TWISTED BRIN–THOMPSON GROUPS

JAMES BELK AND MATTHEW C. B. ZAREMSKY

ABSTRACT. We construct a family of infinite simple groups that we call *twisted Brin–Thompson groups*, generalizing Brin's higher-dimensional Thompson groups sV ($s \in \mathbb{N}$). We use twisted Brin–Thompson groups to prove a variety of results regarding simple groups. For example, we prove that every finitely generated group embeds quasi-isometrically as a subgroup of a two-generated simple group, strengthening a result of Bridson. We also produce examples of simple groups that contain every sV and hence every right-angled Artin group, including examples of type F_∞ and a family of examples of type F_{n-1} but not of type F_n, for arbitrary $n \in \mathbb{N}$. This provides the second known infinite family of simple groups distinguished by their finiteness properties.

INTRODUCTION

The Brin–Thompson groups sV ($s \in \mathbb{N}$) are a family of groups introduced by Brin in [Bri04] as higher-dimensional generalizations of the classical Thompson group $V = 1V$. Each sV acts by homeomorphisms on the Cantor space $\mathfrak{C}^s$, where $\mathfrak{C} = \{0,1\}^\omega$ is the usual Cantor set. Brin proved that all of these groups are finitely generated and simple [Bri10], and that $2V$ is finitely presented [Bri05]. Hennig and Matucci later showed that all of the groups sV are finitely presented [HM12], and Bleak and Lanoue proved that sV is not isomorphic to $s'V$ for $s \neq s'$ [BL10].

In this paper we introduce a new family of infinite simple groups that we call *twisted Brin–Thompson groups*. Given any group G acting faithfully on a countable set S, the associated twisted Brin–Thompson group SV_G acts by homeomorphisms on the Cantor space $\mathfrak{C}^S$. The Brin–Thompson groups themselves correspond to the case where S is finite and G is trivial.

When S is infinite, the group SV_G is "large" in the sense that it contains sV for all $s \in \mathbb{N}$. The first author, Bleak, and Matucci have proven that every finitely generated right-angled Artin group embeds into some sV [BBM20], and therefore SV_G contains every finitely generated right-angled Artin group when S is infinite. It follows moreover from [BBM20] that such SV_G contain all virtually special groups, for example all finitely generated Coxeter groups and many word hyperbolic groups. Despite this "largeness", twisted Brin–Thompson groups can have surprisingly good finiteness properties. For example, in Section 3 we prove:

Theorem A. *The group SV_G is finitely generated if and only if G is finitely generated and the action of G on S has finitely many orbits.*

In particular, if the action of G on S has m orbits and G has n generators, then we prove that SV_G is generated by two copies of G and at most $m + n - 1$ copies of the Brin–Thompson group $2V$. The groups SV_G are also vigorous in the sense of Bleak, Elliott, and Hyde [Hyd17, BEH20], so it follows that SV_G can be generated by two elements whenever it is finitely generated. Indeed, a related generalization of the Brin–Thompson groups has been investigated in this context by Hyde [Hyd17].

There is a natural monomorphism $\iota_\varnothing \colon G \to SV_G$, corresponding to the canonical action of G on the Cantor space $\mathfrak{C}^S$. For G and SV_G finitely generated, we prove in Section 4 that this embedding is geometrically well-behaved:

Theorem B. *Suppose G and SV_G are finitely generated. Then there exists a coarse Lipschitz map $\rho \colon SV_G \to G$ such that $\rho \circ \iota_\varnothing$ is the identity. In particular, the group SV_G quasi-retracts onto G, and $\iota_\varnothing$ is a quasi-isometric embedding of G into SV_G.*

See Section 4 for the relevant definitions. Since every finitely generated group G acts faithfully and transitively on *some* countable set S (e.g., $S = G$), and since SV_G is simple (Theorem 3.4), Theorems A and B immediately yield the following:

Theorem C. *Every finitely generated group embeds quasi-isometrically as a subgroup of a finitely generated (indeed, two-generated) simple group.* ☐

Hall proved that every finitely generated group embeds into a finitely generated simple group [Hal74]. Schupp [Sch76] and independently Goryushkin [Gor74] showed that the simple group can be taken to be two-generated, with Schupp's generators having orders two and three. However, it was not previously known that such embeddings could be made quasi-isometric. By the result of [BEH20], the two generators for SV_G can be taken to have finite order, with one having order two, but we do not know whether the second generator can be made to have order three. The above result also strengthens a theorem of Bridson, who proved that every finitely generated group can be quasi-isometrically embedded as a subgroup of a finitely generated group that has no proper finite index subgroups [Bri98].

