AGT

Algebraic & Geometric Topology

Volume 26 (2026)

Finiteness properties of some groups of piecewise projective homeomorphisms

DANIEL S. FARLEY

msp

Algebraic & Geometric Topology 26:4 (2026) 1395–1449
DOI: 10.2140/agt.2026.26.1395
Published: April 25, 2026

Finiteness properties of some groups of piecewise projective homeomorphisms

DANIEL S. FARLEY

The Lodha–Moore group G is an F_∞ counterexample to von Neumann's conjecture. The group G acts on the real line via piecewise projective homeomorphisms.

We will describe groups $F(S_i)$, $F(S_i')$, $T(S_i)$, $V(S_i)$, and $V(S_i')$ for $i = 2$ and 3. All of these are groups of piecewise projective homeomorphisms that are modelled on Thompson's groups F, T, and V (respectively); each is "locally determined" by one of four inverse semigroups, which we denote by S_i or S_i' ($i = 2, 3$). Following a method developed by Hughes and the author, we will show that all ten groups have type F_∞.

The Lodha–Moore group G is an ascending HNN extension of $F(S_2')$, and thus our results give a new proof that G has type F_∞.

1 Introduction

Monod [11] produced a large family of counterexamples to von Neumann's conjecture; i.e., nonamenable groups with no free subgroups. Corollary 3 from [11] further noted the existence of finitely generated nonamenable groups with no free subgroups, although the method of proof was nonconstructive. Lodha and Moore [9] considered a subgroup G of one of Monod's groups. Their group G, the *Lodha–Moore group*, could be generated by three elements, and was shown in [9] to be finitely presented; indeed, G admits a presentation with three generators and nine relators. In later work, Lodha [8] showed that G has type F_∞. The Lodha–Moore group is also nonamenable, making it an especially economical

MSC2020: primary 20F65, 20J05; secondary 20M18.

finitely presented counterexample to von Neumann's conjecture, and the first F_∞ counterexample to von Neumann's conjecture. (The paper [13] provided the first finitely presented nonamenable groups with no free subgroups, but the groups in question had many more generators and relations. The higher finiteness properties of the groups from [13] remain unknown to the best of this author's knowledge.)

In [6], Hughes and the author described a general approach to establishing finiteness properties for generalised Thompson groups with "piecewise" definitions. The basic theory of [6] was produced in the hope that it would unify the existing proofs of finiteness properties for such groups. The Lodha–Moore group G offers a useful test case.

We now briefly recall the methods of [6]. In the setting of [6], an inverse semigroup S is a set of partial bijections of some set X that is closed under compositions and inverses. By a *partial bijection* of X, we mean a bijection between two subsets of X. We define a group Γ_S, the *group locally determined by* S, to be the collection of all bijections of X that are finite unions of partial bijections from S. (I.e., $\gamma \in \Gamma_S$ if, for some $n \in \mathbb{N}$, there are elements $s_1, s_2, \dots, s_n \in S$ such that the domains of the s_i, denoted by D_i, form a partition of X, the images $s_i(D_i)$ are also a partition of X, and $\gamma_{|D_i} = s_i$, for $i = 1, \dots, n$.)

The construction of classifying spaces for the groups Γ_S depends upon a sequence of choices. The choice of S determines a collection of *domains* $\mathcal{D}_S$, which are simply the domains of the elements $s \in S$. The set of nonempty domains is denoted by $\mathcal{D}_S^+$. The second choice is that of an S-*structure*, which is a function $\mathbb{S} : \mathcal{D}_S^+ \times \mathcal{D}_S^+ \to \mathcal{P}(S)$ assigning a (possibly empty) collection of transformations from S to each pair (D_1, D_2) of nonempty domains. The sets $\mathbb{S}(D_1, D_2)$ are required to satisfy various "groupoid-like" properties. The specific properties that we need are summarised in Proposition 4.4. The S-structure $\mathbb{S}$ determines a set $\mathcal{V}_{\mathbb{S}}$ with a partial order $\leq$ called "expansion" (Definition 4.7). The expansion partial order is, roughly speaking, determined by the subdivision of a given domain into (finitely many) smaller domains. Under appropriate hypotheses, the partially ordered set $(\mathcal{V}_{\mathbb{S}}, \leq)$ becomes a directed Γ_S-set. The simplicial realisation $\Delta_{\mathbb{S}}$ of $\mathcal{V}_{\mathbb{S}}$ is therefore a contractible simplicial complex upon which Γ_S acts simplicially. The construction of $\mathcal{V}_{\mathbb{S}}$ is very much like the one introduced by Brown [3], where he proved the F_∞ property for a wide variety of generalised Thompson groups.

The simplicial complex $\Delta_{\mathbb{S}}$ often has undesirable properties, however. For instance, it almost always fails to be locally finite. It can be helpful to replace $\Delta_{\mathbb{S}}$ with something smaller. A (third) choice of an *expansion scheme* $\mathcal{E}$ (Definition 6.15) determines a subcomplex $\Delta_{\mathbb{S}}^{\mathcal{E}} \subset \Delta_{\mathbb{S}}$. (The complex $\Delta_{\mathbb{S}}^{\mathcal{E}}$ should be thought of as a generalisation of the complexes introduced by Stein [15], who created locally finite models for various groups of piecewise linear homeomorphisms of the line, including Thompson's group F.) The complex $\Delta_{\mathbb{S}}^{\mathcal{E}}$ can be anything from a discrete set of points to the complex $\Delta_{\mathbb{S}}$ itself, depending on the size of the expansion scheme. Given an appropriate choice of $\mathcal{E}$, the main results of [6] show how to deduce the F_∞ property for Γ_S, by applying Brown's finiteness criterion to the complex $\Delta_{\mathbb{S}}^{\mathcal{E}}$.

The main goal of this paper is to describe a family of ten groups of piecewise projective homeomorphisms of $[0, 1)$, S^1, and $[0, \infty)$, and prove that each group has type F_∞. We will also describe the Lodha–Moore group G as an ascending HNN extension of one of the groups, and therefore obtain a new proof that G has type F_∞. Our approach follows the method from [6]:

(1) First, we consider four inverse semigroups: S_2, S_3, S_2', and S_3'. The generators of S_i, for $i = 2, 3$, are A, B, and C_i, where the domain of each transformation is the interval $I = [0, 1)$, and

$$A(x) = \frac{x}{x+1}; \quad B(x) = \frac{1}{2-x}; \quad C_2(x) = \frac{2x}{x+1}; \quad C_3(x) = \frac{3x}{2x+1}.$$

The inverse semigroups S_2' and S_3' have the additional generator $T : [0, \infty) \to [1, \infty)$, defined by the rule $T(x) = x + 1$. The groups

$$F(S_i), \quad F(S_i'), \quad T(S_i), \quad V(S_i), \quad V(S_i')$$

are then defined to be the groups that are "locally determined" (in the sense of [6]) by the semigroup S_i or S_i'. The "F" groups are homeomorphisms of the line or interval, the "T" groups are homeomorphisms of the circle, and the "V" groups are groups of right-continuous bijections. (The notation is intended to recall the definitions of Thompson's group F, T, and V, as presented in (for instance) [5].)

(2) The domains of S_i and S_i' are not very tractable for our purposes. We will therefore restrict the domains under consideration to what we call a set of "generating domains" $\mathcal{D}^+_{\text{gen}}$ (Definition 3.1), which, for us, are simply the forward iterates of $I = [0, 1)$ under the transformations A and B. (The decision to work with a proper subset of $\mathcal{D}^+_S$ represents the most important departure from [6].)

For every pair of domains (D_1, D_2), we then define $\mathbb{S}(D_1, D_2)$ as the set of all transformations from S_i or S_i' having D_1 as the domain and D_2 as the range, where D_1 and D_2 are arbitrary members of $\mathcal{D}^+_{\text{gen}}$. The sets $\mathbb{S}(D_1, D_2)$ enable us to define an "expansion" operation. The expansions (Definition 4.7) from the pair $[\mathrm{id}_I, I]$ (Definition 4.5) can be usefully described by numbered binary trees, exactly as was done in [9].

(3) The directed sets from (2) are too large. In search of more tractable complexes, we define the expansion schemes $\mathcal{E}_i$ and $\mathcal{E}_i'$ as in Example 6.22. The expansion schemes systematically restrict the types of expansions that are allowed in the complexes $\Delta^{\mathcal{E}_i}_{\mathbb{S}}$ and $\Delta^{\mathcal{E}_i'}_{\mathbb{S}}$. The ascending stars in the resulting complexes are isomorphic to products of a simplicial cone on a cellulated line — see Figure 5 on page 1420. The burden of the rest of the argument is to show that $\mathcal{E}_i$ and $\mathcal{E}_i'$ are "n-connected expansion schemes" (Definition 6.18). The proof occupies Sections 7 and 8, and represents the technical heart of the paper. The material from Sections 7 and 8 is heavily indebted to the argument from [9], but generalises that work in what we consider to be interesting ways. For instance, our argument also shows that the monoid generated by the linear fractional transformations $\{A, B, C_2, c_2\}$ (even without the above restrictions on their domains) admits a finite complete rewrite system.

With (1)–(3) complete, the proof that the groups $F(S_i)$, $F(S_i')$, $T(S_i)$, $V(S_i)$, and $V(S_i')$ $(i = 2, 3)$ all have type F_∞ follows by a standard argument (a variant of the main argument of [6]). This standard argument is summarised in Section 9. The relationship between $F(S_2')$ and G is described in Section 10; specifically, we show that G is an ascending HNN extension of $F(S_2')$, which directly implies that G has type F_∞.

The author had originally wanted to offer a proof of F_∞ for an infinite family of generalisations of the Lodha–Moore group G. One approach to producing an infinite family of similar groups is, for a given n, to replace the transformation C_2 with a transformation $C_n : I \to I$, defined as

$$C_n(x) = \frac{nx}{(n-1)x+1}.$$

We can then define S_n to be the inverse semigroup generated by $\{A, B, C_n\}$ (and their inverses), and define the groups $F(S_n)$, $F(S_n')$, $T(S_n)$, $V(S_n)$, and $V(S_n')$ for arbitrary n. We find, however, that our method fails for $n \geq 4$; in fact, we are not even able to build useful directed sets (along the lines sketched for $n = 2, 3$). Our difficulties are summarised in Section 5.

Let us briefly describe the structure of the paper. In Section 2, we define the inverse semigroups S_i and S_i' $(i = 2, 3)$ and the groups that are locally defined by these semigroups. In Section 3, we define the "generating domains" that we need. This section also includes a proof of the "eventual invariance" property, which will be used to construct directed sets later. In Section 4, we build directed sets with Γ actions (for Γ as above) and compute vertex stabilisers. We show, in particular, that the vertex stabilisers in all of our complexes are virtually free abelian of finite rank. In Section 5, we describe an algorithm that analyses various potential generalisations of our piecewise projective homeomorphism groups. The conclusion of the section is that such generalisations are surprisingly very thin on the ground. In Section 6, we review expansion schemes, introduce expansion schemes $\mathcal{E}_i$ and $\mathcal{E}_i'$ $(i = 2, 3)$, and also introduce subdivision trees, which are used to describe expansions.This section also describes an equivalence relation on subdivision trees. In Section 7, we compute finite complete semigroup presentations for the inverse semigroups S_2 and S_3. These presentations are vital in understanding the equivalence relation on subdivision trees. In Section 8, we prove an "intermediate value theorem", which is what we need in order to show that the expansion schemes $\mathcal{E}_i$ and $\mathcal{E}_i'$ define contractible complexes. The proof of the latter is in Section 9, which also assembles all of the other ingredients of the proof that the groups have type F_∞. Section 10 establishes the connection between G and $F(S_2')$.

2 A family of inverse semigroups

We consider the usual action of $\mathrm{PSL}_2(\mathbb{R})$ on the upper half-space model of the hyperbolic plane $\mathbb{H}^2 \subseteq \mathbb{C}$. A 2×2 matrix

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

acts as a linear fractional transformation $f_M : \mathbb{H}^2 \to \mathbb{H}^2$, where

$$f_M(z) = \frac{az+b}{cz+d} \quad (z \in \mathbb{C}; \mathrm{Im} z > 0).$$

It is well-known [14] that the assignment $M \mapsto f_M$ induces an isomorphism between $\mathrm{PSL}_2(\mathbb{R})$ and the group of all orientation-preserving isometries of $\mathbb{H}^2$, denoted by $\mathrm{Isom}^+(\mathbb{H}^2)$. In what follows, we will make no distinction between M and f_M, referring to either one by the matrix M.

In practice, we will be concerned primarily with the action of $\mathrm{PSL}_2(\mathbb{R})$ on $\partial\mathbb{H}^2$, which we identify with $\mathbb{R}\cup\{\infty\}$. The inverse semigroups alluded to in this section's title act as partial bijections of $\partial\mathbb{H}^2$ via (restrictions of) linear fractional transformations.

Definition 2.1 (partial bijections; inverse semigroups; domains) Let X be a set. A *partial bijection* of X is a bijection $h: A_h \to B_h$ between subsets A_h and B_h of X. The composition of two partial bijections is defined on "overlaps": if $g: A_g \to B_g$ and $h: A_h \to B_h$ are partial bijections of X, then $g\circ h$ is a bijection from $h^{-1}(A_g)$ to $g(B_h\cap A_g)$.

A collection S of partial bijections of X is called an *inverse semigroup* if S is closed under inverses and compositions. We may also refer to such an S as an *inverse semigroup acting on* S.

If S is an inverse semigroup and $h: A_h \to B_h$, then we refer to A_h as a *domain* of S. Note that B_h is also a domain, since S is closed under inverses. We let $\mathcal{D}_S$ denote the set of all domains A_h, as h ranges over all $h\in S$. We let $\mathcal{D}_S^+ = \mathcal{D}_S - \{\varnothing\}$ (i.e., the set of all nonempty domains).

Remark 2.2 Let S be an inverse semigroup. We note two basic properties:

(1) If $s\in S$ and D is a domain of S that is contained in the domain of s, then $s_{|D}\in S$. Indeed, let D be the domain of $t\in S$. Then $s_{|D} = st^{-1}t$.

(2) If D is a domain of S, then $\mathrm{id}_D\in S$. Indeed, if $t\in S$ has D as its domain, then $t^{-1}t=\mathrm{id}_D$.

(3) If D_1 and D_2 are domains of S, then $D_1\cap D_2$ is also a domain of S. Indeed, letting D_1 be the domain of $t_1\in S$, and D_2 be the domain of $t_2\in S$, we find that the domain of $t_1^{-1}t_1t_2^{-1}t_2$ is $D_1\cap D_2$.

Remark 2.3 Inverse semigroups can also be defined abstractly (see [7]). The Preston–Wagner theorem states that any inverse semigroup can be realized as a collection of partial bijections (in the above sense). The proof parallels that of Cayley's theorem, which states that every group can be realized as a group of permutations.

Remark 2.4 Some readers may be familiar with the theory of étale groupoids. Inverse semigroups seem to be closely related (see, for instance, [4], where a form of equivalence between étale groupoids and inverse semigroups is established). The precise nature of the relationship between the methods of [6] (and this paper) and the broader literature of étale groupoids is unclear to the author, who has no expertise in the latter area.

Remark 2.5 In Definition 2.1, the function with empty domain and codomain plays the role of a 0.

Definition 2.6 (the inverse semigroups S_n and S_n') Let $A:[0,1)\to[0,1/2)$ be the restriction of the linear fractional transformation

$$A=\begin{pmatrix}1&0\\1&1\end{pmatrix}.$$

Let $B:[0,1)\to[1/2,1)$ be the restriction of

$$B=\begin{pmatrix}0&1\\-1&2\end{pmatrix}.$$

For $n = 2, 3$, let $C_n : [0, 1) \to [0, 1)$ be the restriction of

$$C_n = \begin{pmatrix} n & 0 \\ n-1 & 1 \end{pmatrix}.$$

Let $T : [0, \infty) \to [1, \infty)$ be the restriction of

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

We will use lower-case letters to denote the inverses of the above transformations ($a = A^{-1}$, etc.).

For $n = 2, 3$, we let

$$S_n = \langle A, B, C_n \rangle; \quad S'_n = \langle A, B, C_n, T \rangle,$$

where the brackets indicate the inverse semigroup generated by the bracketed transformations; i.e., S_n and S'_n are closed under compositions and inverses.

Definition 2.7 (locally determined by S; the inverse semigroup $\widehat{S}$) Let S be an inverse semigroup acting on a set X. Let A and B be subsets of X. A bijective function $f : A \to B$ is *locally determined by* S if there is a finite partition $\mathcal{P} = \{D_1, \dots, D_m\}$ of A into domains (i.e., $\mathcal{P} \subseteq \mathcal{D}^+_S$) such that $f_{|D_i} \in S$, for each i.

We let $\widehat{S}$ denote the collection of all functions that are locally determined by S. The set $\widehat{S}$ is an inverse semigroup under the operation of composition.

Definition 2.8 Let $n = 2$ or 3. Let

- $F(S_n)$ be the group of homeomorphisms of $[0, 1)$ that are locally determined by S_n;
- $F(S'_n)$ be the group of homeomorphisms of $[0, \infty)$ that are locally determined by S'_n;
- $T(S_n)$ be the group of homeomorphisms of the circle $[0, 1]/\sim$ that are locally determined by S_n;
- $V(S_n)$ be the group of right-continuous bijections of $[0, 1)$ that are locally determined by S_n;
- $V(S'_n)$ be the group of right-continuous bijections of $[0, \infty)$ that are locally determined by S'_n.

Remark 2.9 The group $T(S_n)$ can also be described as the subgroup of $V(S_n)$ that preserves a cyclic ordering on $[0, 1)$.

3 A generating set of domains for S_i and S'_i

The set $\mathcal{D}^+_S$ (Definition 2.1) will be far too big when $S = S_i$ or $S = S'_i$. In this section, we define a subcollection $\mathcal{D}^+_{S,\mathrm{gen}} \subseteq \mathcal{D}^+_S$, which will be sufficient for the constructions of later sections. We note that this is in contrast with [6], which always uses the full set $\mathcal{D}^+_S$.

Definition 3.1 (generating sets of domains) Let $\{A, B\}^*$ denote the set of all positive words in the alphabet $\{A, B\}$, including the empty word. Let

$$\mathcal{D}^+_{S',\mathrm{gen}} = \{T^\alpha \omega \cdot [0, 1) \mid \omega \in \{A, B\}^*; \alpha \geq 0\} \cup \{T^\alpha \cdot [0, \infty) \mid \alpha \geq 0\}$$

and

$$\mathcal{D}^{+}_{S,\text{gen}} = \{\omega \cdot [0,1) \mid \omega \in \{A, B\}^*\}.$$

We will often refer to $\mathcal{D}^{+}_{S,\text{gen}}$ or $\mathcal{D}^{+}_{S',\text{gen}}$ by the notation $\mathcal{D}^{+}_{\text{gen}}$ if doing so should cause no ambiguity.

We may sometimes refer to the members of $\mathcal{D}^{+}_{S,\text{gen}}$ and $\mathcal{D}^{+}_{S',\text{gen}}$ as *generating domains*.

Remark 3.2 It will be convenient to write I in place of $[0,1)$, and to write ωI in place of $\omega \cdot I$.

The half-open intervals ωI of $\mathcal{D}^{+}_{S,\text{gen}}$ are in one-to-one correspondence with the vertices of an infinite binary tree. The intervals $\omega A I$ and $\omega B I$ correspond to the left and right children (respectively) of ωI. In particular, $\omega' I$ contains ωI if and only if ω' is a prefix of ω, and the intervals are disjoint if neither ω' nor ω is a prefix of the other.

Note, however, that the intervals ωI are very far from being the standard dyadic intervals when the length of ω is two or more. For instance, $ABAI = [1/3, 2/5)$ and $BAI = [1/2, 2/3)$. It appears that ωI, for $\omega \in \{A, B\}^*$, is always an interval between consecutive Farey fractions (as noted in [9]), although we will not need to use this fact. The intervals $T^{\alpha}\omega I$ are simply the translates of the intervals ωI by nonnegative integers.

It will be useful to keep in mind that the products aB and bA are 0 in what follows.

Lemma 3.3 (an eventual invariance property) *Let $s \in S_i$, where $i = 2$ or 3. Let $D \in \mathcal{D}^{+}_{S,\text{gen}}$ be such that D is contained in the domain of s. There is a finite partition $\mathcal{P} \subseteq \mathcal{D}^{+}_{S,\text{gen}}$ of D such that $sP \in \mathcal{D}^{+}_{S,\text{gen}}$, for each $P \in \mathcal{P}$.*

The analogous statement also holds true for S'_n, $i = 2, 3$.

Proof Let $D = \omega I$, where $\omega \in \{A, B\}^*$. By induction on the length of s, it suffices to prove the lemma in the case $s \in \{A, B, C_2, C_3, a, b, c_2, c_3\}$.

Suppose first that $s = A$. It follows directly that $sD = A\omega I$, so $sD \in \mathcal{D}^{+}_{S,\text{gen}}$ and we may set $\mathcal{P} = \{D\}$. If $s = a$, it must be that $\omega = A\omega'$, for some $\omega' \in \{A, B\}^*$, and therefore $sD = \omega' I$. We can therefore again let $\mathcal{P} = \{D\}$.

If $s = B$ or $s = b$, the proof is very similar.

Let $s = C_2$. A straightforward check shows that

$$\begin{aligned} C_2AA &= AC_2; \\ C_2AB &= BAc_2; \\ C_2B &= BBC_2. \end{aligned}$$

We use these identities to "push" C_2 as close to the end of the word ω as possible. (The inverse c_2 can appear during this process. This poses no problems, since we can use the same identities to push c_2 forward as well.) In doing so, we can arrange that

$$C_2\omega = \omega' C_2^{\epsilon} \omega'',$$

where $\omega', \omega'' \in \{A, B\}^*$, $\epsilon = \pm 1$, and

(1) ω'' is an empty word, or

(2) $\omega'' = A$ if $\epsilon = 1$, or

(3) $\omega'' = B$ if $\epsilon = -1$.

If ω'' is not the empty word, we then let $\mathcal{P} = \{\omega AI, \omega BI\}$; if ω'' is empty, we set $\mathcal{P} = \{D\}$; these are the required partitions. (For instance, if $\epsilon = 1$ and $\omega'' = A$, we have

$$C_2\omega AI = \omega' C_2 AAI = \omega' AC_2 I = \omega' AI \in \mathcal{D}^+_{\text{gen}}$$

and

$$C_2\omega BI = \omega' C_2 ABI = \omega' BAc_2 I = \omega' BAI \in \mathcal{D}^+_{\text{gen}}.$$

Similar checking handles the remaining case.)

The case in which $s = c_2$ is very similar, and features the same identities, suitably rewritten so that c_2A, c_2BA, and c_2BB appear on the left-hand sides of the equations.

Now suppose that $s = C_3$. We have the matrix identities

$$\begin{aligned} C_3AAA &= AC_3; \\ C_3AAB &= BAAc_3; \\ C_3ABA &= BABC_3; \\ C_3ABB &= BBAc_3; \\ C_3B &= BBBC_3. \end{aligned}$$

We can then follow the same strategy as we did in the case $s = C_2$. "Push" C_3 as close to the end of ω as possible. The result is $\omega' C_3^{\epsilon} \omega''$, where $\omega' \in \{A, B\}^*$ and

(1) ω'' is empty, or

(2) $\epsilon = 1$ and $\omega'' \in \{A, AA, AB\}$, or

(3) $\epsilon = -1$ and $\omega'' \in \{B, BA, BB\}$.

If ω'' is empty, then $\mathcal{P} = \{D\}$ is the required partition of D. If $\epsilon = 1$ and $\omega'' = A$, then we set $\mathcal{P} = \{\omega AAI, \omega ABI, \omega BAI, \omega BBI\}$. This is the required partition; indeed,

$$C_3\omega AAI = \omega' C_3 AAAI = \omega' AC_3 I = \omega' AI \in \mathcal{D}^+_{\text{gen}},$$

and similar calculations show that $C_3(\mathcal{P}) \subseteq \mathcal{D}^+_{\text{gen}}$. If $\omega'' \in \{AA, AB\}$, then the required partition is $\mathcal{P} = \{\omega AI, \omega BI\}$. If $\epsilon = -1$ and $\omega'' \in \{B, BA, BB\}$, then one proceeds similarly. The required partitions are $\{\omega AAI, \omega ABI, \omega BAI, \omega BBI\}$ (in the first case, when $\omega'' = B$) and $\{\omega AI, \omega BI\}$ (when $\omega'' = BA$ or BB).

