–498. [Zbl 0563.57022](#) [MR 782231](#)
- [5] J. W. Cannon, W. J. Floyd, and W. R. Parry, Introductory notes on Richard Thompson's groups. *Enseign. Math. (2)* **42** (1996), 215–256. [Zbl 0880.20027](#) [MR 1426438](#)
- [6] J. Crisp and B. Wiest, Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. *Algebr. Geom. Topol.* **4** (2004), 439–472. [Zbl 1057.20028](#) [MR 2077673](#)

- [7] M. M. Day, Amenable semigroups. *Illinois J. Math.* **1** (1957), 509–544. [Zbl 0078.29402](#) [MR 0092128](#)
- [8] D. S. Farley, Finiteness and CAT(0) properties of diagram groups. *Topology* **42** (2003), 1065–1082. [Zbl 1044.20023](#) [MR 1978047](#)
- [9] D. S. Farley, Actions of picture groups on CAT(0) cubical complexes. *Geom. Dedicata* **110** (2005), 221–242. [Zbl 02165498](#) [MR 2136028](#)
- [10] R. Geoghegan (ed.), Open problems in infinite-dimensional topology. In *The proceedings of the 1979 topology conference* (Ohio Univ., Athens, Ohio, 1979), Topology Proc. 4 (1979), no. 1, 287–338. [Zbl 0448.57001](#) [MR 0583711](#)
- [11] R. I. Grigorchuk, Degrees of growth of finitely generated groups and the theory of invariant means. *Izv. Akad. Nauk SSSR Ser. Mat.* **48** (1984), 939–985; English transl. *Math. USSR-Izv.* **25** (1985), 259–300. [Zbl 0583.20023](#) [MR 764305](#)
- [12] R. I. Grigorchuk, An example of a finitely presented amenable group not belonging to the class EG. *Mat. Sb.* **189** (1998), 79–100; English transl. *Sb. Math.* **189** (1998), 75–95. [Zbl 0931.43003](#) [MR 1616436](#)
- [13] V. Guba and M. Sapir, Diagram groups. *Mem. Amer. Math. Soc.* **130** (1997), no. 620. [Zbl 0930.20033](#) [MR 1396957](#)
- [14] D. Guralnik, Coarse decompositions of boundaries for CAT(0) groups. Preprint 2006. [arXiv:math/0611006](#)
- [15] A. Yu. Ol’shanskii, On the problem of the existence of an invariant mean on a group. *Uspekhi Mat. Nauk* **35** (4) (1980), 199–200; English transl. *Russ. Math. Surveys* **35** (4) (1980), 180–181. [Zbl 0465.20030](#) [MR 586204](#)
- [16] A. Yu. Ol’shanskii and M. V. Sapir, Non-amenable finitely presented torsion-by-cyclic groups. *Publ. Math. Inst. Hautes Études Sci.* **96** (2002), 43–169 (2003). [Zbl 1050.20019](#) [MR 1985031](#)
- [17] M. Sageev, Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.* (3) **71** (1995), 585–617. [Zbl 0861.20041](#) [MR 1347406](#)

Received August 9, 2007; revised October 14, 2007

D. Farley, Department of Mathematics, Miami University, Oxford, OH 45056, U.S.A.
E-mail: farleyds@muohio.edu