As a remark, in our proof of Theorem B, for our particular choices of finite generating sets, the quasi-isometric embedding is actually an isometric embedding. The fact that every finitely generated group embeds isometrically as a subgroup of a finitely generated simple group has also been proved independently by Darbinyan and Steenbock in [DS20].

We can also prove some further finiteness properties for SV_G in the case where the action of G on S is sufficiently transitive. Specifically, recall that a group G of permutations of a set S is *oligomorphic* if for every $k \geq 1$ there are only finitely many orbits of k-element subsets of S [Cam90]. The proof of the following theorem occupies Sections 5, 6, and 7:

Theorem D. *Let G be an oligomorphic group of permutations of a countable set S. Let $n \in \mathbb{N} \cup \{\infty\}$, and suppose that*

    (i) *G is of type F_n, and*

    (ii) *The stabilizer in G of every finite subset of S is of type F_n.*

Then SV_G is of type F_n as well.

Here we say that a group is of *type F_n* if it acts geometrically (i.e., properly and cocompactly) on some $(n-1)$-connected CW complex. (A more standard definition is that a group is of type F_n if it admits a classifying space with finite n-skeleton, but this is equivalent, for example thanks to Brown's Criterion [Bro87].) For example, a group is of type F_1 if and only if it is finitely generated, and a group is of type F_2 if and only if it is finitely presented. A group is of *type F_∞* if it is of type F_n for all $n \in \mathbb{N}$.

Brown and Geoghegan proved that Thompson's group F is of type F_∞ [BG84], and Brown proved the same result for Thompson's groups T and V [Bro87]. Kochloukova, Martínez-Pérez, and Nucinkis proved that the Brin–Thompson groups $2V$ and $3V$ are of type F_∞ [KMPN13], and Fluch, Marschler, Witzel, and the second author extended this result to all of the Brin–Thompson groups [FMWZ13]. Our proof of the above theorem generalizes their techniques, and involves constructing a similar complex and analyzing the descending links, though the fact that S may be infinite complicates the situation considerably.

We can use the above theorem to produce many new simple groups of type F_∞. The following corollary gives one example, which to the best of our knowledge is the first example of a finitely presented simple group that contains every right-angled Artin group.

Corollary E. *If G is Thompson's group F and S is the set of dyadic rationals in $(0,1)$, then SV_G is a simple group of type F_∞ that contains sV for all $s \in \mathbb{N}$, and hence contains every right-angled Artin group.*

Proof. Thompson's group F is of type F_∞ [BG84], and it acts transitively on k-element sets of dyadic rationals in $(0,1)$ for every k [CFP96, Lemma 4.2]. The stabilizer in F of any set of k dyadic rationals is isomorphic to a direct product of $k+1$ copies of F, and therefore such stabilizers are of type F_∞ as well. The result now follows from Theorem D. $\square$

Note also that SV_G is of type F_∞ whenever the set S is finite. When G is non-trivial, the methods of Bleak and Lanoue from [BL10] can be used to show that such a group does not belong to the Brin–Thompson family (since the corresponding groups of germs are non-abelian), but it would be interesting to classify such groups up to isomorphism.

Alonso proved that any quasi-retract of a group of type F_n must be of type F_n [Alo94]. By Theorem B, it follows that G must be of type F_n whenever SV_G is. Combining this with Theorem D, we obtain the following:

Theorem F. *Let G be an oligomorphic group of permutations of a set S. Let $n \in \mathbb{N}$, and suppose that*

(i) *G is of type F_{n-1} but not of type F_n, and*

(ii) *The stabilizer in G of any finite subset of S is of type F_{n-1}.*

Then SV_G is of type F_{n-1} but not of type F_n. $\square$

We can use this theorem to produce new simple groups that are of type F_{n-1} but not of type F_n.

Corollary G. *If G is the nth Houghton group H_n and $S = \{1, \ldots, n\} \times \mathbb{N}$, then SV_G is a simple group of type F_{n-1} that is not of type F_n.*

Proof. Brown proved that H_n is of type F_{n-1} but not of type F_n [Bro87]. It is easy to see that H_n acts transitively on k-element subsets of S for every k, and the stabilizer of any k-element subset is isomorphic to $H_n \times \Sigma_k$, where Σ_k denotes the kth symmetric group. In particular, such stabilizers are all of type F_{n-1}. The result now follows from Theorem F. $\square$

The family of groups in the corollary above is the second known family $\{G_n\}$ of simple groups such that each G_n is of type F_{n-1} but not F_n, the first being the family arising from Röver–Nekrashevych groups described by Skipper, Witzel, and the second author in [SWZ19]. This also gives the third known infinite family of pairwise non-quasi-isometric simple groups, with the first such example found by Caprace and Rémy [CR09, CR10], and the second being the Skipper–Witzel–Zaremsky family.