The extension of these arguments to S_i' is straightforward, but we consider the case of S_2' by way of example. In this case, the domain D has one of the forms

$$T^{\alpha}\omega \cdot [0, 1) \quad \text{or} \quad T^{\alpha} \cdot [0, \infty),$$

where $\omega \in \{A, B, C, a, b, c\}^*$ and $\alpha \geq 0$. We note that, in the former case, $D \subseteq [\alpha, \alpha + 1)$, while $D = [\alpha, \infty)$ in the latter case. The semigroup element $s \in S_2'$ has the form

$$T^m \widehat{\omega} T^{-n},$$

for some word $\widehat{\omega} \in \{A, B, C, a, b, c\}^*$, where m and n are nonnegative. (The proof is as follows. We know that s is not the zero element by hypothesis. The products XT and $T^{-1}X$ are always zero when $X \in \{A, B, C, a, b, c\}$. We also have the identity $T^{-1}T = \mathrm{id}_{[0,\infty)}$. It follows that, after suitable reductions, no occurrence of T^{-1} can occur immediately before a different generator, and no occurrence of T can occur immediately after a different generator. Thus, the given form describes the only possibilities.)

Now we consider cases. If $D = [\alpha, \infty)$, then s necessarily has the form $T^m T^{-n}$ by domain considerations. (If $\widehat{\omega} \neq 1$, then the domain in question is an interval of finite length. This is impossible, since D must be contained in the domain of s.) It follows that we can simply let $\mathcal{P} = \{D\}$. Suppose that $D = T^\alpha \omega \cdot [0, 1)$; we must consider the possible cases for n. It is not possible for n to be greater than α, since this would mean that the domain of s is contained in $[n, \infty)$, and thus result in D not being a subset of the domain of s. If $n < \alpha$, then s must take the form $T^m T^{-n}$ (otherwise, if $\widehat{\omega} \neq 1$, we would conclude that the domain of s is contained in $[n, n + 1)$, which is disjoint from D). We apply the lemma to the case of S_2, temporarily letting $s = \mathrm{id}_{[0,1)}$ and $D = \omega \cdot [0, 1)$, to find that there is partition $\widehat{\mathcal{P}}$ of $\omega \cdot [0, 1)$ into generating domains. We can then let $\mathcal{P} = T^\alpha \cdot \widehat{\mathcal{P}}$. Noting that the property of being a generating domain is unchanged after an application of T, we see that this is the required partition. If $n = \alpha$, then we apply the lemma with $s = \widehat{\omega}$ and $D = \omega \cdot [0, 1)$ to get a partition $\widehat{\mathcal{P}}$ of D into generating domains with the additional property that $\widehat{\omega} \cdot P$ is a generating domain, for each $P \in \widehat{\mathcal{P}}$. It then follows that $\mathcal{P} = T^\alpha \cdot \widehat{\mathcal{P}}$ is the required partition. □

Remark 3.4 A further application of Lemma 3.3 is that every domain $D \in \mathcal{D}_S^+$ or $\mathcal{D}_{S'}^+$ can be partitioned into finitely many generating domains. Indeed, each such D is the domain of some word ω in the generators of S_i or S_i'. We can then show that D can be partitioned into generating domains by induction on the length of ω. The base case (in which the length of ω is 1) is trivial, since all of the domains in question are necessarily generating domains. The inductive step is then handled by applying Lemma 3.3.

4 A directed set construction

In this section, we will specify an S-structure $\mathbb{S}$ for $S \in \{S_2, S_3, S_2', S_3'\}$, in essentially the sense of [6]. In fact, the only difference is that we will define our S-structure using the domains $\mathcal{D}_{S,\mathrm{gen}}^+$ and $\mathcal{D}_{S',\mathrm{gen}}^+$, rather than the entire collection $\mathcal{D}^+$, as required in [6]. The S-structure leads to a directed set construction of a contractible simplicial complex, exactly as in [6]. We will first consider these directed set constructions for $V(S_i)$ and $V(S_i')$ $(i = 2, 3)$ in Section 4.1. The simplicial complexes for the related groups can then be obtained as subcomplexes; this is spelled out in Section 4.3.

It will be useful to let $\mathcal{D}_{\mathrm{gen}}^+$ denote either $\mathcal{D}_{S,\mathrm{gen}}^+$ or $\mathcal{D}_{S',\mathrm{gen}}^+$, depending on the context.

4.1 The directed set constructions for $V(S_i)$ and $V(S_i')$

We will first show how to make $V(S_i)$ and $V(S_i')$ act on directed sets, and (therefore) on contractible simplicial complexes. The basic approach follows [6], but we are able to use a simplified version of the basic theory, with suitable modifications. All of the results in this subsection work in the same way for all $\Gamma \in \{V(S) \mid S \in \{S_2, S_3, S_2', S_3'\}\}$, so we will use the generic notation Γ to refer to any group from the latter collection.

Definition 4.1 (structure sets; domain types) Let $S \in \{S_2, S_3, S_2', S_3'\}$. Let $D_1, D_2 \in \mathcal{D}^+_{\text{gen}}$. We set

$$\mathbb{S}(D_1, D_2) = \{s \in S \mid \text{the domain of } s \text{ is } D_1 \text{ and the range is } D_2\}.$$

Two domains D_1 and D_2 have the *same type* if $\mathbb{S}(D_1, D_2) \neq \varnothing$; i.e., if there is some $s \in S$ such that D_1 is the domain of s and D_2 is the image of s.

Remark 4.2 (description of domain types) There are one or two domain types, depending on whether $S \in \{S_2, S_3\}$, on the one hand, or $S \in \{S_2', S_3'\}$, on the other. The first of the domain types consists of those sets of the form ωI, where $\omega \in \{A, B\}^*$. This domain type occurs in both S_i and S_i', and it is the only type if $S \in \{S_2, S_3\}$. The second of the domain types (present only when $S \in \{S_2', S_3'\}$) consists of domains of the form $[n, \infty)$, where n is a nonnegative integer.

Theorem 4.3 (explicit description of structure sets) *Let $S = S_2$ or S_3. Given ωI and $\omega' I \in \mathcal{D}^+_{\text{gen}}$ ($\omega, \omega' \in \{A, B\}^*$), the associated structure set takes the form*

$$\mathbb{S}(\omega I, \omega' I) = \{\omega' C^k \omega^{-1} \mid k \in \mathbb{Z}\}.$$

Let $S = S_2'$ or S_3'. Given ωI and $\omega' I \in \mathcal{D}^+_{\text{gen}}$ ($\omega, \omega' \in \{A, B, T\}^*$), the associated structure set takes the form

$$\mathbb{S}(\omega I, \omega' I) = \{\omega' C^k \omega^{-1} \mid k \in \mathbb{Z}\}.$$

The set $\mathbb{S}([m, \infty), [n, \infty))$ is $\{T^{n-m}\}$, when m, n are nonnegative integers.

Proof We first consider the case of $\mathbb{S}(I, I)$ when $S = S_2$.

Let $\omega \in \mathbb{S}(I, I)$. Thus, $\omega I = I$. Let G be the group generated by the linear fractional transformations A, B, and C, each viewed as a transformation of $\mathbb{H}^2$ or the projective line $\mathbb{R} \cup \{\infty\}$. We note that the inverses of A, B, and C may be represented by the matrices

$$a = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}; \quad b = \begin{pmatrix} 2 & -1 \\ 1 & 0 \end{pmatrix}; \quad c = \begin{pmatrix} 1 & 0 \\ -1 & 2 \end{pmatrix}.$$

It follows that, if ω is expressed as a product of matrices, then $\det(\omega) = 2^n$, for some nonnegative integer n. We note also that ω fixes the points 0 and 1 on the projective line, by our assumptions.

We let

$$\omega = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix},$$

where α, β, γ, and δ are integers. The equality $\omega(0) = 0$ directly implies that $\beta = 0$. The equality $\omega(1) = 1$ then implies that $\alpha = \gamma + \delta$. Computing determinants, we find that

$$(\gamma + \delta)\delta = 2^n.$$

It follows that $\gamma + \delta = 2^k$ and $\delta = 2^\ell$, where $k + \ell = n$ and k and ℓ are nonnegative integers. Either $k \le \ell$ or $\ell \le k$; in the first case,

$$\begin{pmatrix} 2^k & 0 \\ 2^k - 2^\ell & 2^\ell \end{pmatrix} \sim \begin{pmatrix} 1 & 0 \\ 1 - 2^{\ell-k} & 2^{\ell-k} \end{pmatrix} = c^{\ell-k}.$$

Similarly, if $\ell \le k$, then $\omega = C^{k-\ell}$. In either case, $\omega = C^\alpha$, for appropriate α.

If $S = S_3$, then the set $\mathbb{S}(I, I)$ takes the same form. Here

$$C = \begin{pmatrix} 3 & 0 \\ 2 & 1 \end{pmatrix}; \quad c = \begin{pmatrix} 1 & 0 \\ -2 & 3 \end{pmatrix},$$

and therefore $\det(\omega) = 3^n$, for some $n \in \mathbb{Z}$. The remainder of the argument differs from the case of S_2 primarily in the fact that it involves powers of 3, rather than powers of 2.

Now we consider a general structure set $\mathbb{S}(\omega I, \omega' I)$, where $\omega, \omega' \in \{A, B\}^*$ and S may be any of the semigroups S_2, S_3, S_2', or S_3'. Let $\sigma \in \mathbb{S}(\omega I, \omega' I)$. It follows that $(\omega')^{-1}\sigma\omega \in \mathbb{S}(I, I)$, so $(\omega')^{-1}\sigma\omega = C^k$, for some $k \in \mathbb{Z}$. Thus, $\sigma = \omega' C^k \omega^{-1}$, as claimed. Conversely, it is clear that any transformation of the form $\omega' C^k \omega^{-1}$ is in $\mathbb{S}(\omega I, \omega' I)$.

The final statement, about $\mathbb{S}([m, \infty), [n, \infty))$ follows from the fact that T^{n-m} is the only inverse semigroup element of S_i' with the given domain and codomain. □

Proposition 4.4 (closure properties of $\mathbb{S}$) *Let $D_1, D_2 \in \mathcal{D}^+_{\text{gen}}$.*

(1) ***Compositions*** *If $h \in \mathbb{S}(D_1, D_2)$ and $g \in \mathbb{S}(D_2, D_3)$, then $gh \in \mathbb{S}(D_1, D_3)$.*

(2) ***Inverses*** *If $h \in \mathbb{S}(D_1, D_2)$, then $h^{-1} \in \mathbb{S}(D_2, D_1)$.*

(3) ***Identities*** $\mathrm{id}_{D_1} \in \mathbb{S}(D_1, D_1)$.

Proof All of these properties follow directly from the definition of the set $\mathbb{S}(D_1, D_2)$. □

Definition 4.5 (the set $\mathcal{B}$) Let

$$\mathcal{A} = \{(f, D) \mid f \in \widehat{S}; D \in \mathcal{D}^+_{\text{gen}}; D \text{ is contained in the domain of } f\}.$$

(Recall that $\widehat{S}$ is the inverse semigroup of functions that are locally determined by S (Definition 2.7).) We write $(f_1, D_1) \sim (f_2, D_2)$ if there is some $h \in \mathbb{S}(D_1, D_2)$ such that $f_1 = f_2 h$. It is easily checked that $\sim$ is an equivalence relation on $\mathcal{A}$, using Proposition 4.4. We let $\mathcal{B}$ denote the set of all equivalence classes. The equivalence class of (f, D) will be denoted by $[f, D]$.

Definition 4.6 (vertices; the type of a vertex) A finite subset

$$\{[f_1, D_1], \ldots, [f_m, D_m]\} \subseteq \mathcal{B}$$

is a *vertex* if

$$\bigcup_{i=1}^{m} f_i(D_i) = \mathbb{R}^+ \quad \text{or} \quad \bigcup_{i=1}^{m} f_i(D_i) = [0,1),$$

depending upon whether the underlying semigroup is S'_n or S_n, respectively. (Here $\mathbb{R}^+$ is the set of nonnegative real numbers and m may be any natural number.)

We let $\mathcal{V}_S$ denote the set of all vertices, where $S \in \{S_2, S_3, S'_2, S'_3\}$. We may sometimes write $\mathcal{V}$ in place of $\mathcal{V}_S$ if this will result in no ambiguity.

Two vertices $\{[f_1, D_1], \dots, [f_m, D_m]\}$ and $\{[g_1, E_1], \dots, [g_n, E_n]\}$ have the *same type* if the multisets $\{[D_1], \dots, [D_m]\}$ and $\{[E_1], \dots, [E_n]\}$ are identical; i.e., $m = n$ and $[D_j] = [E_j]$, for $j = 1, \dots, m$.

Definition 4.7 (expansion; contraction) Let $v = \{[f_1, D_1], \dots, [f_n, D_n]\}$ be a vertex. We say that a vertex v' is obtained from v by *expansion at* $[f_i, D_i]$ if there is some $h \in \mathbb{S}(D_i, D_i)$ and a finite partition $\mathcal{P} \subseteq \mathcal{D}^+_{\text{gen}}$ of D_i into domains such that

$$v' = (v - \{[f_i, D_i]\}) \cup \{[f_i h, P] \mid P \in \mathcal{P}\}.$$

We write $v \nearrow v'$. We also say that v is the result of *contraction* from v'.

We let $\leq$ be the reflexive, transitive closure of $\nearrow$.

Remark 4.8 (an explicit description of expansion) Consider $[f, \omega I]$, where $\omega \in \{A, B\}^*$. We note that $[f, \omega I] = [f\omega, I]$ by the definition of $\mathcal{B}$ (Definition 4.5) and because $\omega \in \mathbb{S}(I, \omega I)$. An arbitrary partition of I into generating domains takes the form

$$\{\tau I \mid \tau \in \mathcal{C}\},$$

where $\mathcal{C} \subseteq \{A, B\}^*$ is a cut set (in the sense of Section 5). It follows directly that an expansion at $[f, \omega I]$ (equivalently, $[f\omega, I]$) involves replacing $[f, \omega I]$ by the members of

$$\{[f\omega C^k \tau, I] \mid \tau \in \mathcal{C}\},$$

for some $k \in \mathbb{Z}$ and some cut set $\mathcal{C}$. (This is by Definition 4.7 and Theorem 4.3.)

The above description is particularly simple when $\mathcal{C}$ is the cut set $\{A, B\}$. It then follows that the expansion replaces $[f, \omega I]$ with the pairs

$$[f\omega C^k A, I] \quad \text{and} \quad [f\omega C^k B, I],$$

for some $k \in \mathbb{Z}$. Moreover, this is essentially the general case, since any expansion can be realized as a sequence of such expansions.

An expansion at a pair $[f, D]$, where $D = [m, \infty)$, is much more straightforward to describe: such an expansion simply replaces $[f, D]$ with the members of

$$\{[f, P] \mid P \in \mathcal{P}\},$$

where $\mathcal{P}$ is a finite partition of D into domains from $\mathcal{D}^+_{\text{gen}}$. This is because $\mathbb{S}(D, D)$ is trivial.

Proposition 4.9 *Expansion is well-defined and* Γ*-invariant*:

(1) *If* $v=\{[f_1,D_1],\dots,[f_m,D_m]\}$, $\hat{v}=\{[g_1,E_1],\dots,[g_m,E_m]\}$, $v=\hat{v}$, *and* v' *is the result of expansion from* v *at* $[f_i,D_i]$, *then* v' *is also the result of expansion from* $\hat{v}$ *at some* $[g_j,E_j]$.

(2) *If* $v \nearrow v'$ (*where* v *and* v' *are as above*) *and* $\gamma\in\Gamma$, *then* $\gamma\cdot v \nearrow \gamma\cdot v'$.

Proof We prove (1). Assume, without loss of generality, that $[f_k,D_k]=[g_k,E_k]$ for $k=1,\dots,m$. We suppose that v' is the result of expansion from v at $[f_i,D_i]$; thus, there is some $h\in\mathbb{S}(D_i,D_i)$ and a finite partition $\mathcal{P}\subseteq\mathcal{D}^+_{\mathrm{gen}}$ of D_i such that

$$v'=(v-\{[f_i,D_i]\})\cup\{[f_ih,P]\mid P\in\mathcal{P}\}.$$

Choose $j\in\mathbb{S}(D_i,E_i)$ such that $j(\mathcal{P})$ is a finite partition of E_i by members of $\mathcal{D}^+_{\mathrm{gen}}$. (For instance, if $D_i=\omega_iI$ and $E_i=\omega_i'I$, for some $\omega_i,\omega_i'\in\{A,B\}^*$, then we can set $j=\omega_i'\omega_i^{-1}$. The only remaining case is when D_i and E_i are both rays. In this case, there is only one member j of $\mathbb{S}(D_i,E_i)$, and this j satisfies the required property.) Since $[f_i,D_i]=[g_i,E_i]$, there is also some $j_1\in\mathbb{S}(D_i,E_i)$ such that $g_ij_1=f_i$, by Definition 4.5.

We claim that

$$\{[f_ih,P]\mid P\in\mathcal{P}\}=\{[g_ij_1hj^{-1},j(P)]\mid P\in\mathcal{P}\}.$$

Indeed, for each i, $(g_ij_1h^{-1}j^{-1})\circ j=f_ih$, so $[f_ih,P]=[g_ij_1hj^{-1},j(P)]$ by Definition 4.5. This proves (1).

The proof of (2) is straightforward. Indeed,

$$\gamma\cdot v=\{[\gamma f_1,D_1],\dots,[\gamma f_m,D_m]\}$$

and

$$\gamma\cdot v'=(\gamma\cdot v-\{[\gamma f_i,D_i]\})\cup\{[\gamma f_ih,P]\mid P\in\mathcal{P}\},$$

from which it directly follows that $\gamma\cdot v'$ is obtained from $\gamma\cdot v$ via expansion at $[\gamma f_i,D_i]$ (with respect to the same choices of h and $\mathcal{P}$). □

Corollary 4.10 (the partial order on vertices) *The relation* $\le$ *is a partial order on* $\mathcal{V}$. *The group* Γ *acts on* $(\mathcal{V},\le)$ *in an order-preserving fashion.*

Proof This follows directly from Proposition 4.9. □

Definition 4.11 (simplicial realisation; the complexes $\Delta(S_n)$ and $\Delta(S_n')$) Let $\mathcal{P}$ be a partially ordered set. The *simplicial realisation* of P is the simplicial complex whose vertex set is P and whose simplices are finite ascending chains in P.

We let $\Delta(S_n)$ and $\Delta(S_n')$ denote the simplicial realisations of $V(S_n)$ and $V(S_n')$, respectively.

Theorem 4.12 (the directed Γ-set of vertices) *The relation* $\le$ *is a partial order on* $\mathcal{V}$, *and* $\mathcal{V}$ *is a directed set with respect to* $\le$. *The group* Γ *acts on* $(\mathcal{V},\le)$ *in an order-preserving fashion.*

In particular, the simplicial realisations $\Delta(S_i)$ *and* $\Delta(S_i')$ *are contractible* Γ*-complexes.*

Proof It is already clear from Proposition 4.9 and Definition 4.7 that $(\mathcal{V}, \leq)$ is a partially ordered set on which Γ acts in an order-preserving fashion.

Let $S = S_i$ or S'_i, for $i = 2$ or 3. We must show that $(\mathcal{V}, \leq)$ is a directed set. The main step is to show that any vertex

$$v = \{[f_1, D_1], \ldots, [f_m, D_m]\}$$

can be expanded into a vertex of the form

$$\hat{v} = \{[\mathrm{id}_{E_1}, E_1], \ldots, [\mathrm{id}_{E_n}, E_n]\}.$$

Note that each D_i can be partitioned into finitely many elements of $\mathcal{D}^+_{\mathrm{gen}}$ in such a way that the restriction of f_i to each piece acts as a member of S (see Definition 2.7). Thus, we may assume, possibly after expansion, that v already has this property. Consider the pair $[f_1, D_1]$. By Lemma 3.3, there is a finite partition $\mathcal{P} \subseteq \mathcal{D}^+_{\mathrm{gen}}$ of D_1 such that $f_1(P) \in \mathcal{D}^+_{\mathrm{gen}}$, for each $P \in \mathcal{P}$. We note that $f_{1|P} \in S$ by Remark 2.2. It follows that $f_1 \in \mathbb{S}(P, f_1(P))$, for each $P \in \mathcal{P}$, so

$$\{[f_1, P] \mid P \in \mathcal{P}\} = \{[\mathrm{id}, f_1(P)] \mid P \in \mathcal{P}\},$$

by the definition of the equivalence relation on pairs (see Definition 4.5). Note that the act of replacing $[f_1, D_1]$ by the collection $\{[f_i, P] \mid P \in \mathcal{P}\}$ is an expansion at $[f_1, D_1]$. By performing similar expansions at the remaining $[f_i, D_i]$ $(i = 2, \ldots, m)$, we arrive at the required $\hat{v}$.

Now suppose that v_1 and v_2 are any two vertices. By the argument of the previous paragraph, we can find $\hat{v}_1$ and $\hat{v}_2$ such that $v_1 \leq \hat{v}_1$ and $v_2 \leq \hat{v}_2$, and both $\hat{v}_1$ and $\hat{v}_2$ have the general form of the vertex $\hat{v}$; i.e., each pair in $\hat{v}_i$ $(i = 1, 2)$ has the form $[\mathrm{id}, E]$, where id denotes the identity function on E and $E \in \mathcal{D}^+_{\mathrm{gen}}$. Thus, we can identify $\hat{v}_i$ $(i = 1, 2)$ with a partition of the nonnegative real numbers. (Under this identification, $\hat{v}$ would correspond to the partition $\{E_1, \ldots, E_n\}$.)

Finally, we observe that the partitions determined by the $\hat{v}_i$ have a common finite refinement $\mathcal{P}'$ that is also a subset of $\mathcal{D}^+_{\mathrm{gen}}$. Letting $\tilde{v}$ denote the vertex corresponding to $\mathcal{P}'$, we find that $\hat{v}_i \leq \tilde{v}$, for $i = 1, 2$. Thus, $v_1, v_2 \leq \tilde{v}$, from which it follows that $(\mathcal{V}, \leq)$ is a directed set.

The final statement is standard. □

4.2 Vertex stabilisers

In this subsection, we consider the stabiliser Γ_v, where v is a vertex and Γ is one of the groups $V(S_i)$ or $V(S'_i)$ $(i = 2$ or 3). We will largely follow the proof of Proposition 5.3 in [6]. We include the proof for the reader's convenience.

Proposition 4.13 (virtually free abelian vertex stabilisers) *Let*

$$v = \{[f_1, D_1], \ldots, [f_m, D_m]\},$$

where $v \in \Delta(S_n)$ or $v \in \Delta(S'_n)$ $(n = 2$ or 3). Let $\Gamma = V(S_n)$ or $V(S'_n)$ (respectively).

The stabiliser group Γ_v is virtually free abelian of rank at most m.

Proof The elements of the group Γ_v permute the elements of v. That is, for each $\gamma \in \Gamma_v$, there is a permutation $\sigma_\gamma \in S_m$ such that

$$\gamma \cdot [f_j, D_j] = [\gamma \circ f_j, D_j] = [f_{\sigma_\gamma(j)}, D_{\sigma_\gamma(j)}].$$

The assignment $\gamma \mapsto \sigma_\gamma$ is a homomorphism from Γ_v to S_m, the symmetric group on m symbols. The kernel K of the latter homomorphism thus has finite index in Γ_v. Each $\gamma \in K$ fixes the members of v pointwise; i.e., $\gamma \cdot [f_j, D_j] = [\gamma \circ f_j, D_j] = [f_j, D_j]$, for $j = 1, \dots, m$. It follows, from the definition of the equivalence relation, that there are $h_j \in \mathbb{S}(D_j, D_j)$ such that $\gamma_{|f_j(D_j)} = f_j h_j f_j^{-1}$, for $j = 1, \dots, m$. The latter equalities determine an injective homomorphism

$$\Phi : K \to \prod_{j=1}^{m} \mathbb{S}(D_j, D_j)$$

defined by the rule $\gamma \mapsto (h_1, \dots, h_m)$. Since each of the groups $\mathbb{S}(D_j, D_j)$ is either infinite cyclic or trivial by Theorem 4.3, the proposition follows. □

4.3 The directed set constructions for "F" and "T" groups

The "F" and "T" groups act on a subcomplex of the complexes for $\Delta(S_i)$ and $\Delta(S_i')$.