Finally, we should mention that Theorem D does not appear to be sharp. We conjecture the following precise characterization of the finiteness properties of SV_G, which is true for $n = 1$ as seen in Theorem A, and is reminiscent of the analogous characterization for the wreath products considered in [BdCK15].

Conjecture H. *The group SV_G is of type F_n if and only if the following conditions hold:*

  (i) *The action of G on S has finitely many orbits of n-element subsets.*

  (ii) *G is of type F_n.*

  (iii) *For each $1 \leq k < n$, the stabilizer in G of any k-element subset of S is of type F_{n-k}.*

For example, we conjecture that if G is finitely presented, the action of G on S has finitely many orbits of pairs, and the stabilizer in G of each point of S is finitely generated, then SV_G is finitely presented.

1. The groups

In this section we define the family of twisted Brin–Thompson groups. We begin by briefly reviewing the definition of Brin's groups sV, which we generalize slightly to a family of groups SV, where S is a set of any cardinality. In the case where S is a finite set with s elements, our group SV is isomorphic to Brin's group sV.

The groups SV. Let $\mathfrak{C}$ be the Cantor set $\{0,1\}^\omega$. Given a set S, the associated *Cantor cube* $\mathfrak{C}^S$ is the product space $\prod_{s \in S} \mathfrak{C}$, i.e., the set of all functions $S \to \mathfrak{C}$.

Let $\{0,1\}^*$ be the set of all finite binary sequences, including the empty sequence $\varnothing$. We say that a function $\psi \colon S \to \{0,1\}^*$ has *finite support* if $\psi(s) = \varnothing$ for all but finitely many $s \in S$. Given such a function, the associated *dyadic brick* in S is the set

$$B(\psi) = \{\kappa \in \mathfrak{C}^S \mid \psi(s) \text{ is a prefix of } \kappa(s) \text{ for each } s \in S\}.$$

Note that there is a *canonical homeomorphism* $h_\psi \colon \mathfrak{C}^S \to B(\psi)$ defined by

$$h_\psi(\kappa)(s) = \psi(s) \cdot \kappa(s)$$

for each $\kappa \in \mathfrak{C}^S$ and $s \in S$, where $\cdot$ denotes concatenation. More generally, if $B(\varphi)$ and $B(\psi)$ are dyadic bricks, we refer to the composition $h_\psi \circ h_\varphi^{-1}$ as the *canonical homeomorphism* from $B(\varphi)$ to $B(\psi)$.

Given any two partitions $B(\varphi_1), \ldots, B(\varphi_n)$ and $B(\psi_1), \ldots, B(\psi_n)$ of $\mathfrak{C}^S$ into the same number of dyadic bricks, we can define a homeomorphism $\mathfrak{C}^S \to \mathfrak{C}^S$ by mapping each $B(\varphi_i)$ to the corresponding $B(\psi_i)$ by a canonical homeomorphism. The group of all homeomorphisms of $\mathfrak{C}^S$ defined in this fashion is the *Brin–Thompson group SV*.

Remark 1.1. If $\mathcal{P}$ is a partition of $\mathfrak{C}^S$ into dyadic bricks, a *very elementary expansion* of $\mathcal{P}$ is a partition $\mathcal{P}'$ with the property that each dyadic brick in $\mathcal{P}$ is the union of at most two dyadic bricks in $\mathcal{P}'$. A *dyadic partition* of $\mathfrak{C}^S$ is any partition that can be obtained from the trivial one by a sequence of very elementary expansions.

The difference between arbitrary partitions into dyadic bricks and dyadic partitions has been a source of some confusion in the literature on Brin–Thompson groups. Every partition of $\mathfrak{C}^S$ into dyadic bricks is a dyadic partition as long as $|S| \leq 2$, but for $|S| \geq 3$ non-dyadic partitions exist, e.g., the partition

$$B(0,0,0), \ B(1,1,1), \ B(0,1,\varnothing), \ B(\varnothing,0,1), \ B(1,\varnothing,0)$$

when $|S| = 3$. However, if $\mathcal{P}$ is any partition of $\mathfrak{C}^S$ into dyadic bricks, there always exists a dyadic partition $\mathcal{P}'$ of $\mathfrak{C}^S$ which can be obtained from $\mathcal{P}$ by a sequence of very elementary expansions.

If g is an element of SV, then it is always possible to find domain and range partitions for g that are dyadic. Something stronger is true: if $\mathcal{P}_1$ and $\mathcal{P}_2$ are any domain and range partitions for g, then there exist domain and range dyadic partitions $\mathcal{P}'_1$ and $\mathcal{P}'_2$ for g so that each $\mathcal{P}'_i$ can be obtained from $\mathcal{P}_i$ by a sequence of very elementary expansions.

In particular, the confusion in the literature does not result in any actual problems. One can define SV either using partitions into dyadic bricks or using dyadic partitions, and obtain the same group.

The twisted group SV_G. Now let G be any group of permutations of the set S. For each $\gamma \in G$, let $\tau_\gamma \colon \mathfrak{C}^S \to \mathfrak{C}^S$ be the *twist homeomorphism* that permutes the coordinates of points in $\mathfrak{C}^S$ according to γ. That is, τ_γ is the homeomorphism of $\mathfrak{C}^S$ defined by

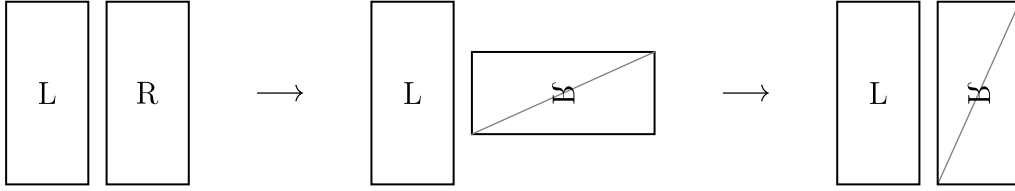$$\tau_\gamma(\kappa)(s) = \kappa(\gamma^{-1}s)$$

FIGURE 1. An example of an element of SV_G, where $S = \{1, 2\}$ and G is cyclic of order two.

for all $\kappa \in \mathfrak{C}^S$ and $s \in S$. Note that this definition is the natural one since $\tau_\gamma(\kappa)(\gamma s) = \kappa(s)$ for all $s \in S$, i.e., the γs coordinate of $\tau_\gamma(\kappa)$ is the same as the s coordinate of κ. Note also that $\tau_{\gamma\gamma'} = \tau_\gamma \tau_{\gamma'}$ for all $\gamma, \gamma' \in G$.

More generally, given any two dyadic bricks $B(\varphi)$ and $B(\psi)$ and an element $\gamma \in G$, the associated *twist homeomorphism* $B(\varphi) \to B(\psi)$ is the composition $h_\psi \circ \tau_\gamma \circ h_\varphi^{-1}$, where $h_\varphi \colon \mathfrak{C}^S \to B(\varphi)$ and $h_\psi \colon \mathfrak{C}^S \to B(\psi)$ are the canonical homeomorphisms.

Example 1.2. If S is the two-element set $\{1, 2\}$, then $\mathfrak{C}^S = \mathfrak{C} \times \mathfrak{C}$ is known as the *Cantor square*. In this case, the twist homeomorphism $\tau_{(1\,2)}$ associated with the transposition $(1\,2)$ is the "diagonal flip" of the Cantor square, which switches the first and second coordinates of each point. The transposition $(1\,2)$ also induces a twist homeomorphism between any two dyadic bricks in the Cantor square, which consists of a diagonal flip followed by translation and scaling.

Now, if $B(\varphi_1), \ldots, B(\varphi_n)$ and $B(\psi_1), \ldots, B(\psi_n)$ are any two partitions of $\mathfrak{C}^S$ into the same number of dyadic bricks and $\gamma_1, \ldots, \gamma_n \in G$, we can define a homeomorphism of $\mathfrak{C}^S$ by mapping each $B(\varphi_i)$ to $B(\psi_i)$ via the twist homeomorphism associated to γ_i. The group of all homeomorphisms of $\mathfrak{C}^S$ defined in this fashion is the *twisted Brin–Thompson group* SV_G. For example, Figure 1 shows an element of SV_G, where $S = \{1, 2\}$ and G is cyclic of order two.