Definition 4.14 Let $\Gamma \in \{F(S_i), F(S_i'), T(S_i)\}$. We consider the smallest subcomplex of $\Delta(S_i)$ (or $\Delta(S_i')$ if $\Gamma = F(S_i')$) that contains the vertices $[\gamma, X]$ ($X = [0,1)$ or $[0,\infty)$, respectively), for all $\gamma \in \Gamma$, and is closed under expansion.

We denote this complex by $\Delta_F(S_i)$, $\Delta_F(S_i')$, or $\Delta_T(S_i)$, respectively.

Proposition 4.15 *The vertices of $\Delta_F(S_i)$, $\Delta_F(S_i')$, and $\Delta_T(S_i)$ form directed sets under expansion.*

In particular, the complexes $\Delta_F(S_i)$, $\Delta_F(S_i')$, and $\Delta_T(S_i)$ are contractible Γ-simplicial complexes, where $\Gamma = F(S_i)$, $F(S_i')$, or $T(S_i)$, respectively.

In all of the above cases, the vertex stabiliser groups are virtually free abelian.

Proof The Γ-equivariance of the complexes in question follows from the fact that the expansion relation is Γ-equivariant. The contractibility of these complexes follows from the fact that the vertex sets are still directed, since the vertex sets in question are closed under expansion.

The proof that the vertex stabiliser groups are free abelian follows the general idea of Proposition 4.13. □

5 An algorithm

In this section, we will describe a simple algorithm. The input is a linear fractional transformation C_n of the interval $[0,1)$. The algorithm attempts to derive a collection of equations like those from Lemma 3.3, which are so essential to our main argument.

We first need to set some conventions. Let $A : [0,1) \to [0,1/2)$ and $B : [0,1) \to [1/2,1)$ be defined as in Definition 2.6. The vertices of a rooted infinite binary tree can be labelled by words in the monoid

$\{A, B\}^*$, as follows: The root is labelled by the empty word. If a given vertex v is labelled by $\omega \in \{A, B\}^*$, then the left and right children of v are labelled by ωA and ωB, respectively. Let us denote the label of v by $L(v)$. We can then assign a half-open interval $I(v)$ to each vertex by the rule

$$I(v) = L(v) \cdot I,$$

where I denotes the interval $[0, 1)$. Note that $I(v_1) \subseteq I(v_2)$ if and only if $L(v_1)$ is a prefix of $L(v_2)$.

By a *cut set* of a rooted infinite binary tree, we mean a set C of vertices such that every embedded geodesic ray issuing from the root passes through exactly one member of C. We may also refer to a set of words in $\{A, B\}^*$ as a cut set if the corresponding set of vertices is a cut set in the above sense.

We define, as in the introduction,

$$C_n(x) = \frac{nx}{(n-1)x + 1},$$

where C_n is defined only on the interval $[0, 1)$. The (hoped-for) output is a collection of matrix identities, of the general form

$$\begin{aligned} C_n\omega_1 &= \omega_1' C_n^{\pm}; \\ C_n\omega_2 &= \omega_2' C_n^{\pm}; \\ &\vdots \\ C_n\omega_k &= \omega_k' C_n^{\pm}, \end{aligned}$$

where $\omega_i, \omega_i' \in \{A, B\}^*$ for $i = 1, \dots, k$, and the sets $\{\omega_1, \dots, \omega_k\}$ and $\{\omega_1', \dots, \omega_k'\}$ are cut sets. (Collections of such equations figured prominently in the proof of Lemma 3.3.) Given the above identities and the corresponding cut sets, we can then define directed sets just as we did in Section 4. The groups that are locally determined by $\{A, B, C_n, a, b, c_n\}^*$ would then act on these directed sets exactly as before.

The algorithm works in the following way. Each vertex of the tree is assigned a type. Initially, this type is "u" for all vertices, indicating a vertex of unknown type. (Actually, the program creates new vertex objects during its run time, although we can ignore this detail for the sake of the current discussion.) Each vertex is also assigned a toggle that is initially set to "0". When the program encounters a vertex v, it performs an action depending on the type of the vertex, which is one of n (for "(internal) node"), l (for "leaf"), or u (for "unknown"), and the value of its toggle, which is either 0 or 1. A value of "0" indicates that the program still needs to do some work at or beneath a given vertex, while a "1" indicates the opposite.

If the vertex is of unknown type ("u"), the program runs the following test:

(1) It first appends C_n to the beginning of the string $L(v)$. This initialises the *matrix string product* of v, which we will here denote by $M(v)$. It is a string over the alphabet $\{A, B, C_n, a, b, c_n\}$.

(2) The program interprets $M(v)$ as a product of matrices and computes the interval $M(v) \cdot I$:

(a) If $M(v) \cdot I \subseteq [0, 1/2)$, then the program appends a to the front of $M(v)$; the result is defined to be the new $M(v)$. The program then returns to step (2).

(b) If $M(v) \cdot I \subseteq [1/2, 1)$, then the program appends b to the front of $M(v)$; the result is defined to be the new $M(v)$. The program then returns to step (2).

(c) If $M(v) \cdot I = I$, then $M(v) \in \mathbb{S}(I, I)$, so $M(v)$ is equivalent (as a linear fractional transformation) to a power of C_n by Theorem 4.3 (or by a variant thereof, if $n \neq 2$ or 3). Let us suppose that $M(v) = C_n^k$. In this case, the program appends c_n^k to the front of $M(v)$ (creating a new $M(v)$). The program now classifies the current vertex as a leaf (changing the unknown "u" designation to "l"). The toggle of the current vertex is also set to "1" (changed from "0").

(d) If $M(v) \cdot I$ satisfies none of the above (i.e., $1/2 \in M(v) \cdot I$, but $M(v) \cdot I \neq I$), then v is reclassified as an (internal) node "n". The toggle stays at 0.

At the end of the above process, the vertex v has been reclassified as an internal node ("n") or a leaf ("l"). In the latter case, the toggle has been set to 1 and a certain matrix string product has been produced. By construction, the (final) matrix string product of a leaf necessarily evaluates to the identity matrix when interpreted as a product of matrices.

If the current vertex v is an internal node (i.e., designated by "n") and its toggle value is 0, then the program determines the toggle value of the left child of v. If this value is 0, it moves to this left child. If the toggle value of the left child is 1, but the toggle value of the right child is 0, then the program moves to the right child. If both children have toggle value 1, then the program flips the toggle of v itself to 1.

If the toggle value of v is 1, then the program moves to the parent of v. If there is no such parent (i.e., v is the root), then the program terminates, and records the matrix string products for each leaf. The latter matrix string products, which take the form

$$c_n^k \omega_1 C_n \omega_2,$$

where $m \in \mathbb{Z}$, $\omega_1 \in \{a, b\}^*$, and $\omega_2 \in \{A, B\}^*$, are readily interpretable as a collection of identities having the desired form, indicated above, if $m = \pm 1$. The leaves determine a (finite) cut set. This completes the description of the algorithm.

We omit the proof of the validity of the algorithm — i.e., the proof that the program finds appropriate cut sets and associated matrix identities, if such things exist.

The author's experience of running the program has led to unexpected results. If $n = 2$, then the program finds a cut set with three elements, and returns the three matrix equations ($C_2 AA = AC_2$; etc.) displayed in the proof of Lemma 3.3. If $n = 3$, then the program finds a cut set with five elements, and the five matrix equations associated to C_3, as described in the proof of Lemma 3.3. If $n = 4$, the program fails to terminate, although it finds many leaves during its run time. The same is true for all values of $n \geq 4$ that the author has tried. (It may be worth noting here that the program computes using only integer values, not floating-point numbers, so round-off errors are apparently not a source for the problems that are encountered here.) It follows from this that an analysis of the groups $V(S_n)$ and $V(S_n')$ for $n \geq 4$ (and, indeed, the corresponding "F" and "T" versions of these groups) lies beyond the techniques described in this paper.

It is also possible to run similar tests for different transformations. One might change not only C_n, but also the transformations A and B. The author has run such tests in a few cases, but with no success to date.

6 The expansion schemes $\mathcal{E}_i$ and $\mathcal{E}'_i$

In this section, we will introduce subdivision trees as a device for diagramming expansions, and describe how subdivision trees represent partitions of [0, 1) into subintervals. Similar trees were considered in [9].

We will then describe expansion schemes $\mathcal{E}_i$ and $\mathcal{E}'_i$, which will eventually be used to simplify the directed set constructions from Section 4. In order to establish the required properties of $\mathcal{E}_i$ and $\mathcal{E}'_i$, we will need to understand when two subdivision trees define the same partition. The latter will be the project of Sections 7 and 8.

6.1 Subdivision trees, equivalence, and elementary equivalence

Definition 6.1 (subdivision trees) Let T be a finite rooted binary tree. The vertices of degree one are *leaves*; all other vertices are *nodes*. The topmost node is the *root*. We say that T is a *subdivision tree* if each node is labelled by an integer.

We let T_ℓ and T_r denote the left and right branches of the subdivision tree T.

Remark 6.2 (the subdivision represented by a subdivision tree) Each leaf in a subdivision tree is labelled by a word in the alphabet $\{A, B, C, c\}$. The labelling is obtained as follows. Trace the (unique) path p from the root to a given leaf ℓ. Suppose that

$$v_1, e_1, v_2, e_2, \dots, e_k, v_{k+1}$$

is a complete list of the vertices and edges encountered along the path p, written in the order that they are encountered. Thus, in particular, v_1 is the root of the tree and $v_{k+1} = \ell$. Let $n_1, n_2, \dots, n_k$ be the integers labelling the nodes $v_1, \dots, v_k$; for $i = 1, \dots, k$, let X_i be A if e_i points downward and to the left, and let X_i be B if e_i points downward and to the right. The labelling of the leaf ℓ is then

$$C^{n_1} X_1 C^{n_2} X_2 \dots C^{n_k} X_k.$$

For instance, the leaves of the left tree in Figure 1 are labelled by the words $ACAcA$, $ACAcB$, ACB, and B. The leaves of the right tree are labelled by C^2ACA, C^2ACBC^3A, C^2ACBC^3B, C^2BcA, and C^2BcB. This labelling scheme works the same way, no matter whether we are letting C represent C_2 or C_3.

We obtain a partition of [0, 1) by applying these words to the interval [0, 1). Thus, $ACAcA$ determines the interval $ACAcA \cdot [0, 1)$, and so forth.

The partition of [0, 1) determined by the trivial subdivision tree is {[0, 1)}.

Remark 6.3 (subdivision trees over M_2 and M_3) For $i = 2, 3$, we let $M_i = \{A, B, C_i, c_i\}^*$; i.e., M_i is the monoid consisting of positive (possibly empty) words in the alphabet $\{A, B, C_i, c_i\}$. For $i = 2, 3$, we

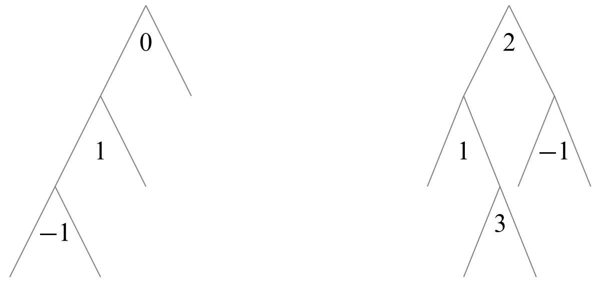

Figure 1: Subdivision trees.

have the proper inclusions $M_i \subsetneq S_i$. In Section 7, we will obtain finite complete presentations of M_2 and M_3, which will aid in analysing subdivision trees.

A subdivision tree T represents one of two subdivisions of the interval [0, 1), depending upon whether the "C" is interpreted as C_2 or C_3. In most contexts, it should be clear which is intended, but, in cases of possible ambiguity, we may refer to T as *a subdivision tree over* M_2 or *over* M_3, as the case may be.

Definition 6.4 (equivalent subdivision trees; the functions n and N) Two subdivision trees T_1 and T_2 (both over either M_2 or M_3) are *equivalent* if they represent the same collection of intervals. We write $T_1 \approx T_2$.

If T is a nontrivial subdivision tree, then $n(T)$ denotes the label of the root. Let

$$N(T) = \{n(T') \mid T' \approx T\}.$$

The set $N(T)$ is empty if T is the trivial subdivision tree.

Lemma 6.5 (finiteness of $N(T)$) *If T is a subdivision tree, then $N(T)$ is a finite set.*

Proof We prove the lemma in the case of M_2, the argument for the case of M_3 being similar.

Note that, if $T' \approx T$ and $n(T') = k$, then the collection of intervals $\mathcal{C}$ determined by T refines $\{[0, 2^k/(2^k+1)), [2^k/(2^k+1), 1)\}$. Thus, if $N(T)$ were infinite, $\mathcal{C}$ would refine an infinite partition of [0, 1), which would force $\mathcal{C}$ to be infinite. This is impossible, since T has only finitely many leaves. □

Lemma 6.6 *Let T and T' be nontrivial subdivision trees (both over M_2 or M_3), and assume that $n(T) = n(T')$. Then $T \approx T'$ if and only if $T_\ell \approx T'_\ell$ and $T_r \approx T'_r$.*

Proof We prove the lemma in the case that T and T' are subdivision trees over M_2; the case of M_3 differs in only minor ways. Assume that T and T' are subdivision trees, and that $n(T) = n(T') = t$.

Assume that $T \approx T'$. Let $\ell_1, \dots, \ell_k$ label the leaves of T_ℓ and let $\ell_{k+1}, \dots, \ell_m$ label the leaves of T_r. (Here, and throughout the proof, the labels of the leaves are read from left to right, so ℓ_1 is the label of the leftmost leaf of T_ℓ, etc.) Let $\ell'_1, \dots, \ell'_{\hat{k}}$ label the leaves of T'_ℓ and let $\ell'_{\hat{k}+1}, \dots, \ell'_m$ label the leaves of T'_r. It follows that

$$C^t A\ell_1, \dots, C^t A\ell_k, C^t B\ell_{k+1}, \dots, C^t B\ell_m$$

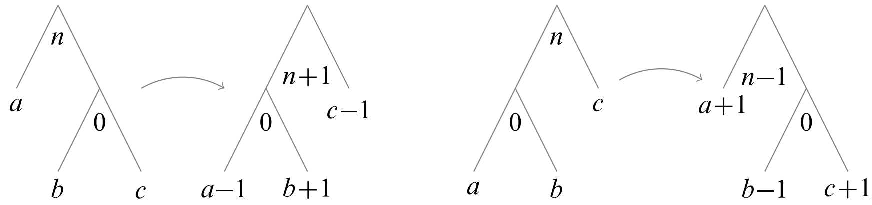

Figure 2: The relations that define elementary equivalence between subdivision trees over M_2.

label the leaves of T and

$$C^t A\ell'_1, \dots, C^t A\ell'_{\hat{k}}, C^t B\ell_{\hat{k}+1}, \dots, C^t B\ell_m$$

label the leaves of T'. It follows that the above labels pairwise determine equal intervals, in the given order: $C^t A\ell_1 \cdot [0,1) = C^t A\ell'_1 \cdot [0,1)$, etc. Since ℓ_k and $\ell'_{\hat{k}}$ label rightmost leaves (of the trees T_ℓ and T'_ℓ, respectively), $C^t A\ell_k \cdot [0,1)$ and $C^t A\ell'_{\hat{k}} \cdot [0,1)$ have the same supremum, namely $2^t/(2^t+1)$ (since $\ell_k \cdot [0,1)$ and $\ell'_{\hat{k}} \cdot [0,1)$ have the supremum 1). It follows directly that $C^t A\ell_k$ and $C^t A\ell'_{\hat{k}}$ determine the same interval; thus, $k = \hat{k}$.

It follows easily that ℓ_j and ℓ'_j determine the same interval, for $j = 1, \dots, k$ (simply cancel $C^t A$ in the relevant products). Thus, $T_\ell \approx T'_\ell$. By similar reasoning, $T_r \approx T'_r$.

Conversely, assuming that $T_\ell \approx T'_\ell$ and $T_r \approx T'_r$, we easily conclude that $T \approx T'$. □

Proposition 6.7 (equality of leaves) *Let $\omega, \omega' \in \{A, B, C, c\}^*$, where $C = C_2$ or C_3. The intervals $\omega \cdot [0,1)$ and $\omega' \cdot [0,1)$ are equal if and only if $\omega = \omega' C^k$, for some $k \in \mathbb{Z}$.*

Proof If $\omega = \omega' C^k$, then

$$\omega \cdot [0,1) = \omega' C^k \cdot [0,1) = \omega' \cdot [0,1),$$

where the final equality follows from the fact that $C \cdot [0,1) = [0,1)$.

Conversely, suppose $\omega I = \omega' I$. It follows that $(\omega')^{-1}\omega I = I$, so $(\omega')^{-1}\omega \in \mathbb{S}(I, I)$, so $(\omega')^{-1}\omega = C^k$, for some $k \in \mathbb{Z}$, by Theorem 4.3. It follows that $\omega = \omega' C^k$. □

Definition 6.8 (elementary equivalence) The two transformations in Figure 2 define *elementary equivalence* between subdivision trees over M_2.

To apply one of the transformations from Figure 2 to a subdivision tree T over M_2 is to replace a subtree of the form on the left with a subtree of the form on the right. Here the labels a, b, c represent the integer labels of the nodes of T that are attached at the leaves labelled by a, b, c (respectively). An application of the given transformation changes the integer labels of these nodes, as indicated on the right-hand tree. If one of the integers a, b, c labels a leaf in T, then that integer is ignored (since leaves of subdivision trees are never labelled by integers). We also say that two subdivision trees T_1 and T_2 are elementary equivalent over M_2 if one can be transformed into the other by a sequence of such transformations.

Elementary equivalence over M_3 is defined by the tree pairs in Figure 3.

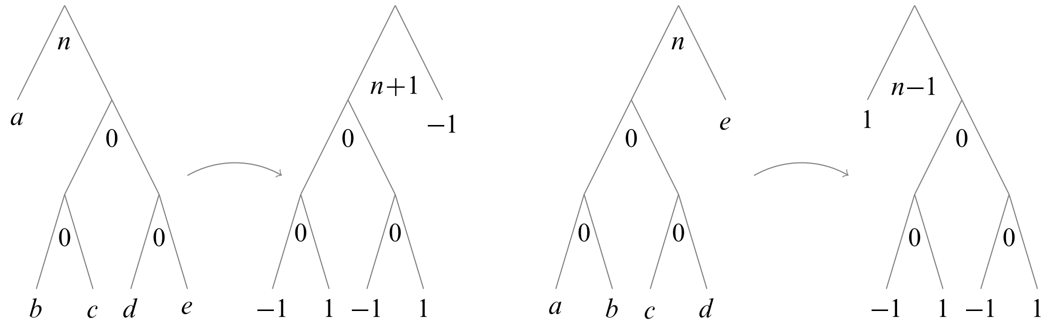

Figure 3: The relations that define elementary equivalence between subdivision trees over M_3.

In Figure 3, the labels on the leaves of the right-hand trees have been abbreviated to avoid creating an over-crowded figure. The leftmost "−1" on the second tree from the left represents "$a-1$", and so on.

Remark 6.9 The transformations in Figure 2 are inverses of each other; similarly for Figure 3.

Example 6.10 Figure 4 depicts an elementary equivalence between two subdivision trees over M_2. The right-hand tree is the result of applying the second relation to the left-hand tree at the node labelled by "4".

One easily checks that the two trees are indeed equivalent.

Lemma 6.11 *If two subdivision trees T_1, T_2 (over M_2 or M_3) are elementary equivalent, then they are equivalent.*

Proof The proof that the left-hand transformation in Figure 2 preserves equivalence relies on the system of equalities

$$C^n A C^a = C^{n+1} A A C^{a-1};$$
$$C^n B A C^b = C^{n+1} B A C^{b+1};$$
$$C^n B B C^c = C^{n+1} B C^{c-1},$$

all of which are easily verified, and from the interpretation of subdivision trees (Remark 6.2). The other three verifications follow similarly. □

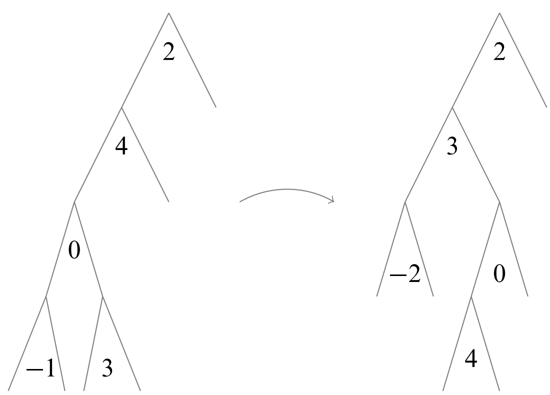

Figure 4: Elementary equivalence between two subdivision trees over M_2.

6.2 The correspondence between subdivision trees and expansions

Theorem 6.12 (subdivision trees and expansions) *We let C denote either C_2 or C_3. Let $v \in \mathcal{V}$ be the result of a sequence of expansions from $\{[\mathrm{id}_I, I]\}$; i.e., $\{[\mathrm{id}, I]\} \le v$. It follows that there is some subdivision tree T such that the set $\mathcal{L}$ of labels on the leaves satisfies*

$$v = \{[\omega, I] \mid \omega \in \mathcal{L}\}.$$

Conversely, any subdivision tree T determines a vertex v by the above equality, and $\{[\mathrm{id}_I, I]\} \le v$ for this v.

If $\{[\mathrm{id}_I, I]\} \le v, v'$ and T and T' are the subdivision trees corresponding to v and v', then $v = v'$ if and only if T and T' are equivalent.

Proof The correspondence between subdivision trees T and vertices v satisfying $\{[\mathrm{id}_I, I]\} \le v$ is straightforward, in view of the discussion in Remark 4.8.

We will now show that $v = v'$ if and only if T is equivalent to T'. Assume first that T and T' are equivalent. Thus,

$$v = \{[\omega_1, I], \dots, [\omega_n, I]\} \quad \text{and} \quad v' = \{[\omega_1', I], \dots, [\omega_n', I]\},$$

where the ω_i are the labels of the leaves of T (listed from left to right) and, similarly, ω_i' are the labels of the leaves of T' (also listed from left to right). Since T and T' are equivalent, we have

$$\omega_i = \omega_i' C^{k_i},$$

for $i = 1, \dots, n$ and for some $k_i \in \mathbb{Z}$, by Definition 6.4 and Proposition 6.7. It follows directly that, for all i, $[\omega_i, I] = [\omega_i', I]$, by Definition 4.5, letting $h = C^{k_i}$ (since $C^{k_i} \in \mathbb{S}(I, I)$). Thus, $v = v'$.

If we carry over the notation from above, the converse essentially follows from the fact that the equality $[\omega_i, I] = [\omega_i', I]$ implies the equality $\omega_i = \omega_i' C^{k_i}$ (for appropriate k_i); this is a direct consequence of Definition 4.5 and the description of $\mathbb{S}(I, I)$ from Theorem 4.3. □

6.3 A discussion of expansion schemes; the expansion schemes $\mathcal{E}_i$ and $\mathcal{E}_i'$

The directed set construction from Section 4, and the generalisations considered in [6], lead to complexes that are often too difficult to analyse when (for instance) attempting to establish finiteness properties for the acting group. One device for simplifying the complexes appeared in [6] under the name of "expansion schemes". An expansion scheme $\mathcal{E}$ assigns to each pair $[f, D] \in \mathcal{B}$ a collection of expansions. This assignment determines a simplicial complex $\Delta^{\mathcal{E}}$ in which the simplices are chains

$$v_1 < v_2 < \cdots < v_n$$

such that the vertices v_i ($i = 2, \dots, n$) are all the result of expansions from v_1 that are allowed by $\mathcal{E}$. Thus, for instance, the trivial expansion scheme, which allows no expansions, results in a discrete set of vertices. At the opposite extreme, an expansion scheme may impose no restraint at all, resulting in the original directed set construction. Since the topology of $\Delta^{\mathcal{E}}$ depends significantly on the choice of $\mathcal{E}$, it would be useful to have a criterion that recognises when the complex $\Delta^{\mathcal{E}}$ is n-connected. The idea of an "n-connected expansion scheme" offers such a criterion. The necessary definitions follow.