2. THE GROUPOIDS

In this section we define certain groupoids associated to twisted Brin–Thompson groups. These groupoids provide a convenient conceptual and notational framework for working with elements of the groups.

The spaces $\mathfrak{C}^S(n)$. As before let $\mathfrak{C}$ be the Cantor space $\{0, 1\}^\omega$ and let $\mathfrak{C}^S$ be the space of functions $S \to \mathfrak{C}$. Let $\mathfrak{C}_1^S, \mathfrak{C}_2^S, \ldots$ be an infinite sequence of disjoint copies of $\mathfrak{C}^S$, and for $m \in \mathbb{N}$ let

$$\mathfrak{C}^S(m) = \mathfrak{C}_1^S \cup \cdots \cup \mathfrak{C}_m^S$$

be a union of m copies of $\mathfrak{C}^S$. A *dyadic brick* in $\mathfrak{C}^S(m)$ is any dyadic brick in any of the individual cubes $\mathfrak{C}_1^S, \ldots, \mathfrak{C}_m^S$. Note that it makes sense to talk about canonical homeomorphisms and twist homeomorphisms between dyadic bricks in different cubes.

In the future, we will identify the standard Cantor cube $\mathfrak{C}^S$ with the first cube $\mathfrak{C}_1^S$, or equivalently with the space $\mathfrak{C}^S(1)$.

The groupoids SV and SV_G. If $B_1, \ldots, B_k$ and $B_1', \ldots, B_k'$ are partitions of $\mathfrak{C}^S(m)$ and $\mathfrak{C}^S(n)$ respectively into the same number of dyadic bricks, we can define a homeomorphism $\mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ by mapping each B_i to B_i' by a canonical homeomorphism. The set of all such homeomorphisms is the *Brin–Thompson groupoid SV*. If we instead map each B_i to B_i' via the twist homeomorphism associated with some $\gamma_i \in G$, then the set of all such homeomorphisms is the *twisted Brin–Thompson groupoid SV_G*.

Note that the Brin–Thompson group SV is precisely the subgroup of SV consisting of all homeomorphisms $\mathfrak{C}^S(1) \to \mathfrak{C}^S(1)$. Similarly, the twisted Brin–Thompson group SV_G is the subgroup of SV_G consisting of all homeomorphisms $\mathfrak{C}^S(1) \to \mathfrak{C}^S(1)$.

Germinal twists. If $h \in SV_G$ and κ is a point in $\mathfrak{C}^S$, then there exists a dyadic brick B containing κ, a dyadic brick B' containing $h(\kappa)$, and an element $\gamma \in G$ so that h maps B to B' via the twist homeomorphism determined by γ. In this case, we say that γ is the *germinal twist* for h at κ, denoted $\mathrm{gtwist}_\kappa(h)$. Note that this does not depend on the bricks B and B' chosen, for if B'' is a sub-brick of B, then $h(B'')$ must be a sub-brick of B', and h maps B'' to $h(B'')$ by the twist homeomorphism associated to the same element $\gamma \in G$.

The notion of a germinal twist extends to the groupoid SV_G in an obvious way. Note that a homeomorphism $h \colon \mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ in SV_G lies in SV if and only if the germinal twist at each point in $\mathfrak{C}^S(m)$ is the identity element of G.

Note that the germinal twist is multiplicative, in the sense that

$$\mathrm{gtwist}_\kappa(h_2 h_1) = \mathrm{gtwist}_{h_1(\kappa)}(h_2) \, \mathrm{gtwist}_\kappa(h_1)$$

for any $h_1, h_2 \in SV_G$ and $\kappa \in \mathfrak{C}^S$. For $h_1, h_2 \in SV_G$, it follows that

 (i) $SV \, h_1 = SV \, h_2$ if and only if $\mathrm{gtwist}_\kappa(h_1) = \mathrm{gtwist}_\kappa(h_2)$ for all $\kappa \in \mathfrak{C}^S$, and
 (ii) $SV \, h_1 \, SV = SV \, h_2 \, SV$ if and only if

$$\{\mathrm{gtwist}_\kappa(h_1) \mid \kappa \in \mathfrak{C}^S\} = \{\mathrm{gtwist}_\kappa(h_2) \mid \kappa \in \mathfrak{C}^S\}.$$

Direct sum. In addition to their groupoid structure, SV and SV_G have an additional operation $\oplus$ that we refer to as the *direct sum*. Given any two elements $h \colon \mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ and $h' \colon \mathfrak{C}^S(m') \to \mathfrak{C}^S(n')$ of SV_G, their direct sum is the homeomorphism

$$h \oplus h' \colon \mathfrak{C}^S(m + m') \to \mathfrak{C}^S(n + n')$$

that maps the first m cubes of $\mathfrak{C}_{m+m'}^S$ to the first n cubes of $\mathfrak{C}_{n+n'}^S$ via h, and maps the m' remaining cubes of $\mathfrak{C}_{m+m'}^S$ to the n' remaining cubes of $\mathfrak{C}_{n+n'}^S$ via h'.