We will begin with a general discussion of expansion schemes; the definitions of $\mathcal{E}_i$ and $\mathcal{E}'_i$ are in Example 6.22. The definition of "pseudovertex" is from Section 4 of [6], while the other definitions and theorems in this subsection are from Section 6 of [6].

Definition 6.13 (pseudovertices) Let $v = \{[f_1, D_1], \dots, [f_m, D_m]\} \subseteq \mathcal{B}$. We say that v is a *pseudovertex* if the sets $f_i(D_i)$ $(i = 1 \dots, m)$ are pairwise disjoint.

Remark 6.14 (the partial order on pseudovertices; the action of $\widehat{S}$ on pseudovertices) The pseudovertices are partially ordered by expansion, which can be defined exactly as it was for vertices (Definition 4.7). The pseudovertices do not form a directed set, since the *support* of a given pseudovertex

$$(f_1(D_1) \cup \dots \cup f_m(D_m))$$

is invariant under expansion. The proof of Theorem 4.12 still shows that any two pseudovertices with the same support have an upper bound. Thus, the simplicial realisation of the set of all pseudovertices is a disjoint union of contractible sets.

There is a (partial) action of $\widehat{S}$ (Definition 2.7) on $\mathcal{B}$, defined by

$$\hat{s} \cdot [f, D] = [\hat{s} f, D].$$

This action is defined for suitable $[f, D]$ and $\hat{s}$; i.e., for all pairs $[f, D]$ and $\hat{s} \in \widehat{S}$ such that $f(D)$ is a subset of the domain of $\hat{s}$.

Definition 6.15 ($\mathcal{E}$-expansion; expansion scheme) Let $\mathcal{PV}$ denote the collection of all pseudovertices. Assume that $\mathcal{E} : \mathcal{B} \to 2^{\mathcal{PV}}$ satisfies (1)–(3), for each $[f, D] \in \mathcal{B}$ (we let b, rather than $[f, D]$, denote a typical member of $\mathcal{B}$ in order to simplify notation):

(1) Each $w \in \mathcal{E}(b)$ is the result of a sequence of expansions from $\{b\}$; i.e., for each $w \in \mathcal{E}([f, D])$, we have $\{[f, D]\} \le w$.

(2) $\{b\} \in \mathcal{E}(b)$.

(3) **$\widehat{S}$-invariance** For each $\hat{s} \in \widehat{S}$, and each $b \in \mathcal{B}$ for which $\hat{s} \cdot b$ is defined, $\hat{s} \cdot \mathcal{E}(b) = \mathcal{E}(\hat{s} \cdot b)$.

Let $v \in \mathcal{PV}$; we write $v = \{b_1, \dots, b_m\}$, where $b_1, \dots, b_m \in \mathcal{B}$. We say that v' is a result of *$\mathcal{E}$-expansion from v* if there are $v'_i \in \mathcal{E}(b_i)$, for $i = 1, \dots, m$, such that

$$v' = \bigcup_{i=1}^{m} v'_i.$$

We say that $\mathcal{E}$ is an *expansion scheme* if

(4) for every $[f, D] \in \mathcal{B}$ and every $w_1, w_2 \in \mathcal{E}([f, D])$ such that $w_1 \le w_2$, w_2 is the result of $\mathcal{E}$-expansion from w_1.

Definition 6.16 (the complex $\Delta^{\mathcal{E}}$) Let $\mathcal{E}$ be an expansion scheme. We let $\Delta^{\mathcal{E}}$ be the subcomplex of the directed set construction made up of *$\mathcal{E}$-simplices*; i.e., simplices

$$v_1 < v_2 < \cdots < v_m$$

such that the vertices v_j $(j \in \{2, \dots, m\})$ are obtained from v_1 by $\mathcal{E}$-expansion.

Definition 6.17 (interval subcomplexes; relative ascending links) Let v' and v'' be pseudovertices such that $v' \leq v''$ (in the sense of the expansion partial order; see Remark 6.14 and Definition 4.7). We let $\Delta^{\mathcal{E}}_{[v',v'']}$ denote the set of all $\mathcal{E}$-simplices

$$v_1 < \cdots < v_m$$

such that $v' \leq v_1 < v_m \leq v''$. This is the *interval subcomplex* determined by v', v'', and $\mathcal{E}$.

The *ascending link of v' relative to v''* is the link of v' in the complex $\Delta^{\mathcal{E}}_{[v',v'']}$.

Definition 6.18 (n-connected expansion schemes) Let $\mathcal{E}$ be an expansion scheme. We say that $\mathcal{E}$ is *n-connected* if, for each $b \in \mathcal{B}$ and each pseudovertex v such that $\{b\} < v$, the ascending link of $\{b\}$ relative to v is $(n-1)$-connected.

Theorem 6.19 (n-connectedness of $\Delta^{\mathcal{E}}$) *If $\mathcal{E}$ is an n-connected expansion scheme, then the complex $\Delta^{\mathcal{E}}$ is n-connected.*

Remark 6.20 (n-connectedness of complexes determined by pseudovertices) Theorem 6.19's conclusion carries over to complexes determined by pseudovertices in a component-by-component fashion; i.e., each connected component is n-connected.

Example 6.21 (the case of Thompson's group V) We consider a basic example of an expansion scheme. Let

$$\mathcal{C} = \prod_{n=1}^{\infty} \{0, 1\}$$

denote the usual binary Cantor set. The elements of $\mathcal{C}$ are infinite binary strings. We let $\mathcal{C}^{\text{fin}}$ denote the set of all finite binary strings. For each $\omega \in \mathcal{C}^{\text{fin}}$, we let D_ω denote the set of all infinite binary strings that begin with the prefix ω. For $\omega_1, \omega_2 \in \mathcal{C}^{\text{fin}}$, the transformation $\sigma_{\omega_1,\omega_2} : D_{\omega_1} \to D_{\omega_2}$ removes the prefix ω_1 from the input and adds the prefix ω_2 in its place. We let

$$S_V = \{\sigma_{\omega_1,\omega_2} \mid \omega_1, \omega_2 \in \mathcal{C}^{\text{fin}}\} \cup \{0\},$$

where 0 represents the empty function. The set S_V is an inverse monoid under composition. The associated set of domains $\mathcal{D}^+_{S_V}$ consists of all of the sets D_ω, where $\omega \in \mathcal{C}^{\text{fin}}$.

The set of all bijections $\gamma : \mathcal{C} \to \mathcal{C}$ that are locally determined by S_V make up a group, which we denote by V. This is Thompson's well-known group V, as described in [5]. We define an S_V structure as follows. For each pair $(D_{\omega_1}, D_{\omega_2})$, we define

$$\mathbb{S}(D_{\omega_1}, D_{\omega_2}) = \{\sigma_{\omega_1,\omega_2}\}.$$

The verification that this assignment does, indeed, define an S_V-structure is routine. (For the sake of this discussion, we can use the properties from Proposition 4.4 as the definition of S-structure. The reader is referred to [6] for a more complete definition. We note, however, that the longer definition from the latter source is designed to address numerous complications that do not arise in the case of V.)

We now define an expansion scheme $\mathcal{E}$. For each $[f, D_\omega]$, let

$$\mathcal{E}([f, D_\omega]) = \{\{[f, D_\omega]\}, \{[f, D_{\omega 0}], [f, D_{\omega 1}]\}\}.$$

Thus, the set $\mathcal{E}([f, D_\omega])$ consists of two pseudovertices: the base pseudovertex $\{[f, D_\omega]\}$, and the pseudovertex obtained by performing the simplest possible expansion at $[f, D_\omega]$, namely the expansion that subdivides D_ω into left and right halves ($D_{\omega 0}$ and $D_{\omega 1}$, respectively). It is straightforward to check that the assignment $\mathcal{E}$ satisfies the conditions of Definition 6.15.

A simplex in $\Delta^{\mathcal{E}}$ is a chain

$$v_1 < v_2 < \cdots < v_m,$$

where $v_1 = \{[f_1, D_{\omega_1}], \dots, [f_n, D_{\omega_n}]\}$, and each vertex v_j $(2 \le j \le m)$ can be obtained from v_1 by, for a given $i \in \{1, \dots, n\}$, either replacing $[f_i, D_{\omega_i}]$ with its left and right halves (in the sense described above), or leaving $[f_i, D_{\omega_i}]$ unchanged.

It is also straightforward to check that the expansion scheme $\mathcal{E}$ is n-connected, for all n. Indeed, let $b \in \mathcal{B}$ and let $\{b\} < v$. There is a unique $\mathcal{E}$-expansion from b, and thus the 1-simplex connecting $\{b\}$ to $\{b_\ell, b_r\}$ is the star of $\{b\}$ in $\Delta^{\mathcal{E}}_{[\{b\}, v]}$. The ascending link of $\{b\}$ relative to v is therefore always a point. It follows from Theorem 6.19 that $\Delta^{\mathcal{E}}$ is contractible.

Example 6.22 We now return to the main examples of this paper. Let $S = S_2$, S_3, S_2', or S_3', and let the S-structure $\mathbb{S}$ be defined as in Definition 4.1. We will define expansion schemes $\mathcal{E}_i$ and $\mathcal{E}_i'$, for $i = 2, 3$. In order to do so, we must first introduce some useful notation.

For each $k \in \mathbb{Z}$, we let u_k denote the vertex that corresponds to the subdivision tree consisting of a single caret in which the root is numbered k. (The correspondence in question is that of Theorem 6.12.) Thus,

$$u_k = \{[C^k A, I], [C^k B, I]\},$$

where $C = C_2$ or C_3, depending on the semigroup S in question. If $S = S_2$ or S_2', we let $u_{k-\frac{1}{2}}$ be the vertex corresponding to the subdivision tree consisting of two carets: a top caret (with root labelled k), and a second caret, attached to the left child of the root, labelled 0. Thus,

$$u_{k-\frac{1}{2}} = \{[C^k AA, I], [C^k AB, I], [C^k B, I]\}.$$

If $S = S_3$ or S_3', then $u_{k-\frac{1}{2}}$ is the vertex represented by the subdivision tree consisting of a top caret (with root labelled k), and a complete depth-two binary tree, attached at the left child of the root, in which each node is labelled 0. Thus,

$$u_{k-\frac{1}{2}} = \{[C^k AAA, I], [C^k AAB, I], [C^k ABA, I], [C^k ABB, I], [C^k B, I]\}.$$

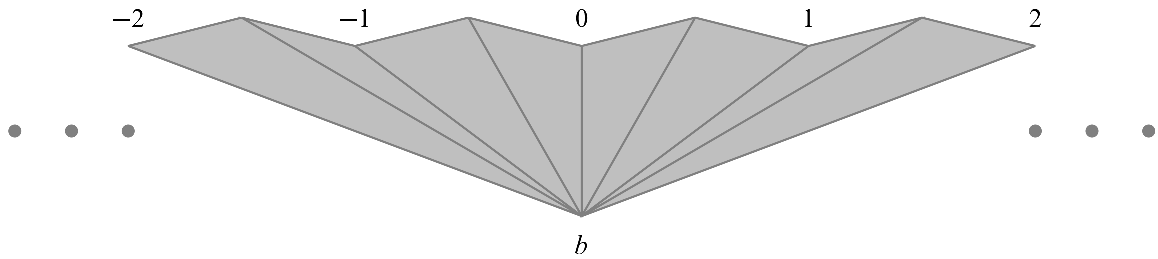

Figure 5: Above we have depicted the simplicial complex $\mathcal{E}(b)$ associated to $b = [\mathrm{id}_I, I]$ by the expansion schemes $\mathcal{E}_i$ and $\mathcal{E}_i'$. An integer k refers to the vertex u_k.

With the above conventions, we can set

$$\mathcal{E}_i([\mathrm{id}_I, I]) = \mathcal{E}_i'([\mathrm{id}_I, I]) = \{\{[\mathrm{id}_I, I]\}\} \cup \{u_{k/2} \mid k \in \mathbb{Z}\},$$

for $i = 2$ or 3. By extending $\widehat{S}$-equivariantly, we arrive at a definition of $\mathcal{E}_i(b) = \mathcal{E}_i'(b)$, for any $b = [f, I]$ and for $i = 2$ or 3, where I is contained in the domain of $f \in \widehat{S}$:

$$\mathcal{E}_i([f, I]) = \mathcal{E}_i'([f, I]) = \{\{[f, I]\}\} \cup \{f \cdot u_{k/2} \mid k \in \mathbb{Z}\}.$$

The well-definedness of this assignment is easy to check.

It is straightforward to check that $u_k \le u_{k-\frac{1}{2}}$ and $u_{k-1} \le u_{k-\frac{1}{2}}$, for each integer k. (The first inequality is clear; the second inequality follows directly after applying an elementary equivalence.) Moreover, no two of the vertices u_{k_1} and u_{k_2} are comparable and no two of the vertices $u_{k_1-\frac{1}{2}}$ and $u_{k_2-\frac{1}{2}}$ are comparable (if $k_1 \ne k_2$). It follows that the simplicial realisations of $\mathcal{E}_i(b)$ and $\mathcal{E}_i'(b)$ take the form indicated in Figure 5.

We recall that $[f, \omega I] = [f\omega, I]$ when $\omega \in \{A, B\}^*$ (see the beginning of Remark 4.8). It follows that the description of $\mathcal{E}_i$ is complete for $i = 2$ and 3.

We next define

$$\mathcal{E}_i'([f, [m, \infty)) = \{\{[f, [m, \infty)]\}, \{[f, [m, m+1)], [f, [m+1, \infty)]\}\}.$$

This completes the definition of $\mathcal{E}_i'$, for $i = 2$ and 3.

Observe that, if $\{[f, [m, \infty)]\} < v$, then

$$v = \{[f, D] \mid D \in \mathcal{P}\},$$

where $\mathcal{P} \subseteq \mathcal{D}^+_{\mathrm{gen}}$ is a finite partition of $[m, \infty)$ into generating domains (see Remark 4.8). The partition $\mathcal{P}$ is necessarily a proper refinement of $\{[m, m+1), [m+1, \infty)\}$ if $\{[f, [m, \infty)]\} < v$. It follows that the ascending link of $\{[f, [m, \infty)]\}$ relative to such v is always a point, and thus contractible. Thus, the expansion scheme $\mathcal{E}_i'$ is n-connected for a given n if and only if $\mathcal{E}_i$ is. We may therefore concentrate on $\mathcal{E}_i$ in what follows.

7 Finite complete presentations of semigroups

In order to understand equivalence between subdivision trees, we will need a full analysis of the monoids M_2 and M_3, which, by definition, are generated by the linear fractional transformations that we have denoted by A, B, C_i, and c_i (for $i = 2, 3$).

In this section, in contrast to our usual practice, the letters A, B, and C will be used as formal symbols. We will define abstract monoid presentations $\mathcal{P}_i$ and $\widehat{\mathcal{P}}_i$ ($i = 2, 3$), with the ultimate goal of proving that the abstract monoid $M(\mathcal{P}_i)$ defined by $\mathcal{P}_i$ is isomorphic to M_i. (The monoids $M(\widehat{\mathcal{P}}_i)$ represent a necessary intermediate device.)

The arguments in this section parallel those from Section 5 of [9].

7.1 Monoid presentations and string-rewriting systems

Definition 7.1 (monoid presentations) Let Σ be a set. The *free monoid on* Σ, denoted by Σ^*, is the set of all positive (possibly empty) words in Σ, with the operation of concatenation. The empty word is denoted by 1. We write $\omega_1 \equiv \omega_2$ if $\omega_1, \omega_2 \in \Sigma^*$ are identical as words.

Let $\mathcal{R}$ be a set of ordered pairs $(r_1, r_2) \in \Sigma^* \times \Sigma^*$. We view such a pair as an equality between words in Σ^*, writing $r_1 = r_2$ if either (r_1, r_2) or (r_2, r_1) is in $\mathcal{R}$. The pair $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ is called a *monoid presentation*; the set $\mathcal{R}$ is the set of *relations*. These relations determine an equivalence relation on Σ^* in the following way. If $\omega_1, \omega_2 \in \Sigma^*$, then we write $\omega_1 \approx \omega_2$ if $\omega_1 \equiv \alpha r_1 \beta$ and $\omega_2 \equiv \alpha r_2 \beta$ for some words $\alpha, \beta \in \Sigma^*$, and $(r_1, r_2) \in \mathcal{R}$. The symmetric, transitive closure of $\approx$, denoted by $=$, is an equivalence relation on Σ^*. We sometimes denote the equivalence class of a word ω by $[\omega]$.

The concatenation operation on Σ^* determines a well-defined associative operation on the set of equivalence classes $\Sigma^*/{=}$. We let $M(\mathcal{P})$ denote the set of these equivalence classes, with the operation induced by concatenation. The set $M(\mathcal{P})$ is a monoid with respect to this operation, called the *monoid determined by* $\mathcal{P}$.

Definition 7.2 (rewrite systems; string-rewriting systems) A *rewrite system* is a directed graph Γ. We allow loops and multiple edges. If v_1 and v_2 are vertices of Γ, we write $v_1 \to v_2$ if there is a directed edge issuing from v_1 and terminating at v_2. We write $v_1 \stackrel{\cdot}{\longrightarrow} v_2$ if there is a directed edge path from v_1 to v_2. Equivalently, $\stackrel{\cdot}{\longrightarrow}$ is the transitive closure of $\to$.

Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a monoid presentation. We define a rewrite system $\Gamma(\mathcal{P})$ as follows. The vertex set of $\Gamma(\mathcal{P})$ is Σ^*. There is a directed edge from ω_1 to ω_2 if $\omega_1 \equiv \alpha r_1 \beta$ and $\omega_2 \equiv \alpha r_2 \beta$, where $\alpha, \beta \in \Sigma^*$ and $(r_1, r_2) \in \mathcal{R}$. The directed graph $\Gamma(\mathcal{P})$ is called the *string-rewriting system* associated to the monoid presentation $\mathcal{P}$.

Remark 7.3 Let $\stackrel{\cdot}{\longleftrightarrow}$ denote the symmetric, transitive closure of $\to$. Thus, $\stackrel{\cdot}{\longleftrightarrow}$ is an equivalence relation on the vertices of $\Gamma(\mathcal{P})$. The above definitions easily show that the relation $\stackrel{\cdot}{\longleftrightarrow}$ coincides with $=$ on Σ^*. In other words, equivalence classes of words in Σ^* modulo $=$ are in one-to-one correspondence with (undirected) path components of $\Gamma(\mathcal{P})$.

$C_2AA \to AC_2$	$C_2B \to BBC_2$	$C_2AB \to BAc_2$
$c_2A \to AAc_2$	$c_2BB \to Bc_2$	$c_2BA \to ABC_2$
$C_2c_2 \to 1$	$c_2C_2 \to 1$	
$C_2aa \to aC_2$	$C_2b \to bbC_2$	$C_2ab \to bac_2$
$c_2a \to aac_2$	$c_2bb \to bc_2$	$bA \to 0$
$aB \to 0$	$bB \to 1$	$aA \to 1$
$0X \to 0$	$X0 \to 0$	

Table 1: The relations of $\widehat{\mathcal{R}}_2$. The relations in the top box are $\mathcal{R}_2$. The "X" stands for any of the generators.

In view of this close identification between the monoid $M(\mathcal{P})$ and the string-rewriting system $\Gamma(\mathcal{P})$, it causes no harm to write $r_1 \to r_2$ for a relation $(r_1, r_2) \in \mathcal{R}$.

Definition 7.4 (terminating; confluent; locally confluent; reduced) A rewrite system Γ is *terminating* if every sequence of vertices $v_1 \to v_2 \to v_3 \to \cdots$ is finite. We say Γ is *confluent* if whenever $v_1 \dot{\longrightarrow} w_1$ and $v_1 \dot{\longrightarrow} w_2$, there is some v_2 such that $w_1 \dot{\longrightarrow} v_2$ and $w_2 \dot{\longrightarrow} v_2$. We say Γ is *locally confluent* if whenever $v_1 \to w_1$ and $v_1 \to w_2$, there is some v_2 such that $w_1 \dot{\longrightarrow} v_2$ and $w_2 \dot{\longrightarrow} v_2$.

A vertex v of Γ is called *reduced* if there is no directed edge issuing from v.

A rewrite system is *complete* if it is terminating and confluent. We say that a monoid presentation $\mathcal{P}$ is complete if the associated string-rewriting system $\Gamma(\mathcal{P})$ is complete.

Theorem 7.5 [12] *If the rewrite system Γ is terminating and locally confluent, then Γ is confluent.* □

Corollary 7.6 (unique reduced forms) *If the rewrite system Γ is terminating and locally confluent, then each connected component of Γ contains a unique reduced vertex.*

In particular, if $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ is a complete monoid presentation, then any connected component of $\Gamma(\mathcal{P})$ contains a unique reduced word, and any word $\omega \in \Sigma^*$ is equivalent to a unique reduced word modulo $=$. □

7.2 Basic definitions of the rewrite systems

Definition 7.7 We define monoid presentations, $\mathcal{P}_2$, $\widehat{\mathcal{P}}_2$, $\mathcal{P}_3$, and $\widehat{\mathcal{P}}_3$, as follows.

(1) $\mathcal{P}_2 = \langle A, B, C_2, c_2 \mid \mathcal{R}_2 \rangle$, where $\mathcal{R}_2$ consists of the relations appearing in the top box of Table 1.

(2) $\widehat{\mathcal{P}}_2 = \langle A, B, C_2, a, b, c_2, 0 \mid \widehat{\mathcal{R}}_2 \rangle$, where $\widehat{\mathcal{R}}_2$ consists of the relations that appear in Table 1.

(3) $\mathcal{P}_3 = \langle A, B, C_3, c_3 \mid \mathcal{R}_3 \rangle$, where $\mathcal{R}_3$ consists of the relations appearing in the top box of Table 2.

(4) $\widehat{\mathcal{P}}_3 = \langle A, B, C_3, a, b, c_3, 0 \mid \widehat{\mathcal{R}}_3 \rangle$, where $\widehat{\mathcal{R}}_3$ consists of the relations that appear in Table 2.

We note that the occurrences of "X" in the tables represent arbitrary generators.

Remark 7.8 (the modified rewrite systems $\Gamma^*(\widehat{\mathcal{P}}_i)$) The string-rewriting systems $\Gamma(\widehat{\mathcal{P}}_i)$ $(i = 2, 3)$ both fail to be locally confluent. For instance, the directed edges $c_2bbB \to bc_2B$ and $c_2bbB \to c_2b$ show that $\Gamma(\widehat{\mathcal{P}}_2)$ is not locally confluent, since the words bc_2B and c_2b are both reduced. The directed edges $c_3bbbB \to c_3bb$ and $c_3bbbB \to bc_3B$ similarly show that $\Gamma(\widehat{\mathcal{P}}_3)$ is not locally confluent.

$C_3AAA \to AC_3$	$C_3AAB \to BAAc_3$	$C_3ABA \to BABC_3$
$C_3ABB \to BBAc_3$	$C_3B \to BBBC_3$	$c_3A \to AAAc_3$
$c_3BAA \to AABC_3$	$c_3BAB \to ABAc_3$	$c_3BBA \to ABBC_3$
$c_3BBB \to Bc_3$	$C_3c_3 \to 1$	$c_3C_3 \to 1$
$c_3a \to aaac_3$	$C_3aab \to baac_3$	$c_3bab \to abac_3$
$C_3abb \to bbac_3$	$c_3bbb \to bc_3$	$C_3aaa \to aC_3$
$c_3baa \to aabC_3$	$C_3aba \to babC_3$	$c_3bba \to abbC_3$
$C_3b \to bbbC_3$	$aB \to 0$	$bA \to 0$
$0X \to 0$	$X0 \to 0$	$bB \to 1$
$aA \to 1$		

Table 2: The relations of $\widehat{\mathcal{R}}_3$. The relations in the top box are $\mathcal{R}_3$.

In order to apply Theorem 7.5 (and thus establish the uniqueness of reduced forms via Corollary 7.6), we will create modified rewrite systems $\Gamma^*(\widehat{\mathcal{P}}_i)$ as follows. The modified rewrite systems have the same vertex sets as the original string-rewriting systems (i.e., vertices are words in the alphabets specified in Definition 7.7). If a word ω contains no occurrences of the subwords 0, aB, bA, bB,aA, c_2C_2, c_3C_3, C_2c_2, or C_3c_3, then the directed edges leading from ω are unchanged. If, however, one or more of the above occurs in ω as a subword, then we repeatedly apply the eight rewriting rules (from either Table 1 or 2, as appropriate) that have "1" or "0" on the right side, until the resulting word, denoted by $R(\omega)$, contains no occurrences of the above subwords, or is 0. The sole directed edge leading from ω then connects to $R(\omega)$.