Permutations. Given a permutation σ of $\{1, \ldots, n\}$, the associated *permutation homeomorphism*

$$p_\sigma \colon \mathfrak{C}^S(n) \to \mathfrak{C}^S(n)$$

in SV is the homeomorphism that permutes the n Cantor cubes of $\mathfrak{C}^S(n)$ via σ. That is, p_σ is the homeomorphism of $\mathfrak{C}^S(n)$ mapping each Cantor cube $\mathfrak{C}_i^S$ to $\mathfrak{C}_{\sigma(i)}^S$ via the canonical homeomorphism.

G-twists and twisted permutations. Given $\gamma_1, \ldots, \gamma_n \in G$ with associated twist home-omorphisms $\tau_{\gamma_1}, \ldots, \tau_{\gamma_n}$ of $\mathfrak{C}^S$, the direct sum

$$t = \tau_{\gamma_1} \oplus \cdots \oplus \tau_{\gamma_n}$$

will be referred to as a *G-twist* of $\mathfrak{C}^S(n)$, or simply a *twist* if the group G is clear. The group of all G-twists of $\mathfrak{C}^S(n)$ is isomorphic to the direct product $G^n = G \times \cdots \times G$.

If t is a G-twist of $\mathfrak{C}^S(n)$ and p_σ is a permutation homeomorphism of $\mathfrak{C}^S(n)$, the composition tp_σ will be called a *twisted permutation* of $\mathfrak{C}^S(n)$. Let $\mathcal{G}(n)$ be the group of all twisted permutations of $\mathfrak{C}^S(n)$; this is isomorphic to the wreath product $G \wr \Sigma_n = G^n \rtimes \Sigma_n$, where Σ_n is the symmetric group of rank n.

Simple splits. Given an $s \in S$, the *simple split with color s* is the homeomorphism

$$x_s \colon \mathfrak{C}^S(1) \to \mathfrak{C}^S(2)$$

in SV defined as follows. Let $\{B_{s,0}, B_{s,1}\}$ be the partition of $\mathfrak{C}^S(1) = \mathfrak{C}^S$ into two bricks cut along the s coordinate, so $B_{s,i}$ is the dyadic brick consisting of all $\kappa \in \mathfrak{C}^S$ such that $\kappa(s) = i$. Then the simple split x_s is the homeomorphism that maps $B_{s,0}$ to the first cube $\mathfrak{C}_1^S$ of $\mathfrak{C}^S(2)$ and $B_{s,1}$ to the second cube $\mathfrak{C}_2^S$ of $\mathfrak{C}^S(2)$ by canonical homeomorphisms.

The following lemma records some basic relations in the groupoid SV_G, all of which are easy to check.

Lemma 2.1 (Relations in SV_G).

    (1) *If $\sigma, \sigma' \in \Sigma_n$ then $p_\sigma p_{\sigma'} = p_{\sigma\sigma'}$.*

    (2) *If $\sigma \in \Sigma_n$ and $\sigma' \in \Sigma_{n'}$, then $p_\sigma \oplus p_{\sigma'} = p_{\sigma''}$ for some $\sigma'' \in \Sigma_{n+n'}$.*

    (3) *If $\gamma, \gamma' \in G$ then $\tau_\gamma \tau_{\gamma'} = \tau_{\gamma\gamma'}$.*

    (4) *If $h, h', h'' \in SV_G$ then $(h \oplus h') \oplus h'' = h \oplus (h' \oplus h'')$.*

    (5) *If $h_1, h_1', h_2, h_2' \in SV_G$ and $h_1 h_1'$ and $h_2 h_2'$ are both defined, then $(h_1 \oplus h_2)(h_1' \oplus h_2') = (h_1 h_1') \oplus (h_2 h_2')$.*

    (6) *If $h_1, \ldots, h_n \in SV_G$ and $\sigma \in \Sigma_n$ then $(h_1 \oplus \cdots \oplus h_n)p_\sigma = p_\sigma(h_{\sigma(1)} \oplus \cdots \oplus h_{\sigma(n)})$.*