It is straightforward to check that $R(\omega)$ is indeed uniquely defined; one considers the string-rewriting system that uses only the eight relations specified above. The latter is easily seen to be terminating and locally confluent, so applications of Theorem 7.5 and Corollary 7.6 establish uniqueness.

Proposition 7.9 *The rewrite systems $\Gamma(\mathcal{P}_2)$, $\Gamma(\mathcal{P}_3)$, $\Gamma^*(\widehat{\mathcal{P}}_2)$, and $\Gamma^*(\widehat{\mathcal{P}}_3)$ are locally confluent and terminating. In particular, each word ω has a unique reduced form $r(\omega)$.*

Proof We first consider local confluence in the case of the rewrite systems $\Gamma^*(\widehat{\mathcal{P}}_i)$. Indeed, there is nothing to prove here if only a single directed edge leads away from ω. It therefore suffices to consider only words ω containing no occurrences of 0, aB, bA, bB,aA, c_2C_2, c_3C_3, C_2c_2, or C_3c_3 as subwords. However, in this case local confluence is essentially trivial, since there can be no overlaps between the left sides of the relations that appear as subwords of ω. Thus, assuming that $\ell' \to r'$ and $\ell'' \to r''$ are rewriting rules that are applicable to ω, we can factor ω in the form $\omega_1\ell'\omega_2\ell''\omega_3$ (without loss of generality), where any one of the ω_i could be trivial. The result of applying the first reduction, $\omega_1r'\omega_2\ell''\omega_3$, and the result of applying the second reduction, $\omega_1\ell'\omega_2r''\omega_3$, then both flow to $\omega_1r'\omega_2r''\omega_3$, proving local confluence.

In the case of the string-rewriting systems $\Gamma(\mathcal{P}_i)$, the only overlaps between applications of rewrite rules $\ell' \to r'$ and $\ell'' \to r''$ occur when ℓ' or ℓ'' (or both) take the form cC or Cc. In all of these cases, checking local confluence is straightforward.

The "terminating" condition follows from the fact that every relation either "moves" C closer to the end of the word (possibly changing the occurrence of C to c in the process), or shortens the word. □

Remark 7.10 In the case of the rewrite systems $\Gamma^*(\widehat{\mathcal{P}}_i)$, we can have two (or more) distinct reduced words that are equivalent modulo the monoid presentation $\widehat{\mathcal{P}}_i$. Indeed, bc_2B and c_2b are two such words.

There are, however, no such pairs of words in the case of the string-rewriting systems $\Gamma(\mathcal{P}_i)$, by Remark 7.3.

Definition 7.11 (the monoids M_2 and M_3) Let $T_A : [0,1) \to [0,1/2)$ be the transformation defined by the rule

$$T_A(x) = \frac{x}{x+1}.$$

Thus, T_A is exactly the transformation denoted by A in Section 2. We similarly define T_B, T_{C_2}, and T_{C_3} as B, C_2, and C_3 were defined in Section 2.

For $i = 2, 3$, we let M_i be the monoid generated by the transformations T_A, T_B, T_{C_i}; i.e., the collection of functions generated by these transformations under the operation of concatenation.

For $i = 2, 3$, we let $\widehat{M}_i$ be the inverse monoid generated by the transformations T_A, T_B, T_{C_i}; i.e., the collection of functions generated by these transformations and their inverses, under the operation of composition.

Definition 7.12 (the maps π_2, $\hat{\pi}_2$, π_3, $\hat{\pi}_3$, and π) For each $X \in \{A, B, C_2, C_3\}$, we set $\pi(X) = T_X$. We extend this map to the lower-case letters a, b, c_2, c_3 by sending each to the relevant inverses; i.e., $\pi(a) = T_A^{-1}$, $\pi(b) = T_B^{-1}$, etc. We define $\pi(0)$ to be the empty function (with empty domain and codomain).

For $i = 2, 3$, we define monoid homomorphisms $\pi_i : M(\mathcal{P}_i) \to M_i$ by letting π_i agree with π on the relevant generating sets. We similarly define monoid homomorphisms $\hat{\pi}_i : M(\widehat{\mathcal{P}}_i) \to \widehat{M}_i$ (for $i = 2, 3$) by letting $\hat{\pi}_i$ agree with π on the relevant generating sets.

The homomorphisms $\hat{\pi}_i$ restrict to π_i, for $i = 2, 3$ (and, indeed, $M(\mathcal{P}_i)$ is a submonoid of $M(\widehat{\mathcal{P}}_i)$, as the latter remark implies).

Remark 7.13 The proof that π_i and $\hat{\pi}_i$ $(i = 2, 3)$ are monoid homomorphisms depends on showing that the defining relations of $M(\mathcal{P}_i)$ and $M(\widehat{\mathcal{P}}_i)$ are satisfied by their images in M_i and $\widehat{M}_i$. This verification is routine, and is left to the reader.

It is clear that the maps π_i and $\hat{\pi}_i$ are surjective.

Remark 7.14 Note that, although $T_B^{-1} T_B = 1$ (where "1" here denotes the identity function on [0, 1]), $T_B T_B^{-1} = \mathrm{id}_{[1/2,1)} \neq 1$. Similarly, $T_A T_A^{-1} = \mathrm{id}_{[0,1/2)}$.

Remark 7.15 The rewrite rules $1X \to X$ and $X1 \to X$ are implicit in the definitions of $M(\mathcal{P}_i)$ and $M(\widehat{\mathcal{P}}_i)$. It is technically unnecessary to include them, since "1" is simply notation for the empty string.

7.3 The "no potential cancellations" condition

Throughout this subsection, we will write "C" in place of C_2 or C_3, and similarly write "c" in place of c_2 or c_3.

Definition 7.16 (C-tracks) A subword ω' of $\omega \in \{A, B, C, a, b, c, 0\}^*$ is called a *C-track* if

(1) ω' contains at most one occurrence of C or c (not both);

(2) any occurrence of C or c is at the beginning of the word ω';

(3) ω' is a maximal subword with respect to properties (1) and (2).

Remark 7.17 Any word $\omega \in \{A, B, C, a, b, c, 0\}^*$ has a unique decomposition

$$\omega \equiv \omega_1 \dots \omega_n$$

as a product of C-tracks. For instance, the decomposition of the word

$$CCabCABC$$

into C-tracks is $\omega_1\omega_2\omega_3\omega_4$, where

$$\omega_1 = C, \quad \omega_2 = Cab, \quad \omega_3 = CAB, \quad \omega_4 = C.$$

Definition 7.18 ([9, Definition 5.7], advancing an occurrence of C or c) To *advance* an occurrence of C (or c) is to apply one of the relations from Definition 7.7, other than those of the form $Cc \to 1$, $cC \to 1$, $0X \to 0$, and $X0 \to 0$, to a subword containing that occurrence of C or c.

Definition 7.19 ([9, Definition 5.8], no potential cancellations) Assume that $\omega \in \{A, B, C, c\}^*$. Let

$$\omega \equiv \omega_1 \dots \omega_n$$

be the unique decomposition into C-tracks. We say that ω has *no potential cancellations* if the words

$$r(\omega_1)\omega_2 \dots \omega_n, \quad \omega_1 r(\omega_2) \dots \omega_n, \quad \dots, \quad \omega_1\omega_2 \dots r(\omega_n)$$

contain no occurrences of cC or Cc as subwords. Here $r(\omega)$ (for $\omega \in \{A, B, C, c\}^*$) denotes the reduced form of ω relative to $\Gamma(\mathcal{P}_i)$, for $i = 2$ or 3 (see Definition 7.7 and Proposition 7.9).

Remark 7.20 If ω has no potential cancellations, then ω contains no occurrences of cC or Cc as subwords. This follows directly from the observation that c and C are their own reduced forms; i.e., $r(c) = c$ and $r(C) = C$.

Proposition 7.21 [9, Lemma 5.9] *If $\omega \in \{A, B, C, c\}^*$ has no potential cancellations and ω' is the result of advancing a c or C exactly once, then ω' has no potential cancellations.*

Proof Let $\omega \equiv \omega_1\omega_2 \dots \omega_n$, where the right side of the equation is the unique decomposition of ω into C-tracks. Suppose that ω' is the result of advancing an occurrence of C (or c) exactly once; suppose that

$\ell \to r$	$x \equiv C$ or c	$x \equiv CA$	$x \equiv cB$
$CAA \to AC$	$CAC \dashrightarrow CAC$	$CAAC \dashrightarrow ACC$	$cBAC \dashrightarrow ABCC$
$CB \to BBC$	$CBBC \dashrightarrow B^4C^2$	$CAB^2C \dashrightarrow BAcBC$	$cB^3C \dashrightarrow cB^3C$
$CAB \to BAc$	$CBAc \dashrightarrow BBCCAc$	$CABAc \dashrightarrow BAAAcc$	$cBBAc \dashrightarrow BAAcc$
$cA \to AAc$	$cAAc \dashrightarrow A^4cc$	$CA^3c \dashrightarrow ACAc$	$cBAAc \dashrightarrow ABCAc$
$cBB \to Bc$	$cBc \dashrightarrow cBc$	$CABc \dashrightarrow BAcc$	$cBBc \dashrightarrow Bcc$
$cBA \to ABC$	$cABC \dashrightarrow AAcBC$	$CAABC \dashrightarrow AB^2C^2$	$cBABC \dashrightarrow AB^3C^2$

Table 3: The proof of Proposition 7.21 in the case of $M(\mathcal{P}_2)$.

the advanced occurrence of C or c appears in ω_i, and let $\ell \to r$ be the relation that advances this C or c. Thus $\omega_i \equiv \ell\beta$ for some word β. Let

$$\omega' \equiv \omega'_1\omega'_2 \dots \omega'_n$$

be the unique decomposition of ω' into C-tracks. It follows directly that $\omega'_{i-1}\omega'_i \equiv \omega_{i-1}r\beta$, while $\omega'_j \equiv \omega_j$ if $j \in \{1, \dots, n\} - \{i-1, i\}$. Note that the subword $\omega_{i-1}r$ consists of the C-track of the $(i-1)$-st occurrence of a C (or c) in ω', followed by a C (or c); note also that the only chance of an occurrence of Cc or cC in the words $\omega'_1 \dots r(\omega'_{j-1})\omega'_j \dots \omega'_n$ might occur when $j = i$.

To prove the proposition, it therefore suffices to prove that, during the reduction of the subword $\omega_{i-1}r$, no occurrence of Cc or cC can arise. Note that $\omega_{i-1}r$ begins and ends with occurrences of C (or c), while all intermediate letters are A or B. Thus, after reducing ω_{i-1}, it suffices to show that an occurrence of Cc or cC cannot arise in (further) reducing $r(\omega_{i-1})r$. Finally, we note that the reduced word $r(\omega_{i-1})$ ends in a reduced word x that begins with a C or c. There are only finitely many possibilities for x: indeed, $x \in \{c, C, CA, cB\}$ in the case of $M(\mathcal{P}_2)$, while $x \in \{c, C, CA, CAA, CAB, cB, cBA, cBB\}$ in the case of $M(\mathcal{P}_3)$. Furthermore, in either case, one of the cases $x \equiv C$ or $x \equiv c$ can be ruled out, since $x\ell$ contains no occurrence of Cc or cC by hypothesis.

Thus, in summary, it suffices to show that, for each rewriting rule $\ell \to r$, no occurrence of cC or Cc can appear when reducing the word xr, where x runs over the above possibilities and further satisfies the condition that $x\ell$ itself contains neither cC nor Cc.

The relevant calculations are summarised in Tables 3 and 4. □

Definition 7.22 (negative-to-positive words) Let ω be a word in the alphabet $\{A, B, C, a, b, c, 0, 1\}$. We say that ω is *negative-to-positive* if all occurrences of a and b (if any) occur before any occurrence of either A or B.

Remark 7.23 If $\omega \neq 1, 0$ is a negative-to-positive word containing no occurrences of bB or aA or 1, then each C-track in ω is a word in either $\{a, b, C, c\}$ or $\{A, B, C, c\}$. We call a C-track in the former alphabet *negative*, while a C-track in the latter alphabet is *positive*.

A C-track consisting only of the single symbol C or c can be freely considered negative or positive.

$\ell \to r$	$x \in \{c, C, CA, CAA\}$	$x \in \{CAB, cB, cBA, cBB\}$
$CA^3 \to AC$	$CAC \stackrel{\cdot}{\longrightarrow} CAC$ $CAAC \stackrel{\cdot}{\longrightarrow} CAAC$ $CAAAC \stackrel{\cdot}{\longrightarrow} ACC$	$CABAC \stackrel{\cdot}{\longrightarrow} BABCC$ $cBAC \stackrel{\cdot}{\longrightarrow} cBAC$ $cBAAC \stackrel{\cdot}{\longrightarrow} AABCC$ $cBBAC \stackrel{\cdot}{\longrightarrow} ABBCC$
$CABA \to BABC$	$CBABC \stackrel{\cdot}{\longrightarrow} B^3CABC$ $CABABC \stackrel{\cdot}{\longrightarrow} BAB^4C^2$ $CAABABC \stackrel{\cdot}{\longrightarrow} BA^5cBC$	$CABBABC \stackrel{\cdot}{\longrightarrow} B^2A^4cBC$ $cBBABC \stackrel{\cdot}{\longrightarrow} AB^5C^2$ $cBABABC \stackrel{\cdot}{\longrightarrow} ABA^4cBC$ $cB^3ABC \stackrel{\cdot}{\longrightarrow} BA^3cBC$
$CAAB \to BAAc$	$CBAAc \stackrel{\cdot}{\longrightarrow} B^3CAAc$ $CABAAc \stackrel{\cdot}{\longrightarrow} BABCAc$ $CAABAAc \stackrel{\cdot}{\longrightarrow} BA^8cc$	$CABBAAc \stackrel{\cdot}{\longrightarrow} A^4B^2CAc$ $cBBAAc \stackrel{\cdot}{\longrightarrow} ABBCAc$ $cBABAAc \stackrel{\cdot}{\longrightarrow} ABA^7cc$ $cB^3AAc \stackrel{\cdot}{\longrightarrow} BA^6cc$
$CABB \to BBAc$	$CBBAc \stackrel{\cdot}{\longrightarrow} B^6CAc$ $CABBAc \stackrel{\cdot}{\longrightarrow} B^2A^4cc$ $CAABBAc \stackrel{\cdot}{\longrightarrow} BA^2cBAc$	$CABBBAc \stackrel{\cdot}{\longrightarrow} BBAcBAc$ $cBBBAc \stackrel{\cdot}{\longrightarrow} BA^3cc$ $cBABBAc \stackrel{\cdot}{\longrightarrow} ABAcBAc$ $cB^4Ac \stackrel{\cdot}{\longrightarrow} BcBAc$
$CB \to BBBC$	$CBBBC \stackrel{\cdot}{\longrightarrow} B^9C^2$ $CABBBC \stackrel{\cdot}{\longrightarrow} BBAcBC$ $CAABBBC \stackrel{\cdot}{\longrightarrow} BA^2cBBC$	$CABBBBC \stackrel{\cdot}{\longrightarrow} BBAcBBC$ $cBBBBC \stackrel{\cdot}{\longrightarrow} BcBC$ $cBABBBC \stackrel{\cdot}{\longrightarrow} ABAcBBC$ $cB^5C \stackrel{\cdot}{\longrightarrow} BcBBC$
$cA \to AAAc$	$cAAAc \stackrel{\cdot}{\longrightarrow} A^9c^2$ $CA^4c \stackrel{\cdot}{\longrightarrow} ACAc$ $CA^5c \stackrel{\cdot}{\longrightarrow} ACAAc$	$CABAAAc \stackrel{\cdot}{\longrightarrow} BABCAAc$ $cBAAAc \stackrel{\cdot}{\longrightarrow} AABCAc$ $cBAAAAc \stackrel{\cdot}{\longrightarrow} AABCAAc$ $cBBAAAc \stackrel{\cdot}{\longrightarrow} ABBCAAc$
$cBAA \to AABC$	$cAABC \stackrel{\cdot}{\longrightarrow} A^6cBC$ $CAAABC \stackrel{\cdot}{\longrightarrow} AB^3C^2$ $CA^4BC \stackrel{\cdot}{\longrightarrow} ACABC$	$CABAABC \stackrel{\cdot}{\longrightarrow} BABCABC$ $cBAABC \stackrel{\cdot}{\longrightarrow} A^2B^4C^2$ $cBA^3BC \stackrel{\cdot}{\longrightarrow} AABCABC$ $cBBAABC \stackrel{\cdot}{\longrightarrow} ABBCABC$
$cBAB \to ABAc$	$cABAc \stackrel{\cdot}{\longrightarrow} A^3cBAc$ $CAABAc \stackrel{\cdot}{\longrightarrow} BA^5cc$ $CAAABAc \stackrel{\cdot}{\longrightarrow} AB^3CAc$	$CABABAc \stackrel{\cdot}{\longrightarrow} BAB^4CAc$ $cBABAc \stackrel{\cdot}{\longrightarrow} ABA^4cc$ $cBAABAc \stackrel{\cdot}{\longrightarrow} A^2B^4CAc$ $cBBABAc \stackrel{\cdot}{\longrightarrow} AB^5CAc$
$cBBA \to ABBC$	$cABBC \stackrel{\cdot}{\longrightarrow} A^3cBBC$ $CAABBC \stackrel{\cdot}{\longrightarrow} BAAcBC$ $CA^3BBC \stackrel{\cdot}{\longrightarrow} AB^6CC$	$CABABBC \stackrel{\cdot}{\longrightarrow} BAB^7CC$ $cBABBC \stackrel{\cdot}{\longrightarrow} ABAcBC$ $cBAABBC \stackrel{\cdot}{\longrightarrow} AAB^7CC$ $cBBABBC \stackrel{\cdot}{\longrightarrow} AB^8C^2$
$cB^3 \to Bc$	$cBc \stackrel{\cdot}{\longrightarrow} cBc$ $CABc \stackrel{\cdot}{\longrightarrow} CABc$ $CAABc \stackrel{\cdot}{\longrightarrow} BAAcc$	$CABBc \stackrel{\cdot}{\longrightarrow} BBAcc$ $cBBc \stackrel{\cdot}{\longrightarrow} cBBc$ $cBABc \stackrel{\cdot}{\longrightarrow} ABAcc$ $cBBBc \stackrel{\cdot}{\longrightarrow} Bcc$

Table 4: The proof of Proposition 7.21 in the case of $M(\mathcal{P}_3)$.

Definition 7.24 (no potential cancellations in negative-to-positive words) Let $\omega \in \{A, B, C, a, b, c\}^*$. Assume that the reduced form of ω is not 0, and that ω also contains no occurrences of bB or aA.

Let

$$\omega \equiv \omega_1 \dots \omega_n$$

be the unique decomposition into C-tracks. We say that ω has *no potential cancellations* if the words

$$r(\omega_1)\omega_2 \dots \omega_n, \quad \omega_1 r(\omega_2) \dots \omega_n, \quad \dots, \quad \omega_1\omega_2 \dots r(\omega_n)$$

contain no occurrences of cC or Cc. Here $r(\omega)$ denotes the reduced form of ω relative to the rewrite system $\Gamma^*(\widehat{\mathcal{P}}_i)$ (see Remark 7.8 and Proposition 7.9).

Remark 7.25 If $\omega \equiv \omega_1 \dots \omega_n$ is the decomposition of the negative-to-positive word ω into C-tracks, then the words $\omega_1 \dots r(\omega_j) \dots \omega_n$ (for $j = 1, \dots, n$) need not be accessible from ω by a directed edge-path in $\Gamma^*(\widehat{\mathcal{P}}_i)$. This contrasts with the case of words in the generators $\{A, B, C, c\}$; i.e., words consisting only of positive C-tracks.

Consider the word c_2bC_2BB. We have

$$c_2br(C_2BB) \equiv c_2bBBBBC_2,$$

and there is no directed edge-path from c_2bC_2BB to $c_2bBBBBC_2$, since the application of the relation $C_2B \to BBC_2$ must be followed by an application of the cancellation $bB \to 1$:

$$c_2bC_2BB \to c_2bBBC_2B \to c_2BC_2B \to c_2BBBC_2 \to Bc_2BC_2.$$

Proposition 7.26 *Assume that*

(1) *ω is negative-to-positive*;

(2) *ω has no potential cancellations*;

(3) *ω has no subword of the form aA or bB*;

(4) *the reduced form of ω (in the sense of the rewrite system $\Gamma^*(\widehat{\mathcal{P}}_i)$) is not* 0.

Let ω' be the result of advancing a c or C exactly once, and then removing all occurrences of aA or bB, along with all occurrences of "1". The word ω' is also a negative-to-positive word with no potential cancellations.

Proof The proof is like that of Proposition 7.21. We assume that ω' is the result of advancing the i-th occurrence of C or c in ω exactly once. Let $\omega_1\omega_2 \dots \omega_n$ be the C-track decomposition of ω. There are three cases: ω_{i-1} and ω_i are both positive C-tracks, or ω_{i-1} and ω_i are both negative, or ω_{i-1} is negative and ω_i is positive.

We note that the first case (in which both C-tracks are positive) is already handled by the proof of Proposition 7.21. The second case is also handled by the proof of Proposition 7.21. This follows from the observation that each rewriting rule between words in the alphabet $\{a, b, C, c\}$ corresponds to a rewriting rule in between words in the alphabet $\{A, B, C, c\}$. One need only replace a C with c (or the reverse), an A with a b, and a B with an a. Thus, for instance, the rewriting rule $CAA \to AC$ corresponds to $cbb \to bc$. Using this substitution, we can transform Tables 3 and 4 into tables that prove the negative-to-negative case.

It remains to consider the case in which ω_{i-1} is negative and ω_i is positive. Since ω_{i-1} is a negative C-track, we have $\omega_{i-1} \equiv C^{\pm}u$, where u is a nonempty word in the generators a and b. Let $\ell \to r$ be the rewriting rule that advances the occurrence of C or c in the subword ω_i. There are nonempty words r_1, r_2 such that $r \equiv r_1 r_2$, where $r_1 \in \{A, B\}^*$ and $r_2 = c$ or C. After applying the rule $\ell \to r$ to the subword $\omega_{i-1}\ell$, and before any cancellation, we arrive at a word of the form $C^{\pm}ur_1r_2$. Since $\omega \neq 0$, the subword $C^{\pm}ur_1r_2$ cannot contain any occurrence of aB or bA. It follows that one of the rewriting rules $aA \to 1$ or $bB \to 1$ can be applied at least once to ur_1, and, indeed, that such rules can be applied to ur_1 until an entirely positive word (i.e., a word in the generators $\{A, B\}$) or an entirely negative word remains. Note that, in either case, the word ω' described in the proposition is still negative-to-positive. In fact, only the "no potential cancellations" condition remains to be proved.

The proof involves an analysis of various subcases. We first assume that, after cancellation, a negative word remains. In this subcase, $u \equiv \hat{u}r_1^{-1}$, where $\hat{u}$ is a possibly empty word in the generators $\{a, b\}$. It suffices to show that the word $r(C^{\pm}\hat{u})r_2$ contains no occurrence of the subwords cC or Cc. Suppose, for a contradiction, that there is such an occurrence. The occurrence must be at the end of the word $r(C^{\pm}\hat{u})r_2$, from which it follows that $r(C^{\pm}\hat{u}) \equiv \tilde{u}r_2^{-1}$, for some $\tilde{u} \in \{a, b\}^*$. Next, we consider again the subword $C^{\pm}u\ell$ of ω. We have

$$C^{\pm}u\ell \equiv C^{\pm}\hat{u}r_1^{-1}\ell \overset{\cdot}{\longrightarrow} \tilde{u}r_2^{-1}r_1^{-1}\ell.$$

Since, for each rewrite rule $\ell \to r$ in $\mathcal{P}_i$, the rule $r^{-1} \to \ell^{-1}$ is a rewrite rule in $\widehat{\mathcal{P}}_i$ (see Tables 1 and 2), it follows that

$$\tilde{u}r_2^{-1}r_1^{-1}\ell \overset{\cdot}{\longrightarrow} \tilde{u}\ell^{-1}\ell,$$

which implies that $\omega_1 \dots r(\omega_{i-1})\omega_i \dots \omega_n$ contains an occurrence of cC or Cc (which appears in the subword $\ell^{-1}\ell$). This is a contradiction of the "no potential cancellations" hypothesis.

Now we consider the subcases in which a nonempty positive word remains after cancelling within ur_1. In this subcase, we can list all of the possibilities for the word $\omega_{i-1}r \equiv C^{\pm}ur$, which arises from the subword $\omega_{i-1}\ell$ of ω after advancing the initial "c" or "C" in ω_i via the rewriting rule $\ell \to r$. Indeed, in the case of $i = 2$, the only possibilities are

$$\begin{array}{llll} C_2bBBC_2, & c_2bBBC_2, & C_2bBAc_2, & c_2bBAc_2, \\ C_2aAAc_2, & c_2aAAc_2, & C_2aABC_2, & c_2aABC_2. \end{array}$$

Each case is easily handled. For instance, in the first case, C_2bBBC_2 becomes C_2BC_2 after cancelling bB. The word C_2B is the $(i-1)$-st C-track in ω'. Rewriting, we find that $r(C_2B) = BBC_2$, which shows that ω' still has no potential cancellations.

The case of $i = 3$ involves many more cases, but we can make some useful general observations. We consider the possible forms of the word $R(C_3^{\pm}ur)$, which is the subword of ω' that occurs after advancing the leading "C" symbol in ω_i and after performing any reductions of the form $aA \to 1$ or $bB \to 1$. We note that $R(C_3^{\pm}ur)$ necessarily takes the form $C_3^{\pm}\tilde{u}C_3^{\pm}$, where $\tilde{u} \in \{A, B\}^*$ has length either 1 or 2, and the first and last "C" symbol may have the same or opposite exponents. An examination of Table 2 shows that the only way that the initial "C" symbol in $C_3^{\pm}\tilde{u}C_3^{\pm}$ can be advanced to the end of the word (and thus create a potential cancellation) is if $\tilde{u}$ does not contain both of the symbols A and B. We can therefore assume that $\tilde{u} = A$, AA, B, or BB. Next, we note that, if $\tilde{u}$ ends with A, then r necessarily ended with c_3, while if $\tilde{u}$ ends with B, then r necessarily ended with C_3. (This again follows from Table 2.) Finally, we note that, if there is to be cancellation in $C_3^{\pm}\tilde{u}C_3^{\pm}$ as the result of advancing the initial "C" symbol, then the first and last exponents must be opposite; this now follows because advancing a "C" symbol past a word of the form A, AA, B, or BB never changes the exponent. Thus, the only cases left are

$$C_3Ac_3, \quad C_3AAc_3, \quad c_3BC_3, \quad c_3BBC_3.$$

These words are all reduced, completing the proof. □

Corollary 7.27 *If ω satisfies the hypotheses of Proposition 7.26 and $\omega \dashrightarrow \omega'$ in $\Gamma^*(\widehat{\mathcal{P}}_i)$, then ω' also satisfies the hypotheses of Proposition 7.26, after we remove all occurrences of aA and bB.*

Proof This follows by repeatedly applying Proposition 7.26 to ω, and from the fact that any application of a rewriting rule to ω (under the hypotheses on ω) necessarily entails advancing a "C" symbol. □

7.4 Presentations and normal forms for the monoids M_2 and M_3

Proposition 7.28 [9, Lemma 5.10] *Let ω be a word in the generators $\{A, B, C, c\}$. Assume that ω has no potential cancellations.*

There is a word $\tau \in \{A, B\}^*$ such that $r(\omega\tau) \equiv \widehat{\omega}C^{\epsilon}$, where $\widehat{\omega}$ is a word in $\{A, B\}$ and $\epsilon \geq 0$ is the total exponent of C and c in ω.

Proof The proof is by induction on the (combined) exponent ϵ of C and c in ω. We note that, due to the "no potential cancellations" condition, it is not possible to reduce (or, indeed, increase) the exponent ϵ by applying any of the monoid relations.

Our proof will use the fact that, if $\omega \in \{A, B, C, c\}^*$ has no potential cancellations, then any word of the form $\omega\tau$ ($\tau \in \{A, B\}^*$) also has no potential cancellations. This is an easy consequence of Definition 7.19.

We first consider the case of $M(\mathcal{P}_2)$; assume $\epsilon = 1$, the case $\epsilon = 0$ being trivial. We note that $r(\omega)$ ends with one of the strings C, c, CA, or cB (and the only occurrences of C or c occur in these strings). In the case of C, there is nothing to prove. If $r(\omega)$ ends with c, we can let $\tau \equiv BA$ and then reduce the result. If $r(\omega)$ ends with either CA or cB, we can let $\tau \equiv A$ and then reduce the result. This proves the base case.

Now let $\epsilon > 1$. We can express ω as a product $\omega_1\omega_2$, where the total combined exponent of C and c in ω_2 is $\epsilon - 1$, and ω_1 contains a single occurrence of C or c. By induction, we can find $\tau_1 \in \{A, B\}^*$ such that $r(\omega_2\tau_1) \equiv \widehat{\omega}_2C^{\epsilon-1}$, where $\widehat{\omega}_2 \in \{A, B\}^*$. Thus, after reducing the word $\omega_1\omega_2\tau_1$, we obtain a word $\omega' \in \{A, B, C, c\}^*$ that ends with C^ϵ, $CAC^{\epsilon-1}$, or $cBC^{\epsilon-1}$. (Note that the case $cC^{\epsilon-1}$ is ruled out by the "no potential cancellations" hypothesis.) In the first case, we are finished; set $\tau_2 \equiv 1$. If ω' ends with $CAC^{\epsilon-1}$ or $cBC^{\epsilon-1}$, we can set $\tau_2 \equiv A^{2^{\epsilon-1}}$. After reducing the word $\omega'\tau_2$, we have a string of the required form, so the required τ is $\tau_1\tau_2$.

Now we consider $M(\mathcal{P}_3)$. Let $\omega \in \{A, B, C, c\}$ and define ϵ as before. We first consider the case $\epsilon = 1$. The word $r(\omega)$ ends with c, C, CA, CAB, CAA, cB, cBA, or cBB. If $r(\omega)$ ends with c, we can let $\tau \equiv BBA$ and apply the relation $cBBA \to ABBC$. If $r(\omega)$ ends with C, there is nothing to prove ($\tau \equiv 1$). In the remaining cases, we let $\tau \equiv AA, A, A, AA, A$, or A (respectively).

Now suppose $\epsilon > 1$. We can write ω as the product $\omega_1\omega_2$, where the total combined exponent of C and c in ω_2 is $\epsilon - 1$, and ω_1 contains a single occurrence of either C or c. Proceeding as in the case of $M(\mathcal{P}_2)$, we can right multiply by some $\tau_1 \in \{A, B\}^*$ and reduce to arrive at a word ω' that ends with one of the following strings: C^ϵ, $CAC^{\epsilon-1}$, $CABC^{\epsilon-1}$, $CAAC^{\epsilon-1}$, $cBC^{\epsilon-1}$, $cBAC^{\epsilon-1}$, $cBBC^{\epsilon-1}$. In the first case, there is nothing to prove; let $\tau_2 \equiv 1$. In the remaining cases, we multiply by $\tau_2 \equiv A^{2\cdot 3^{\epsilon-1}}$, $A^{3^{\epsilon-1}}$, $A^{3^{\epsilon-1}}$, $A^{2\cdot 3^{\epsilon-1}}$, $A^{3^{\epsilon-1}}$, or $A^{3^{\epsilon-1}}$, respectively. Thus, the required τ is $\tau_1\tau_2$. □

Proposition 7.29 *Let $\omega \in \{A, B, C, a, b, c\}^*$ be a negative-to-positive word with no potential cancellations. Assume that there is no $\tau \in \{A, B\}^*$ such that $r(\omega\tau) \equiv 0$.*

There is some $\tau' \in \{A, B\}^*$ such that $r(\omega\tau') \equiv \widehat{\omega}C^\epsilon$, where $\widehat{\omega} \in \{A, B\}^*$ and ϵ is the total combined exponent of C and c in ω.

Proof We prove this by induction on the sum k of the combined exponents of a and b in ω. The case $k = 0$ is handled by Proposition 7.28. We will use the fact that, if $\omega \in \{A, B, a, b, C, c\}^*$ is a negative-to-positive word with no potential cancellations, then so is the word $\omega\tau$, where τ is any word in $\{A, B\}^*$. This fact is easily verified from Definition 7.24.

Let $\omega \in \{A, B, C, a, b, c\}^*$, let k be defined as above, and suppose that the proposition is known to be true for smaller k. We can write $\omega \equiv \omega_1\omega_2$, where ω_1 involves no occurrences of A or B, and ω_2 involves no occurrences of a or b. We may further assume that ω_1 ends with an occurrence of a or b, since any occurrence of C or c may be subsumed by ω_2.

By Proposition 7.28, we can find a word $\tau_1 \in \{A, B\}^*$ such that $r(\omega_2\tau_1) \equiv \widehat{\omega}C^{\epsilon_2}$, where ϵ_2 is the total exponent sum of C and c in ω_2 and $\widehat{\omega} \in \{A, B\}^*$. If $\widehat{\omega}$ is not the empty word, then it must be that the initial letter of $\widehat{\omega}$ cancels with the terminal letter of ω_1 in $\omega_1\widehat{\omega}C^{\epsilon_2}$. (This is because occurrences of aB and bA cannot arise, by the hypothesis that $r(\omega\tau)$ is never 0.) After performing all cancellations of the form $aA \to 1$ and $bB \to 1$, we can call the inductive hypothesis, to find τ_2 such that $r(\omega_1\widehat{\omega}C^{\epsilon_2}\tau_2) \equiv \widetilde{\omega}C^{\epsilon_1+\epsilon_2}$, where $\widetilde{\omega} \in \{A, B\}^*$ and ϵ_1 is the total exponent of c and C in ω_1. This completes the induction, under the assumption that $\widehat{\omega}$ is not the empty word. (We note that, to apply the inductive hypothesis, we are implicitly calling Corollary 7.27, and using the completeness of the rewrite system $\Gamma^*(\widehat{\mathcal{P}}_i)$.)

If $\widehat{\omega} \equiv 1$, we simply multiply by a suitable word $\tau_{3/2}$: either $A^{2^{\epsilon_2}}$ or $A^{3^{\epsilon_2}}$ (depending on whether we are considering $M(\mathcal{P}_2)$ or $M(\mathcal{P}_3)$). After reducing, we find that $r(\omega_2\tau_1\tau_{3/2}) \equiv \widehat{\omega}'C^{\epsilon_2}$, where $\widehat{\omega}'$ is nonempty. This reduces us to the previous case, completing the induction and the proof. □

Proposition 7.30 (normal forms in $M(\mathcal{P}_i)$) *The reduced words modulo the presentation $\mathcal{P}_2$ take the form*

$$\omega_1\omega_2\omega_3,$$

where $\omega_1 \in \{A, B\}^*$, $\omega_2 \in \{C^n A, c^m B \mid m, n \in \mathbb{N}\}^*$, and $\omega_3 \in \{C, c\}^*$.

The reduced words modulo the presentation $\mathcal{P}_3$ take the form

$$\omega_1\omega_2\omega_3,$$

where $\omega_1 \in \{A, B\}^*$, $\omega_2 \in \{C^{n_1} A, C^{n_2} AA, C^{n_3} AB, c^{n_4} B, c^{n_5} BA, c^{n_6} BB \mid n_i \in \mathbb{N}\}^*$, and $\omega_3 \in \{C, c\}^*$.

Proof It is clear that the words in question are reduced. Thus, the main point is to show that every word in the generators can be reduced to a word of the given type. This is easily done by induction on the length of the word. □

Theorem 7.31 (monoid presentations for M_2 and M_3) *The monoid homomorphisms $\pi_i : M(\mathcal{P}_i) \to M_i$ are isomorphisms, for $i = 2, 3$.*

In particular, $\mathcal{P}_i$ is a presentation for M_i, for $i = 2, 3$.

Proof In view of Remark 7.13, it suffices to show that π_i is injective, for $i = 2, 3$. We suppose, for a contradiction, that π_2 is not injective. Let

$$S' = \{\{\omega_1, \omega_2\} \mid \omega_1 \not\equiv \omega_2;\ \pi_2(\omega_1) = \pi_2(\omega_2);\ \omega_1 \text{ and } \omega_2 \text{ are reduced}\}.$$

We let

$$S'' = \{\{\omega_1, \omega_2\} \mid \{\omega_1, \omega_2\} \in S';\ \omega_1, \omega_2 \text{ begin with different letters}\}.$$

We note that S' is nonempty by hypothesis, and it follows easily that S'' is also nonempty. (It suffices to cancel the maximal common prefix of the words ω_1, ω_2, where $\{\omega_1, \omega_2\} \in S'$.) Next, we note that if $\{\omega_1, \omega_2\} \in S''$, then one of ω_1 or ω_2 begins with C or c, or is trivial. (The case in which ω_1 begins with "A" and ω_2 begins with "B" (or the reverse) can be ruled out, since $\pi_2(\omega_1)$ cannot be equal to $\pi_2(\omega_2)$ under these conditions.) We will assume (without loss of generality) that it is ω_1 that begins with C or c, or is trivial.

Consider $\{\omega_1, \omega_2\} \in S''$ such that the total exponent of C and c in ω_1 and ω_2 is a minimum. We first assume that ω_1 begins with either C or c. (The case in which ω_1 is trivial is easier, and will be handled in the course of the more difficult argument.) It follows that $\omega_1 \equiv \omega_1' C^k$, where $\omega_1' \in \{C^n A, c^m B \mid n, m \in \mathbb{N}\}^*$ and $k \in \mathbb{Z}$, by Proposition 7.30. Indeed, we can assume that $k = 0$, for if $k \neq 0$, then we simply multiply both ω_1 and ω_2 on the right by c^k. The word $\omega_2' :\equiv \omega_2 c^k$ is then necessarily reduced, by the hypothesis that the total exponent of C and c in ω_1 and ω_2 is a minimum in S''. We can then replace the pair $\{\omega_1, \omega_2\}$ by $\{\omega_1', \omega_2'\}$, where the latter is still in S''.

Thus, we can assume that $\omega_1 \in \{C^n A, c^m B \mid n, m \in \mathbb{N}\}^*$, $\{\omega_1, \omega_2\} \in S''$, and the total combined exponent of C and c in ω_1 and ω_2 is a minimum within S''. We claim that the words ω_1^{-1} and ω_2 have no potential cancellations. This is obvious in the case of ω_2, since it is reduced. In ω_1^{-1}, every occurrence of a is followed by c, every occurrence of b is followed by C, and there are no occurrences of

A, B, Cc, or cC. Now note that no occurrences of Cc or cC can occur when reducing subwords of the form $C^{\pm 1}ac$ or $C^{\pm 1}bC$; from this it follows that ω_1^{-1} has no potential cancellations.

Next we claim that $\omega_1^{-1}\omega_2$ has no potential cancellations. (Here the claim is obvious if ω_1 is the trivial word; thus, the argument from this point is the general case.) Indeed, the only C-track of $\omega_1^{-1}\omega_2$ that could cause a problem is the one that begins with the terminal letter (c or C) of ω_1^{-1}. Assume that $\omega_1 \equiv C\widehat{\omega}_1$ (without loss of generality), and suppose that $\omega_1^{-1}\omega_2$ has a potential cancellation. Replacing $\{\omega_1, \omega_2\}$ by $\{\widehat{\omega}_1, r(c\omega_2)\}$, we find (possibly after cancelling common prefixes) that the latter is a pair in S'' of smaller total exponent in C and c. This contradicts the choice of $\{\omega_1, \omega_2\}$, which proves the claim.

Since $\pi_2(\omega_1) = \pi_2(\omega_2)$, we have $\widehat{\pi}_2(\omega_1^{-1}\omega_2) = \widehat{\pi}(1) = \mathrm{id}_{[0,1)}$. In particular, this means that $\omega_1^{-1}\omega_2$ satisfies the hypotheses of Proposition 7.29. We can therefore find a word $\tau \in \{A, B\}^*$ such that $r(\omega_1^{-1}\omega_2\tau) \equiv \widehat{\omega}C^k$, where $k \geq 0$ is the total combined exponent of C and c in the word $\omega_1^{-1}\omega_2$ and $\widehat{\omega} \in \{A, B\}^*$. Thus, we have $\pi_2(\widehat{\omega}C^k) = \pi_2(\tau)$. We now cancel the maximal common prefix of $\widehat{\omega}C^k$ and τ. We continue to denote the resulting strings by $\widehat{\omega}C^k$ and τ, respectively, but now either $\widehat{\omega}$ or τ is trivial.

If $\widehat{\omega}$ is trivial, but τ is not, then $\pi_2(\widehat{\omega}C^k) = C^k$, while $\pi_2(\tau)$ is a transformation whose range is a proper subinterval of $[0, 1)$. This is a contradiction. If τ is trivial, but not $\widehat{\omega}$, we find that $\pi_2(\tau)$ has the image $[0, 1)$, but $\pi_2(\widehat{\omega}C^k)$ does not, which is also a contradiction. Finally, if both $\widehat{\omega}$ and τ are trivial, we find that $\pi_2(C^k) = \mathrm{id}_{[0,1)}$, which is possible only if $k = 0$. The latter implies that ω_1 is the trivial word, $\omega_2 \in \{A, B\}^*$, and ω_2 is not the trivial word. This leads us to conclude that $\pi_2(\omega_2) = \mathrm{id}_{[0,1)}$, which is impossible since the image of $\pi_2(\omega_2)$ is a proper subinterval of $[0, 1)$. Thus, π_2 is injective.

The case of π_3 is similar. □

Remark 7.32 (presentations for certain submonoids of $\mathrm{Isom}(\mathbb{H}^2)$) Let $\tilde{A}$ denote the transformation of the projective line $\mathbb{P}_1$ ($= \partial\mathbb{H}^2$) that agrees with A (as defined in Definition 2.6) on the $[0, 1)$. Similarly define $\widetilde{B}$, $\widetilde{C}_2$, and so forth. The transformations $\tilde{A}$, $\widetilde{B}$, $\widetilde{C}_2$, $\widetilde{C}_3$, and their inverses may equivalently be considered isometries of $\mathbb{H}^2$. Let $\widetilde{S}_i = \{\tilde{A}, \widetilde{B}, \widetilde{C}_i, \tilde{c}_i\}^*$, for $i = 2, 3$. There are obvious homomorphisms $\phi_i : \widetilde{S}_i \to M_i$ for $i = 2, 3$. It is just as clear that these homomorphisms are surjective. If $\phi_i(\alpha) = \phi_i(\beta)$, but $\alpha \neq \beta$, then α and β are transformations of $\partial\mathbb{H}^2$ that agree on $[0, 1)$. This is impossible, however, since any isometry of $\mathbb{H}^2$ is determined by its effect on any three boundary points. Thus, ϕ_i is injective, for $i = 2, 3$.

It follows from all of this that $\widetilde{S}_2$ and $\widetilde{S}_3$ admit the same presentations and normal forms as do M_2 and M_3.

8 An intermediate value theorem for the expansion scheme $\mathcal{E}_i$

In this section, we will argue that $N(T)$ is always a set of consecutive integers if T is a nontrivial subdivision tree. This will be the main ingredient to our proof that the expansion schemes $\mathcal{E}_i$ and $\mathcal{E}'_i$ are n-connected, for all n. The proof of the latter fact will be assembled in Section 9.

The essential idea, that of a "sufficiently expanded subdivision tree", is drawn from [9, Definition 5.5].

8.1 The case of M_2

In this subsection, we will argue that $N(T)$ is always a set of consecutive integers in the case that T is a subdivision tree over M_2. All of the subdivision trees in question will be subdivision trees over M_2; "C" will refer to C_2, and so forth.

We remind the reader that "node" means "interior node" (see Definition 6.1).

Definition 8.1 (sufficiently expanded subdivision trees) A subdivision tree is *sufficiently expanded* if there is a directed arc p in T from the root ϵ to a leaf ℓ such that

(1) each nonroot node on the arc has a nonzero label;

(2) if p passes through a nonroot node v and the label of v is positive, then p also passes through the left child of v;

(3) if p passes through a nonroot node v and the label of v is negative, then p also passes through the right child of v.

A sufficiently expanded subdivision tree is *left sufficiently expanded* (respectively, *right sufficiently expanded*) if some directed arc p as described above passes through the left (respectively, right) child of the root.

Lemma 8.2 *Let T be a subdivision tree.*

(1) *If $n(T)=k$ and $k-1\notin N(T)$, then there is $T'\approx T$ such that $n(T')=k$ and T' is left sufficiently expanded.*

(2) *If $n(T)=k$ and $k+1\notin N(T)$, then there is $T'\approx T$ such that $n(T')=k$ and T' is right sufficiently expanded.*

Proof We prove both parts simultaneously by induction on the number of carets in the subdivision tree T. The induction begins trivially, since a subdivision tree with a single caret is necessarily both left and right sufficiently expanded.

Now we consider a subdivision tree T, and assume that the lemma is true of all subdivision trees containing fewer carets. We will argue for (1); the argument proving (2) is similar. Thus, we let $n(T)=k$ and assume that $k-1\notin N(T)$. We note that $0\notin N(T_\ell)$; otherwise (up to equivalence) T takes the form in Figure 6. This allows us to apply an elementary equivalence from Definition 6.8, resulting in a $T'\approx T$ such that $n(T')=k-1$. However, this implies that $k-1\in N(T)$, a contradiction. Thus, $0\notin N(T_\ell)$, as claimed.

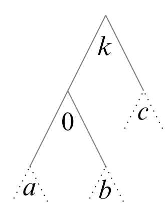

Figure 6: The case in which $0\in N(T_\ell)$.

We let k_1 be either

(i) the smallest positive member of $N(T_\ell)$, or

(ii) the largest negative member of $N(T_\ell)$.

We can assume that $n(T_\ell) = k_1$ (possibly after replacing T_ℓ with an equivalent tree and applying Lemma 6.6). We note that $k_1 - 1 \notin N(T_\ell)$ (in case (i)), or $k_1 + 1 \notin N(T_\ell)$ (in case (ii)); therefore the inductive hypothesis applies, and we conclude that T_ℓ is equivalent to a left sufficiently expanded tree T'_ℓ (in case (i)), or to a right sufficiently expanded tree T'_ℓ (in case (ii)). In either case, we replace T_ℓ by T'_ℓ. We let T' denote the result of replacing T_ℓ by T'_ℓ in the tree T. We note that $n(T') = n(T) = k$ and $T' \approx T$ by Lemma 6.6.

Assume that we are in case (i); case (ii) is similar. Since T'_ℓ is left sufficiently expanded, there is a path p' from the root of T'_ℓ to a leaf of T'_ℓ satisfying the properties in Definition 8.1, such that p' also passes through the left child of the root of T'_ℓ. Let p be the concatenation of e and p', where e is the edge connecting the root of T' to its left child, the root of T'_ℓ. The path p satisfies all of the properties from Definition 8.1 and passes through the left child of the root in T', so T' is left sufficiently expanded. This completes the induction. □

Proposition 8.3 *Let T be a nontrivial subdivision tree.*

(1) *If T is left sufficiently expanded and T' satisfies $n(T') < n(T)$ then $T' \not\approx T$.*

(2) *If T is right sufficiently expanded and T' satisfies $n(T') > n(T)$ then $T' \not\approx T$.*

Proof We first prove (1). Assume that T is left sufficiently expanded and T' is such that $n(T') < n(T)$ and $T' \approx T$. After letting a suitable power of C act at the roots of T and T', we can assume that $n(T) > 0$ and $n(T') = 0$. Since T is left sufficiently expanded, there is a directed arc p from the root of T to a leaf ℓ satisfying the conditions of Definition 8.1; the label of ℓ is a reduced word ω. There is a leaf ℓ' of T' that corresponds to ℓ; let ω' be the label of ℓ'. We have $\omega = \omega' C^k$, for some $k \in \mathbb{Z}$, by Proposition 6.7. After reducing, we find

$$\omega \equiv r(\omega' C^k).$$

However, these words cannot be equal letter-by-letter, since ω necessarily begins with an occurrence of C, but $\omega' C^k$ (and, thus, $r(\omega' C^k)$) begins with either A or B. This is a contradiction to Theorem 7.31.