    (7) *If $s \in S$ and $\gamma \in G$ then $x_{\gamma s}\tau_\gamma = (\tau_\gamma \oplus \tau_\gamma)x_s$.*

    (8) *If $s, t \in S$ and $s \neq t$ then $(x_t \oplus x_t)x_s = p_{(2\ 3)}(x_s \oplus x_s)x_t$, where $(2\ 3) \in \Sigma_4$.*   □

Multicolored trees. A homeomorphism $f \colon \mathfrak{C}^S \to \mathfrak{C}^S(n)$ is called a *multicolored tree* if f maps the bricks of some dyadic partition $B_1, \ldots, B_n$ of $\mathfrak{C}^S$ (see Remark 1.1) to the Cantor cubes $\mathfrak{C}_1^S, \ldots, \mathfrak{C}_n^S$ of $\mathfrak{C}^S(n)$ by canonical homeomorphisms.

Equivalently, f is a multicolored tree if either f is the identity on $\mathfrak{C}^S$ or

$$f = p_\sigma(f_0 \oplus f_1)x_s$$

for some permutation homeomorphism p_σ of $\mathfrak{C}^S(n)$, some multicolored trees f_0, f_1, and some simple split x_s. Note that a given multicolored tree might have more than one description of this form, as in Lemma 2.1(8).

Multicolored forests. A *multicolored forest* is any homeomorphism $\mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ of the form

$$p_\sigma(f_1 \oplus \cdots \oplus f_m)$$

where p_σ is a permutation homeomorphism of $\mathfrak{C}^S(n)$ and $f_1, \ldots, f_m$ are multicolored trees. Note that the set of all multicolored forests is closed under composition and direct sums. Indeed, the set of multicolored forests is precisely the smallest set of elements of $S\mathcal{V}_G$ which contains all permutations and simple splits and is closed under composition and direct sum.

We can also describe multicolored forests in terms of dyadic partitions. We say that a partition $B_1, \ldots, B_n$ of $\mathfrak{C}^S(m)$ into dyadic bricks is a *dyadic partition* if its restriction to each $\mathfrak{C}_i^S$ is a dyadic partition of $\mathfrak{C}_i^S$. Then $f \colon \mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ is a multicolored forest if and only if it maps the bricks of some dyadic partition $B_1, \ldots, B_n$ of $\mathfrak{C}^S(m)$ to the Cantor cubes $\mathfrak{C}_1^S, \ldots, \mathfrak{C}_n^S$ of $\mathfrak{C}^S(n)$ by canonical homeomorphisms. This gives a one-to-one correspondence between multicolored forests $\mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ and ordered dyadic partitions $(B_1, \ldots, B_n)$ of $\mathfrak{C}^S(m)$.

The following lemma describes the structure of arbitrary elements of $S\mathcal{V}_G$.

Lemma 2.2. *Every element of $S\mathcal{V}_G$ can be written as $f_2^{-1} t f_1$ for some multicolored forests f_1, f_2 and some G-twist t. Such an element lies in $S\mathcal{V}$ if and only if t is the identity.*

Proof. Given a homeomorphism $h \colon \mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ in $S\mathcal{V}_G$, let $B_1, \ldots, B_k$ be a domain partition and $B_1', \ldots, B_k'$ a range partition so that h maps each B_i to B_i' via the twist homeomorphism associated to some $\gamma_i \in G$. Refining if necessary as in Remark 1.1, we may assume that $B_1, \ldots, B_k$ and $B_1', \ldots, B_k'$ are dyadic partitions. Then

$$h = f_2^{-1}(\tau_{\gamma_1} \oplus \cdots \oplus \tau_{\gamma_k})f_1$$

where $f_1 \colon \mathfrak{C}^S(m) \to \mathfrak{C}^S(k)$ is the multicolored forest that maps each B_i to $\mathfrak{C}_i^S$ by a canonical homeomorphism, and $f_2 \colon \mathfrak{C}^S(n) \to \mathfrak{C}^S(k)$ is the multicolored forest that maps each B_i' to $\mathfrak{C}_i^S$ by a canonical homeomorphism. Moreover, since h has germinal twist γ_i on each B_i, it lies in $S\mathcal{V}$ if and only if $\gamma_1, \ldots, \gamma_k$ are all the identity, which occurs if and only if $\tau_{\gamma_1} \oplus \cdots \oplus \tau_{\gamma_k}$ is the identity homeomorphism of $\mathfrak{C}^S(k)$. $\qquad\square$

It follows from this lemma that every element of the group $S\mathcal{V}_G$ can be written as $f_2^{-1} t f_1$ for some multicolored trees f_1, f_2 and some G-twist t.