The proof of (2) is similar. One can reduce to the case in which $n(T) < 0$ and $n(T') = 0$, and then argue that ω begins with a c, while $r(\omega' C^k)$ begins with either A or B. □

Proposition 8.4 *Let T be a nontrivial subdivision tree.*

(1) *If $0 \notin N(T_\ell)$, then $n(T) = \min(N(T))$.*

(2) *If $0 \notin N(T_r)$, then $n(T) = \max(N(T))$.*

Proof We prove (1), the proof of (2) being similar.

We can find a subdivision tree $T'_\ell \approx T_\ell$ such that $n(T'_\ell)$ is either the smallest positive number in $N(T_\ell)$ or the largest negative number in $N(T_\ell)$. In either case, the hypothesis of Lemma 8.2 applies, and we can replace T'_ℓ by T''_ℓ, where T''_ℓ is sufficiently expanded. We can then find a directed arc p'' from the root of T''_ℓ to a leaf ℓ, where p'' satisfies the conditions from Definition 8.1. We can then replace the tree T_ℓ by T''_ℓ within the tree T, to create a new T' such that $n(T') = n(T)$, $T' \approx T$, and T''_ℓ is the left branch of the tree T'. Now let $p' = ep''$, where e is the edge connecting the root of T' to the root of T''_ℓ. The path p' satisfies all of the conditions of Definition 8.1, and shows that T' is left sufficiently expanded.

If $T'' \approx T$ and $n(T'') < n(T)$, then $n(T'') < n(T')$ and $T'' \approx T'$, which contradicts Proposition 8.3(1). It follows that $n(T) = \min(N(T))$. □

Theorem 8.5 (the intermediate value theorem for M_2) *If T is a nontrivial subdivision tree, then*

$$N(T) = [m, M] \cap \mathbb{Z},$$

where $m = \min(N(T))$ and $M = \max(N(T))$.

Proof Suppose that $k \in (m, M) \cap \mathbb{Z}$ but $k \notin N(T)$. Assume further that k is the minimal such integer.

There is a subdivision tree $T' \approx T$ such that $n(T') = k - 1$. It must be that $0 \notin N(T'_r)$ (otherwise, we can apply an elementary equivalence to produce a tree $T'' \approx T'$ such that $n(T'') = k$). Thus, $k - 1 = n(T') = \max(N(T')) = \max(N(T)) = M$, a contradiction. □

8.2 The case of M_3

In this subsection, we will argue that $N(T)$ is always a set of consecutive integers in the case that T is a subdivision tree over M_3. All of the subdivision trees in question will be subdivision trees over M_3; "C" will refer to C_3, and so forth.

Definition 8.6 (blocking trees) Let T be a subdivision tree. We say that T is a *blocking tree* if either

(1) both of the children of the root of T are nodes, and the three vertices (the root and its children) are not all labelled by 0, or

(2) one of these three vertices is a leaf, or T is trivial.

Definition 8.7 (sufficiently expanded in M_3) A subdivision tree T over M_3 is *sufficiently expanded* if there is a directed arc p from the root ϵ to some leaf ℓ such that

(1) each nonroot node on the arc p is the root of a blocking (sub)tree;

(2) if a nonroot node v on the arc p has a positive label, then the arc p passes through the left child of v;

(3) if a nonroot node v on the arc p has a negative label, then the arc p passes through the right child of v;

(4) if the arc p passes through a nonroot node v labelled by "0", then the next node along p (if any) has a nonzero label.

A sufficiently expanded subdivision tree is *left sufficiently expanded* (respectively, *right sufficiently expanded*) if p passes through the left (respectively, the right) child of ϵ.

Lemma 8.8 *Let T be a subdivision tree.*

(1) *If $n(T) = k$ and $k-1 \notin N(T)$, then there is $T' \approx T$ such that $n(T') = k$ and T' is left sufficiently expanded.*

(2) *If $n(T) = k$ and $k+1 \notin N(T)$, then there is $T' \approx T$ such that $n(T') = k$ and T' is right sufficiently expanded.*

Proof The proof resembles that of Lemma 8.2. We argue by induction on the number of carets in the subdivision tree T. If T consists of a single caret, then it is necessarily both left sufficiently expanded and right sufficiently expanded; thus, the base case is satisfied.

Now consider an arbitrary subdivision tree T, and suppose that the lemma has been proved for all subdivision trees having fewer carets. We assume that T satisfies (1); the case of (2) is similar. Since $k-1 \notin N(T)$, the left branch T_ℓ of T is a blocking tree. (Indeed, all trees in the equivalence class of T_ℓ are blocking trees, by Lemma 6.6.) There are two possibilities for T_ℓ: either $0 \in N(T_\ell)$ or $0 \notin N(T_\ell)$. In the latter case, we can proceed essentially as in the proof of Lemma 8.2. We therefore assume that $0 \in N(T_\ell)$; indeed, we can assume that $n(T_\ell) = 0$ without loss of generality. Since all trees in the equivalence class of T_ℓ are blocking trees, it must be that either $0 \notin N(T_{\ell\ell})$ or $0 \notin N(T_{\ell r})$. We assume that $0 \notin N(T_{\ell\ell})$. Thus, by induction, we can replace $T_{\ell\ell}$ by a subdivision tree $T'_{\ell\ell} \approx T_{\ell\ell}$ such that there is a path $\hat{p}$ from the root of $T'_{\ell\ell}$ to a leaf of $T'_{\ell\ell}$ that satisfies the conditions of Definition 8.7. We let T' be the result of replacing $T_{\ell\ell}$ with $T'_{\ell\ell}$ in T. Now, letting p_1 denote the path from the root of T' to the root of $T'_{\ell\ell}$ and $p = p_1\hat{p}$, the path p shows that T' is left sufficiently expanded, completing the induction. □

Proposition 8.9 *Let T be a subdivision tree.*

(1) *If T is left sufficiently expanded and T' satisfies $n(T') < n(T)$ then $T' \not\approx T$.*

(2) *If T is right sufficiently expanded and T' satisfies $n(T') > n(T)$ then $T' \not\approx T$.*

Proof The proof is no different from that of Proposition 8.3; again the crucial observation is that the left sufficiently expanded tree T has a leaf ℓ whose label is a reduced word, and the corresponding leaf ℓ' in T' has a leaf whose label, after reduction, cannot be equivalent to that of ℓ. □

Proposition 8.10 *Let T be a subdivision tree.*

(1) *If, whenever $T' \approx T_\ell$, T' is a blocking tree, then $n(T) = \min(N(T))$.*

(2) *If, whenever $T' \approx T_r$, T' is a blocking tree, then $n(T) = \max(N(T))$.*

Proof We prove (1), the proof of (2) being similar.

The proof of Lemma 8.8 allows us to replace T with an equivalent $\widehat{T}$ such that $\widehat{T}$ is left sufficiently expanded and $n(T) = n(\widehat{T})$.

Let $\widetilde{T} \approx T$. Thus, $\widetilde{T} \approx \widehat{T}$, so $n(\widetilde{T}) \geq n(\widehat{T})$, by Proposition 8.9. Thus, $n(\widetilde{T}) \geq n(T)$, which implies $n(T) = \min(N(T))$. □

Theorem 8.11 (the intermediate value theorem for M_3) *If T is a subdivision tree over M_3, then*

$$N(T) = [m, M] \cap \mathbb{Z}.$$

Proof Suppose that $k \in (m, M) \cap \mathbb{Z}$ but $k \notin N(T)$. Assume further that k is the minimal such integer.

There is a subdivision tree $T' \approx T$ such that $n(T') = k-1$. It must be that all subdivision trees that are equivalent to T'_r are blocking trees (otherwise, we can first replace T'_r by a nonblocking equivalent tree T''_r, and then apply an elementary equivalence to produce a $T'' \approx T'$ such that $n(T'') = k$). Thus,

$$k-1 = n(T') = \max(N(T')) = \max(N(T)) = M,$$

a contradiction. □

9 The proof of the F_∞ property

In this section, we will complete the proof that the expansion schemes $\mathcal{E}_i$ and $\mathcal{E}'_i$ ($i = 2, 3$) are n-connected for all n. This involves assembling a few pieces from Section 8.

We will also complete the proofs that the groups $F(S_i)$, $F(S'_i)$, $T(S_i)$, $V(S_i)$, and $V(S'_i)$ have type F_∞, for $i = 2, 3$. These proofs are almost entirely like the ones from [6].

Recall that the approach in this paper departed from that of [6] in using a proper subset $\mathcal{D}^+_{\mathrm{gen}}$ of the domains $\mathcal{D}^+$ as the foundation for the original directed set construction. This makes little difference in the final arguments, but rather than simply referring the reader to [6] (which runs to over sixty pages), we will sketch the necessary changes when it seems appropriate to do so.

9.1 Brown's finiteness criterion

Here we briefly recall Brown's finiteness criterion for the reader's convenience.

Theorem 9.1 ([3] Brown's finiteness criterion) *Let X be a CW-complex. Let G be a group acting on X. If*

(1) *X is $(n-1)$-connected,*

(2) *G acts cellularly on X, and*

(3) *there is a filtration $X_1 \subseteq X_2 \subseteq \cdots \subseteq X_k \subseteq \cdots \subseteq X$ such that*

(a) *$X = \bigcup_{k=1}^{\infty} X_k$,*

(b) *G leaves each $X_k^{(n)}$ invariant and acts cocompactly on each $X_k^{(n)}$,*

(c) *each p-cell stabiliser has type F_{n-p}, and*

(d) *for sufficiently large k, X_k is $(n-1)$-connected,*

then G is of type F_n. □

9.2 Contractibility of the complexes $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}'_i}$

In this subsection, we will prove that the complexes $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}'_i}$ are contractible, for $i = 2, 3$. This completes a line of argument that was begun at the end of Section 6, and extended through Sections 7 and 8.

Recall that the directed set constructions of the classifying spaces for the groups $F(S)$, $T(S)$, $V(S)$ differed in details (see Section 4). We will use the same notation, $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}'_i}$, to denote the subcomplexes determined by $\mathcal{E}$- (or $\mathcal{E}'$-) expansions in all cases, trusting that the precise meaning will always be clear from the context.

Theorem 9.2 (contractibility of the complexes $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}'_i}$) *The complexes $\Delta^{\mathcal{E}_i}$ are contractible, for each of the groups $F(S)$, $T(S)$, and $V(S)$ ($S \in \{S_2, S_3\}$).*

The complexes $\Delta^{\mathcal{E}'_i}$ are contractible, for each of the groups $F(S)$ and $V(S)$ ($S \in \{S'_2, S'_3\}$).

Proof By Theorem 6.19, it suffices to show that the expansion schemes $\mathcal{E}_i$ and $\mathcal{E}'_i$ ($i = 2, 3$) are n-connected for all n. By the discussion at the end of Example 6.22, it suffices to consider the expansion scheme $\mathcal{E}$.

In the groups $F(S_i)$, $T(S_i)$, $V(S_i)$ ($i = 2, 3$), there is just one domain type, namely $[I]$. By equivariance of $\mathcal{E}_i$, it suffices to show that, whenever $\{[\mathrm{id}_I, I]\} < v$, the ascending link of $\{[\mathrm{id}_I, I]\}$ relative to v is contractible.

Since $\{[\mathrm{id}_I, I]\} < v$, v can be represented by a subdivision tree T (by Theorem 6.12). By Theorem 8.5 or 8.11, $N(T) = [m, M] \cap \mathbb{Z}$, for some integers m and M. Thus, $v \geq u_k$ for an integer k if and only if $k \in [m, M]$, where u_k is as defined in Example 6.22).

Now we must determine precisely when $u_{k-1/2} \leq v$. Note first that, if $k - 1/2 \notin [m, M] \subseteq \mathbb{R}$, then $u_{k-1/2} \not\leq v$, since, if it were, we would conclude that $u_{k-1} \leq v$ and $u_k \leq v$ (since $u_{k-1}, u_k \leq u_{k-1/2}$ in the expansion partial order). This contradicts our hypothesis, since at least one of $k-1$ and k is not in $[m, M]$. Now assume that $k - 1/2 \in [m, M]$. It follows from this that $k-1, k \in N(T)$, since $u_{k-1} \leq u_{k-1/2}$ and $u_k \leq u_{k-1/2}$ in the expansion partial order. It follows that $m < k \leq M$ (in the linear order on $\mathbb{R}$). If we are in the case $S = S_2$ or S'_2, then Proposition 8.4 and the inequality $u_k \leq v$ show that $0 \in N(T_\ell)$ (since $k - 1 \in N(T)$. Thus, there is some T', $T' \approx T$, such that the root of T' is labelled by k and the left child of the root is labelled by 0. Since T' represents the vertex v, it follows that $u_{k-1/2} \leq v$. If we are in the case $S = S_3$ or S'_3, then Proposition 8.10 and the inequality $u_k \leq v$ show that there is some T', $T' \approx T$, such that the root of T' is labelled by k and the left branch of T' is not a blocking tree. It now follows directly that $u_{k-1/2} \leq v$ in this case, as well.

Thus, the ascending link of $\{[\mathrm{id}_I, I]\}$ relative to v corresponds exactly to the portion of the cellulated line ℓ between the vertices u_m and u_M, where ℓ is as depicted in Figure 5. The ascending link in question is therefore contractible. □

Remark 9.3 We review some of the relevant ideas from [6].

The basic approach to proving n-connectedness is laid out in Lemma 2.6 from [6]. Let $\widehat{\Delta}$ be a simplicial complex whose vertices are a directed set, which we denote by X. Let h be a height function defined on the vertex set, such that $h(v_1) < h(v_2)$ when $v_1 < v_2$. Assume further that $\widehat{\Delta}$ is a subcomplex of the simplicial realisation of the directed set X. Lemma 2.6 from [6] says that $\widehat{\Delta}$ is n-connected if, for every two vertices v_1, v_2 in $\widehat{\Delta}$ such that $v_1 < v_2$, the ascending link of v_1 relative to v_2 (Definition 6.17) is always $(n-1)$-connected.

Lemma 2.6 from [6] applies to our complexes $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}_i'}$ directly, where the height function h sends a vertex to its cardinality (as in Definition 9.4). It therefore suffices to show that the relative ascending link is n-connected for all n.

We can argue the latter point directly as follows. If

$$v_1 = \{[f_1, D_1], \dots, [f_m, D_m]\}$$

is a vertex of either $\Delta^{\mathcal{E}_i}$ or $\Delta^{\mathcal{E}_i'}$, and $v_1 < v_2$, then we can write

$$v_2 = \bigcup_{k=1}^{m} p_k,$$

where, for $k = 1, \dots, m$, p_k is a pseudovertex having the same support as $\{[f_k, D_k]\}$. The ascending link of v_1 relative to v_2 is homeomorphic to the joins of the ascending links of $\{[f_k, D_k]\}$ relative to p_k, for $k = 1, \dots, m$. (This can be argued exactly as in the proof of Theorem 6.9 from [6].) All of the latter ascending links are contractible if they are nonempty, by the proof of Theorem 9.2 given above. At least one of the latter ascending links is nonempty (since $v_1 \neq v_2$), so the ascending link of v_1 relative to v_2 is contractible, as claimed.

9.3 Γ-finite filtrations of $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}_i'}$

In this subsection, we will describe natural filtrations of the complexes $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}_i'}$. We will denote the acting group by Γ; here Γ might be any of the groups $\{F(S), T(S), V(S)\}$, where $S \in \{S_2, S_3, S_2', S_3'\}$.

Definition 9.4 (Γ-finite filtrations of the complexes $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}_i'}$) Let v be a vertex in $\Delta^{\mathcal{E}_i}$ or $\Delta^{\mathcal{E}_i'}$, or a pseudovertex. We let $|v|$ denote the cardinality of v, which we will call the *height* of v. For $n \geq 1$, we let $\Delta_n^{\mathcal{E}_i}$ denote the subcomplex of $\Delta^{\mathcal{E}_i}$ spanned by vertices of height n or less. Similarly define the subcomplexes $\Delta_n^{\mathcal{E}_i'}$ of $\Delta^{\mathcal{E}_i'}$.

Proposition 9.5 *The group Γ acts on each $\Delta_n^{\mathcal{E}_i}$ (or $\Delta_n^{\mathcal{E}_i'}$, as the case may be) cocompactly, and*

$$\Delta^{\mathcal{E}_i} = \bigcup_{n=1}^{\infty} \Delta_n^{\mathcal{E}_i}.$$

A similar equality is true of $\Delta^{\mathcal{E}_i'}$ and the subcomplexes $\Delta_n^{\mathcal{E}_i'}$.

Proof Let us first note that the Γ-action preserves height, and therefore acts on $\Delta_n^{\mathcal{E}_i}$ (or $\Delta_n^{\mathcal{E}_i'}$). It is easy to see that $\Delta^{\mathcal{E}_i}$ is the union of the subcomplexes in the filtration.

We temporarily let $\mathcal{E}$ denote an arbitrary expansion scheme. Definition 6.12 from [6] describes an action $\star$ of $\mathbb{S}(D, D)$ ($D \in \mathcal{D}_S^+$, or $D \in \mathcal{D}_{\mathrm{gen}}^+$, as in our case) on the set $\mathcal{E}([f, D])$ as

$$h \star v = (f h f^{-1}) \cdot v,$$

where $h \in \mathbb{S}(D, D)$. If the action of $\mathbb{S}(D, D)$ on the simplicial realisation of $\mathcal{E}([f, D])$ is always cocompact, for all D, then $\mathcal{E}$ is said to be $\mathbb{S}$-*finite*.

In the current situation, the group $\mathbb{S}(D, D)$ is isomorphic either to $\mathbb{Z}$ or to the trivial group (when $[D] = [I]$ or $[D] = [[0, \infty)]$, respectively). In either case, the action of $\mathbb{S}(D, D)$ is cocompact. Indeed, the action of $\mathbb{Z}$ on $\mathcal{E}([\mathrm{id}_I, I])$ is by translation (i.e., the integer n moves a vertex u_k to u_{n+k}, where the vertices u_j are as described in Example 6.22). This is clearly cocompact; see Figure 5. This reasoning applies equally to all $[f, D]$ such that $[D] = [I]$ due to the equivariance of the expansion scheme $\mathcal{E}_i$ (or $\mathcal{E}_i'$). If $[D] = [[0, \infty)]$, there is nothing to prove, since the set $\mathcal{E}_i'([f, D])$ is compact. It follows that both $\mathcal{E}_i$ and $\mathcal{E}_i'$ are $\mathbb{S}$-finite.

We can now apply Proposition 6.13 from [6], which says that when an expansion scheme $\mathcal{E}$ is $\mathbb{S}$-finite and $\mathbb{S}$ has finitely many domain types, then the action of Γ on each subcomplex $\Delta_n^{\mathcal{E}}$ is cocompact; this proves that the action of Γ on the filtration is cocompact.

The final equality in the proposition is clear. □

Remark 9.6 We sketch a more direct proof that Γ acts cocompactly.

We assume that $\Gamma = V(S_i)$, the proofs for the other groups being similar. Two vertices v_1 and v_2 are in the same Γ-orbit if and only if they have the same type (Definition 4.6). In the current context, the latter condition is equivalent to having the same height.

Now assume that Γ fails to act cocompactly on $\Delta_n^{\mathcal{E}_i}$, for some n. We note that the dimension of $\Delta_n^{\mathcal{E}_i}$ is no more than $n-1$, since a simplex in $\Delta_n^{\mathcal{E}_i}$ is an ascending chain

$$v_0 < v_1 < v_2 < \cdots < v_k,$$

and the height function strictly increases along such chains. Thus, assuming that the action of Γ is not cocompact, there are infinitely many Γ-orbits of k-simplices, for some k. Since there are only finitely many Γ-orbits of vertices, this implies that there is a vertex $v' = \{b_1, \dots, b_\ell\}$ such that infinitely many Γ-orbits of k-simplices have v' as their minimal vertex. This, however, sets up the contradiction, since all of the k-simplices in question are obtained by $\mathcal{E}_i$-expansion from v', and there are only finitely many such $\mathcal{E}_i$-expansions modulo the action $\star$.

The details of the remainder of the argument follow that of the proof of Proposition 6.13 from [6].

9.4 The F_∞ property for $V(S_n)$ and $V(S_n')$

Definition 9.7 (contracting pseudovertices) Let $\mathcal{E}$ be an arbitrary expansion scheme. We say that a pseudovertex v is *contracting relative to* $\mathcal{E}$ if v has the same type as some $w \in \mathcal{E}(b)$, where $b \in \mathcal{B}$. (Recall that "same type" was defined in Definition 4.6.)

Definition 9.8 (rich in contractions) Let $\mathcal{E}$ be an expansion scheme. We say that $\mathcal{E}$ is *rich in contractions* if there is some constant C such that, if v is a pseudovertex of height at least C, then there is some contracting pseudovertex v' such that $v' \subseteq v$.

Theorem 9.9 ([6, Theorem 8.2], groups of type F_∞) *Let $\mathbb{S}$ be an S-structure with finitely many domain types, such that the group $\mathbb{S}(D, D)$ has type F_∞ for $D \in \mathcal{D}^+$. Let $\mathcal{E}$ be an expansion scheme such that*

(1) *$\mathcal{E}$ is n-connected for all n;*

(2) *$\mathcal{E}$ is rich in contractions;*

(3) *each set $\mathcal{E}(b)$ $(b \in \mathcal{B})$ is finite.*

The group Γ_S has type F_∞.

Theorem 9.10 *The groups $V(S_i)$ and $V(S_i')$ are of type F_∞, for $i = 2, 3$.*

Proof Our strategy is to apply the proof of Theorem 9.9 (Theorem 8.2 from [6]) to the groups Γ. (We note that the groups Γ_S under consideration in Theorem 9.9 are analogous to Thompson's group V, in that there is no assumption that Γ_S preserves a linear or cyclic order.) Let us note that condition (3) is violated, since the sets $\mathcal{E}_i(b)$ and $\mathcal{E}_i'(b)$ are not finite when $b = [f, D]$ and $[D] = [I]$, so the statement does not apply directly.

We have already seen that $\mathcal{E}_i$ and $\mathcal{E}_i'$ are n-connected expansion schemes for all n.

We claim that the expansion scheme $\mathcal{E}_i$ is rich in contractions with constant $C = 2$ when $i = 2$ or 3. Let $\{[f_1, D_1], [f_2, D_2]\} \subseteq \mathcal{B}$ be a pseudovertex. Since $D_1, D_2 \in \mathcal{D}^+_{\text{gen}}$, we have $D_1 = \omega_1 I$ and $D_2 = \omega_2 I$, for some words $\omega_1, \omega_2 \in \{A, B\}^*$. Thus,

$$[f_n, D_n] = [f_n, \omega_n I] = [f_n \omega_n, I],$$

for $n = 1, 2$. Define g on $[0, 1)$ by the rule

$$g(x) = \begin{cases} f_1 \omega_1 a(x) & \text{if } x \in [0, 1/2), \\ f_2 \omega_2 b(x) & \text{if } x \in [1/2, 1). \end{cases}$$

The pseudovertex $\{[g, I]\}$ expands to

$$\begin{aligned} \{[g, AI], [g, BI]\} &= \{[f_1 \omega_1 a, AI], [f_2 \omega_2 b, BI]\} \\ &= \{[f_1 \omega_1, I], [f_2 \omega_2, I]\} \\ &= \{[f_1, D_1], [f_2, D_2]\}. \end{aligned}$$

This proves the claim.

The expansion scheme $\mathcal{E}_i'$ is also rich in contractions with constant $C = 2$. If $\{[f_1, D_1], [f_2, D_2]\} \subseteq \mathcal{B}$ is a pseudovertex and D_1, D_2 have the same domain type as I, then the proof of the previous paragraph shows that a contraction can be performed on $\{[f_1, D_1], [f_2, D_2]\}$. The only remaining case to consider is when $[D_1] = [I]$ and $[D_2] = [[0, \infty)]$. We will write R in place of $[0, \infty)$, to simplify notation. In this case, $D_1 = \omega_1 I$ and $D_2 = T^m R$, where $\omega_1 \in \{A, B, T\}^*$ and $m \geq 0$. We have

$$[f_1, \omega_1 I] = [f_1\omega_1, I] \quad \text{and} \quad [f_2, T^m R] = [f_2 T^m, R].$$

Define $g : [0,\infty) \to [0,\infty)$ as

$$g(x) = \begin{cases} f_1\omega_1(x) & \text{if } x \in [0,1), \\ f_2 T^{m-1}(x) & \text{if } x \in [1,\infty). \end{cases}$$

The pseudovertex $\{[g, R]\}$ expands to

$$\begin{aligned} \{[g,R]\} &= \{[g,I],[g,TR]\} \\ &= \{[f_1\omega_1, I],[f_2T^{m-1}, TR]\} \\ &= \{[f_1, D_1],[f_2, T^m R]\} \\ &= \{[f_1, D_1],[f_2, D_2]\}. \end{aligned}$$

It follows that $\{[f_1, D_1],[f_2, D_2]\}$ is also a contracting vertex relative to $\mathcal{E}_i'$.

The assumption that $\mathcal{E}(b)$ is always finite is used in the proof of Theorem 9.9 in three ways:

(1) to prove that Γ acts cocompactly on the complexes $\Delta_n^{\mathcal{E}}$;

(2) to prove that the cell stabilisers have type F_∞, and

(3) to define a certain constant C_0.

We have already established (1) and (2) by other means: indeed, cell stabilisers are virtually finitely generated free abelian groups, and therefore have type F_∞, and the cocompactness of the actions on the complexes $\Delta_n^{\mathcal{E}_i}$ and $\Delta_n^{\mathcal{E}_i'}$ was proved as part of Proposition 9.5. The constant C_0 is the largest height (i.e., cardinality) of a contracting pseudovertex. Clearly we have an independent bound of $C_0 = 3$ when $S \in \{S_2, S_2'\}$, or $C_0 = 5$ when $S \in \{S_3, S_3'\}$. (Refer to the definitions of $\mathcal{E}_i$ and $\mathcal{E}_i'$ in Example 6.22.) □

Remark 9.11 We will offer a sketch of the argument here. This sketch is intended to remove some of the dependence on Theorem 9.9.

We check the hypotheses of Brown's finiteness criterion (Theorem 9.1). First, we note that $\Delta^{\mathcal{E}_i}$ and $\Delta^{\mathcal{E}_i'}$ are contractible by Theorem 9.2. It is clear that the relevant actions are cellular. Properties (3)(a) and (b) are settled in Proposition 9.5. Cell stabilisers are virtually finitely generated free abelian (and therefore of type F_∞) by Proposition 4.13.

This leaves only (3)(d) to check; i.e., we must show, for each $n \in \mathbb{N}$, that $\Delta_k^{\mathcal{E}_i}$ and $\Delta_k^{\mathcal{E}_i'}$ are $(n-1)$-connected for sufficiently large k. The proof of the latter follows a now-standard strategy: we show that the descending link of a vertex becomes highly connected as the height of the vertex increases. A few basics of this strategy are summarised in Subsection 7.2 of [6], although the methods of argument go back to [3; 1].

We consider $\Delta^{\mathcal{E}_2}$; the other cases are similar. Let v be a vertex of height k in $\Delta^{\mathcal{E}_2}$. The *descending link* of v is its link in $\Delta_k^{\mathcal{E}_2}$. Our analysis of the descending link uses the nerve theorem (as it appears in [2]; the nerve theorem is also Theorem 2.10 in [6]). Let

$$v = \{b_1, \dots, b_k\}.$$

We cover the descending link of v by a number of subcomplexes, called *partitioned downward links*, which are each determined by a partition of v, and which we now define.

Let $\mathcal{P}$ be a partition of v. The *partitioned downward star* $\mathrm{st}_\downarrow(v_\mathcal{P})$ (Definition 7.7 from [6]), is the subcomplex of $\Delta_k^{\mathcal{E}_2}$ consisting of the vertex v and all simplices resulting from $\mathcal{E}_2$-contractions that are supported within members of $\mathcal{P}$. For instance, if

$$\mathcal{P} = \{\{b_1, b_2\}, \{b_3, \dots, b_k\}\},$$

then a contraction supported on the subset $\{b_1, b_2\}$, or on the subset $\{b_3, b_4, b_7\}$ (if $k \geq 7$) (or indeed a combination of such contractions), results in a simplex of $\mathrm{st}_\downarrow(v_\mathcal{P})$, but a contraction supported on $\{b_2, b_3\}$ would not. We then define the partitioned downward link $\mathrm{lk}_\downarrow(v_\mathcal{P})$ as the link of v in $\mathrm{st}_\downarrow(v_\mathcal{P})$.

For each contracting pseudovertex $w \subseteq v$, we let

$$\mathcal{P}_w = \{v - w, w\}.$$

(We note that, in the current context, "contracting pseudovertex" is the same as "pseudovertex with two or three members", by the description of $\mathcal{E}_2$ from Example 6.22.) The collection

$$\mathcal{C} = \{\mathrm{lk}_\downarrow(v_{\mathcal{P}_w}) \mid w \text{ is a contracting pseudovertex}\}$$

is a cover of $\mathrm{lk}_\downarrow(v)$. We apply the nerve theorem to $\mathcal{C}$. The intersection of two members of $\mathcal{C}$ is another partitioned downward link,

$$\mathrm{lk}_\downarrow(v_{\mathcal{P}_{w'}}) \cap \mathrm{lk}_\downarrow(v_{\mathcal{P}_{w''}}) = \mathrm{lk}_\downarrow(v_{\mathcal{P}_{w'} \wedge \mathcal{P}_{w''}}),$$

where $\mathcal{P}_{w'} \wedge \mathcal{P}_{w''}$ is the coarsest common refinement of $\mathcal{P}_{w'}$ and $\mathcal{P}_{w''}$. The generalisation to finite intersections is straightforward.

For a partition $\mathcal{P} = \{P_1, \dots, P_\ell\}$ of v, there is a natural join structure (see Corollary 7.9 from [6]),

$$\mathrm{lk}_\downarrow(v_\mathcal{P}) \cong \mathop{\ast}_{j=1}^{\ell} \mathrm{lk}_\downarrow(P_j),$$

where the latter descending links depend only on the types of the pseudovertices P_j. (In the current case, the type is entirely determined by the cardinality.) Recall that, if X_1 and X_2 are n_1-connected and n_2-connected complexes (respectively), then the join $X_1 * X_2$ is (n_1+n_2+2)-connected. It follows that $\mathrm{lk}_\downarrow(v_\mathcal{P})$ is at least as connected as the most highly connected factor $\mathrm{lk}_\downarrow(P_j)$.

Finally, we note that a pseudovertex of height two or more has a nonempty descending link (since every such pseudovertex contains a contracting pseudovertex). This gives us the base case of an induction; the above considerations allow us to prove inductively that pseudovertices of increasing height have increasing connectivity. The actual induction is done in the proof of Theorem 8.2 from [6]; we omit further details. We will consider a similar induction in more detail in the next subsection.

9.5 The F_∞ property for the remaining groups

Theorem 9.12 *The groups $F(S)$, where $S \in \{S_2, S_3, S_2', S_3'\}$, and $T(S)$, where $S \in \{S_2, S_3\}$, have type F_∞.*

Proof We consider the group $F(S_2)$. The proofs that the other groups have type F_∞ differ in minor details.

We turn to an analysis of the descending link; all of the other ingredients of the proof can be assembled exactly as in Remark 9.11. Let

$$v = \{b_1, b_2, \dots, b_k\}$$

be either a vertex of $\Delta^{\mathcal{E}_2}$, or a pseudovertex. We assume that the b_i are linearly ordered, in the following sense: Each $b_i = [f_i, D_i]$, for appropriate f_i and $D_i \in \mathcal{D}^+_{\text{gen}}$, where $f_i : D_i \to [0,1)$ is a locally S_2-embedding that is, moreover, continuous and increasing. We assume that $f_1(D_1), f_2(D_2), \dots, f_k(D_k)$ are arranged from left to right. With this assumption, each $\mathcal{E}_2$-contraction must be performed on two or three consecutive b_i. For a subset $K \subseteq \{1, \dots, k\}$, we define

$$\mathcal{P}_K = \{\{b_j \mid j \in K\}, \{b_j \mid j \notin K\}\}.$$

We then define

$$\mathcal{C} = \bigl\{\mathrm{lk}_\downarrow(v_{\mathcal{P}_K}) \mid K \in \{\{1,2\}, \{2,3\}, \{1,2,3\}, \{2,3,4\}, \{3,4,5\}\}\bigr\}.$$

(If $k < 5$, then the possible subsets K are restricted accordingly.)

We claim that $\mathcal{C}$ is a cover of $\mathrm{lk}_\downarrow(v)$. Indeed, let σ be a simplex in $\mathrm{lk}_\downarrow(v)$. Thus, there is an increasing sequence

$$v_0 < v_1 < v_2 < \cdots < v_{\ell-1} < v_\ell = v,$$

where each v_α is obtained by $\mathcal{E}_2$-expansion from v_0, and

$$\sigma = v_0 < v_1 < \cdots < v_{\ell-1}.$$

If $v_0 = \{b_0', b_1', \dots, b_q'\}$, where the members are linearly ordered, then there is a leftmost b_β' that is expanded when we pass from v_0 to v. In expanding at b_β', we replace b_β' with either two or three pairs from $\mathcal{B}$. The latter will occur consecutively in v. Thus, the result of expanding at b_β' will contribute either $\{b_\alpha, b_{\alpha+1}, b_{\alpha+2}\}$ or $\{b_\alpha, b_{\alpha+1}\}$ to v, for some α. All other expansions from v_0 to v will contribute a disjoint subset of b_i's to v. It follows easily from this that σ is contained in at least one member of the cover $\mathcal{C}$.

For instance, if the expansion at b_β' contributes b_1 and b_2 to v, any other expansion from v_0 must contribute some subset of $\{b_3, \dots, b_k\}$. Thus, in this case, $\sigma \subseteq \mathrm{lk}_\downarrow(v_{\mathcal{P}_{\{1,2\}}})$. If b_β' contributes b_2 and b_3, then $\sigma \subseteq \mathrm{lk}_\downarrow(v_{\mathcal{P}_{\{2,3\}}})$. If b_β' contributes b_m and b_{m+1}, for some $m \geq 3$, then $\sigma \subseteq \mathrm{lk}_\downarrow(v_{\mathcal{P}_{\{1,2\}}})$ (since, indeed, all expansions contribute some subset of $\{b_3, \dots, b_k\}$ under this hypothesis).

Now we establish the connectivity of the descending link, as a function of the height k. We note first that $\mathrm{lk}_{\downarrow}(v)$ is nonempty provided that $k \geq 2$. It follows from this that each $\mathrm{lk}_{\downarrow}(v_{P_K})$ is connected when $k \geq 7$, since each is a join of two nonempty complexes. Now, if $k \geq 7$, then $\mathrm{lk}_{\downarrow}(v)$ is connected, since it is covered by a collection $\mathcal{C}$ of nonempty subcomplexes, which have a nonempty intersection. (A contraction at $\{b_6, b_7\}$ lies in all of the partitioned descending links simultaneously.)

In general, $\mathrm{lk}_{\downarrow}(v)$ is n-connected provided that $k \geq 5n+7$. We have proved this already for $n=-1$ and $n=0$. Assume that the result is true for n. We consider a vertex v of height k at least $5n+12$. Each $\mathrm{lk}_{\downarrow}(v_{P_K})$ is $(n+1)$-connected, since each is a join of two complexes: one nonempty and one isomorphic to the descending link of a vertex of height at least $5n+7$, and therefore n-connected by induction. Moreover, any subcollection of $\mathcal{C}$ containing two of more members intersects in a subcomplex that is at least n-connected. (Any such intersection is a join, and one of the factors of the join is the descending link on $\{b_\gamma, \dots, b_k\}$, where $\gamma \leq 6$.)

By the nerve theorem [2], $\mathrm{lk}_{\downarrow}(v)$ is $(n+1)$-connected if t-fold intersections of the cover are $(n-t+2)$-connected and the nerve of the cover is $(n+1)$-connected. Since the nerve is easily seen to be a four-dimensional simplex, and t-fold intersections have the required connectivity (by the previous paragraph), $\mathrm{lk}_{\downarrow}(v)$ is $(n+1)$-connected, completing the induction.

By well-established principles (as summarised in Proposition 7.6 from [6], for instance), the connectivity of the subcomplex $\Delta_k^{\mathcal{E}_i}$ tends to infinity as k increases, completing the proof. □

10 The case of the Lodha–Moore group

Recall that $F(S_2')$ is the group of homeomorphisms of $[0,\infty)$ that is locally determined by the inverse semigroup generated by the set

$$\{A, B, C_2, T\},$$

where all of these are as defined in Definition 2.6. In this section, we will prove the following theorem:

Theorem 10.1 *The Lodha–Moore group G is isomorphic to an ascending HNN extension of $F(S_2')$ in which the stable letter is the translation $t \mapsto t+1$.*

In particular, G has type F_∞.

Throughout this section, we let $F(S_2')$ act on the entire real line, by simply defining each element of $F(S_2')$ to be the identity on $(-\infty, 0]$.

Definition 10.2 [8; 9] The *Lodha–Moore group* G is the group of homeomorphisms of the real line generated by three transformations, denoted by a, b, and c, and defined as

$$a(t) = t+1, \qquad b(t) = \begin{cases} t & \text{if } t \leq 0, \\ \frac{t}{1-t} & \text{if } 0 \leq t \leq \frac{1}{2}, \\ 3 - \frac{1}{t} & \text{if } \frac{1}{2} \leq t \leq 1, \\ t+1 & \text{if } 1 \leq t, \end{cases} \qquad c(t) = \begin{cases} t & \text{if } t \leq 0, \\ \frac{2t}{t+1} & \text{if } 0 \leq t \leq 1, \\ t & \text{if } 1 \leq t. \end{cases}$$

Remark 10.3 We are using the notation a, b, and c for the generators of the Lodha–Moore group G, following the practice in [8; 9]. Of course, this notation contradicts our own practice of using lowercase "x" to denote the inverse of the partial transformation X (for $X \in \{A, B, C_2\}$). We will therefore continue to use the uppercase letters to denote our inverse semigroup generators, but will use "X^{-1}" to denote the inverse of X. For the remainder of the paper, a, b, and c refer to the generators in Definition 10.2.

Let us note that

$$b = A^{-1} \cup TB^{-1} \cup T^2T^{-1} \quad \text{and} \quad c = C \cup TT^{-1},$$

where union is interpreted in an obvious sense: b agrees with A^{-1} on $[0, 1/2)$ (the domain of A^{-1}), with TB^{-1} on $[1/2, 1)$ (the domain of TB^{-1}), etc. (The trivial action on negative numbers is implied in both of these definitions of b and c).

It follows directly that $\langle b, c \rangle \leq F(S_2')$.

Lemma 10.4 *Let $\mathcal{P}$ be a partition of $[0, \infty)$ into generating domains (Definition 3.1). There is an element $g \in \langle b, c \rangle$ such that either*

(1) *the leftmost member of $g(\mathcal{P})$ is $[0, 1)$, or*

(2) *the leftmost member of $g(\mathcal{P})$ is $[0, \infty)$.*

In either case, it can be arranged that the restriction $g_{|P} : P \to g(P)$ is in S_2', for each $P \in \mathcal{P}$, and that $g(\mathcal{P})$ is also a partition of $[0, \infty)$ into generating domains.

Proof It is rather clear from Definition 3.1 that a partition of $[0, \infty)$ by generating domains has members of two types:

(i) a (necessarily unique) member of the form $[n, \infty)$, for some integer $n \geq 0$, and

(ii) a collection of generating domains of the form $T^{\alpha}\omega \cdot [0, 1)$, where $\omega \in \{A, B\}^*$ (i.e., ω is a positive, possibly empty, word in the alphabet $\{A, B\}$), and $0 \leq \alpha < n$. Each domain of this form is contained in $[\alpha, \alpha + 1)$.

If $n = 0$, then $\mathcal{P} = \{[0, \infty)\}$, and we can simply let $g = \mathrm{id}_{\mathbb{R}}$. It is clear that g satisfies all of the required properties. (This uses the fact that the restriction of id to each generating domain is a member of S_2', which follows from Remark 2.2.)

Now suppose that $n > 0$. We prove the lemma by induction on the number m of generating domains of $\mathcal{P}$ that are contained in $[0, 1)$. If $m = 1$, then we can let $g = \mathrm{id}_{\mathbb{R}}$. Now suppose that $m > 1$. There are two types of generating domains of $\mathcal{P}$ that are contained in $[0, 1)$: those of the form $A\omega \cdot [0, 1)$, and those of the form $B\omega \cdot [0, 1)$. Both types must be present, since those of the first form are contained in $[0, 1/2)$, while those of the second form are contained in $[1/2, 1)$. We apply the transformation b from Definition 10.2. Using the description of b from Remark 10.3, we find that each domain $A\omega \cdot [0, 1)$ is carried to $\omega \cdot [0, 1)$, and each domain $B\omega \cdot [0, 1)$ is sent to $T\omega \cdot [0, 1)$. The element b acts on every generating domain of $\mathcal{P}$ in $[1, \infty)$ by the translation T. It follows that b carries the set $\mathcal{P}$ of generating domains to another set of generating domains, $b(\mathcal{P})$. We note that the restriction of b to each member of $\mathcal{P}$

is a member of the inverse semigroup S_2', and that the number m is reduced in the process (since each domain of the form $B\omega \cdot [0,1)$ is carried outside of $[0,1)$). We can then apply the inductive hypothesis to the partition $b(\mathcal{P})$ to produce an element $g \in \langle b, c \rangle$ such that $gb(\mathcal{P})$ has the required form, while the restriction of gb to each member of $\mathcal{P}$ is a member of S_2'. This completes the induction. □

Remark 10.5 The element g produced in the proof of Lemma 10.4 is always a nonnegative power of b.

Proposition 10.6 *The group $F(S_2')$ is a subgroup of the Lodha–Moore group G.*

Proof Let $f \in F(S_2')$. There is a partition $\mathcal{P}_1$ of $[0,\infty)$ into finitely many generating domains such that $f_{|P} \in S_2'$ for each $P \in \mathcal{P}_1$, and such that $f(\mathcal{P}) := \mathcal{P}_2$ is also a partition of $[0,\infty)$ into generating domains. (This follows from Definition 2.8 and Remark 3.4.)

We will prove that $f \in G$ by induction on $|\mathcal{P}_1|$. If $|\mathcal{P}_1| = 1$, then $f = \mathrm{id}_{\mathbb{R}}$, and $f \in G$. Let $|\mathcal{P}_1| = m$. Lemma 10.4 allows us to find $g_1, g_2 \in G$ such that, for $i = 1, 2$,

(1) $g_i(\mathcal{P}_i)$ is a partition of $[0,\infty)$ into generating domains;

(2) $g_{i|P} \in S_2'$, for each $P \in \mathcal{P}_i$;

(3) the leftmost member of $g_i(\mathcal{P}_i)$ is $[0,1)$.

It follows from this that the element $g_2 f g_1^{-1}$ carries the generating domain $[0,1)$ to the generating domain $[0,1)$ by a member of S_2'. By the characterisation of $\mathbb{S}(I,I)$ (Theorem 4.3), the restriction of $g_2 f g_1^{-1}$ to $[0,1)$ is C^k, for some $k \in \mathbb{Z}$. It follows that $c^{-k} g_2 f g_1^{-1}$ is equal to the identity on $[0,1)$. (Here we are using "c" to refer to the generator of G, as in Definition 10.2.) We note that $c^{-k} g_2 f g_1^{-1} \in F(S_2')$ by construction, and that the domain and range of $c^{-k} g_2 f g_1^{-1}$ are both partitioned into m pieces, each of which is a generating domain, and such that $c^{-k} g_2 f g_1^{-1}$ matches these pieces by members of S_2'.

It follows from this that $a^{-1} c^{-k} g_2 f g_1^{-1} a$ is a member of $F(S_2')$ that is similarly defined on $m-1$ pieces. It follows by induction that $a^{-1} c^{-k} g_2 f g_1^{-1} a = g_3$, for some $g_3 \in G$. Solving the latter equation for f, we find that $f \in G$, completing the induction. □

Proof of Theorem 10.1 Since $\langle b, c \rangle \le F(S_2')$ by Remark 10.3 and $F(S_2') \le G$ by Proposition 10.6, $\langle a, F(S_2') \rangle = G$.

The map $g \mapsto aga^{-1}$ determines an injective endomorphism of $F(S_2')$. There is an induced homomorphism from the resulting HNN extension $F(S_2')_{*a}$ to G. Clearly this map is surjective by the previous paragraph. Injectivity follows from the fact that each nontrivial normal form $a^{-k} f a^{\ell}$ $(k, \ell \ge 0)$ maps to a nontrivial element of G. This proves that G is the required HNN extension of $F(S_2')$.

The fact that G has type F_∞ now follows from the fact that $F(S_2')$ has type F_∞ by Theorem 9.12 and from the fact that the ascending HNN extension of a type F_∞ group also has type F_∞. □

Acknowledgements

I would like to thank Andy Moawad for suggesting corrections to an earlier version of this manuscript. Matt Zaremsky also suggested corrections, and pointed me to his article [10] with Yash Lodha, in which the authors proved that a certain group of piecewise projective homeomorphisms of the circle has type F_∞.

Their results anticipate the proof offered here that $T(S_2)$ has type F_∞. The question of whether a certain "V-like" Lodha–Moore group has type F_∞ is a conjecture in [10], which our proof that $V(S_2)$ has type F_∞ appears to resolve. The referee notes that, since the Lodha–Moore group is an ascending HNN extension of $F(S_2')$, the latter group must also be nonamenable, and have no free subgroups. Theorem 9.12 proves that $F(S_2')$ also has type F_∞. I do not know whether $F(S_2')$ is a *new* group with the indicated properties, however.

I thank the referee for suggesting numerous corrections and clarifications to an earlier draft of this paper.

References

[1] **M Bestvina**, **N Brady**, *Morse theory and finiteness properties of groups*, Invent. Math. 129:3 (1997) 445–470 MR

[2] **A Björner**, *Topological methods*, from "Handbook of combinatorics, II" (R L Graham, M Grötschel, L Lovász, editors), Elsevier Sci. B. V., Amsterdam (1995) 1819–1872 MR

[3] **K S Brown**, *Finiteness properties of groups*, J. Pure Appl. Algebra 44:1-3 (1987) 45–75 MR

[4] **A Buss**, **R Exel**, **R Meyer**, *Inverse semigroup actions as groupoid actions*, Semigroup Forum 85:2 (2012) 227–243 MR

[5] **J W Cannon**, **W J Floyd**, **W R Parry**, *Introductory notes on Richard Thompson's groups*, Enseign. Math. (2) 42:3-4 (1996) 215–256 MR

[6] **D S Farley**, **B Hughes**, *Finiteness properties of locally defined groups* (2020) arXiv 2010.08035

[7] **J M Howie**, *An introduction to semigroup theory*, L. M. S. Monographs 7, Academic, London (1976) MR

[8] **Y Lodha**, *A nonamenable type* F_∞ *group of piecewise projective homeomorphisms*, J. Topol. 13:4 (2020) 1767–1838 MR

[9] **Y Lodha**, **J T Moore**, *A nonamenable finitely presented group of piecewise projective homeomorphisms*, Groups Geom. Dyn. 10:1 (2016) 177–200 MR

[10] **Y Lodha**, **M C B Zaremsky**, *The BNSR-invariants of the Lodha–Moore groups, and an exotic simple group of type* F_∞, Math. Proc. Cambridge Philos. Soc. 174:1 (2023) 25–48 MR

[11] **N Monod**, *Groups of piecewise projective homeomorphisms*, Proc. Natl. Acad. Sci. USA 110:12 (2013) 4524–4527 MR

[12] **M H A Newman**, *On theories with a combinatorial definition of "equivalence"*, Ann. of Math. (2) 43 (1942) 223–243 MR

[13] **A Y Olshanskii**, **M V Sapir**, *Non-amenable finitely presented torsion-by-cyclic groups*, Publ. Math. Inst. Hautes Études Sci. 96 (2002) 43–169 MR

[14] **J G Ratcliffe**, *Foundations of hyperbolic manifolds*, 2nd edition, Graduate Texts in Mathematics 149, Springer (2006) MR

[15] **M Stein**, *Groups of piecewise linear homeomorphisms*, Trans. Amer. Math. Soc. 332:2 (1992) 477–514 MR

DANIEL S. FARLEY farleyds@miamioh.edu

Department of Mathematics and Statistics, Miami University, Oxford, OH, United States

Received: March 5, 2023 Revised: June 14, 2024

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2026 is US $795/year for the electronic version, and $1170/year (+$80, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 26 Issue 4 (pages 1229–1596) 2026