It will be very important in what follows to understand how multicolored forests interact with twisted permutations. The following lemma establishes the relationship.

Lemma 2.3. *If f is a multicolored forest and g is a twisted permutation such that the composition fg is defined, then $fg = g'f'$ for some twisted permutation g' and some multicolored forest f'. Moreover, if f is a multicolored tree then g' is a direct sum of copies of g.*

Proof. By the relations in parts (6) and (7) of Lemma 2.1, this holds whenever f is a permutation homeomorphism or a simple split. But if the statement holds for multicolored forests f_1 and f_2, then it clearly holds for $f_1 f_2$ if this is defined, and moreover it holds for $f_1 \oplus f_2$ by relation (5) in Lemma 2.1. Since every multicolored forest can be obtained from
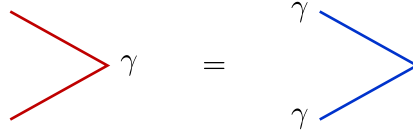
FIGURE 2. With $S = \{r, b\}$ and γ satisfying $r = \gamma b$, a picture of the equality $x_r \tau_\gamma = (\tau_\gamma \oplus \tau_\gamma) x_b$.

permutation homeomorphisms and simple splits by taking iterated compositions and direct sums, the result follows. $\qquad\square$

See Figure 2 for an example of Lemma 2.3.

Remark 2.4. For readers familiar with [WZ18, Wit19], the formula $fg = g'f'$ provides a Zappa–Szép structure, which could be developed into a notion of "multicolored cloning systems" using multicolored forests. A number of aspects of the general framework break down when there are infinitely many colors though, so it is unlikely this generalization would be helpful for our current setup. In the future, developing the frameworks in [WZ18, Wit19] to allow for multicolored forests could be interesting.

Subsets of S and spectrum. Each subset T of S corresponds to a subgroup $TV^{(S)}$ of SV which is isomorphic to TV. Specifically, if we decompose $\mathfrak{C}^S$ as a product $\mathfrak{C}^T \times \mathfrak{C}^{S \setminus T}$, then $TV^{(S)}$ is the subgroup of elements of SV that act as an element of TV on the $\mathfrak{C}^T$ factor and act as the identity on the $\mathfrak{C}^{S \setminus T}$ factor. Similarly, the groupoid $S\mathcal{V}$ has a natural subgroupoid $T\mathcal{V}^{(S)}$ which is isomorphic to $T\mathcal{V}$.

If $h \in S\mathcal{V}$, the *spectrum* of h is the finite set $\mathrm{Spec}(h) \subseteq S$ of all coordinates that h uses in an essential way. That is, $\mathrm{Spec}(h)$ is the smallest set $T \subseteq S$ such that $h \in T\mathcal{V}^{(S)}$. This definition has the following properties:

(i) $\mathrm{Spec}(h)$ is empty if and only if h is a permutation homeomorphism.
(ii) $\mathrm{Spec}(x_s) = \{s\}$ for any simple split x_s.
(iii) $\mathrm{Spec}(h_1 \oplus h_2) = \mathrm{Spec}(h_1) \cup \mathrm{Spec}(h_2)$ for any $h_1, h_2 \in S\mathcal{V}$.
(iv) $\mathrm{Spec}(h_1 h_2) \subseteq \mathrm{Spec}(h_1) \cup \mathrm{Spec}(h_2)$ for any $h_1, h_2 \in S\mathcal{V}$ such that $h_1 h_2$ is defined, with equality if h_1 and h_2 are multicolored forests.

Note that these properties can be used to inductively compute the spectrum of any multicolored forest. For example,

$$\mathrm{Spec}\big(p_{(2\ 3\ 5)}((x_s \oplus \mathrm{id}_1) x_t \oplus x_u) x_v\big) = \{s, t, u, v\}$$

for any $s, t, u, v \in S$, where id_1 denote the identity map on $\mathfrak{C}^S$.

3. Generators and simplicity

We will need two embeddings of G into SV_G. First, let $\iota_\varnothing \colon G \to SV_G$ be the embedding

$$\iota_\varnothing(\gamma) = \tau_\gamma.$$

Note that this is a homomorphism by part (3) of Lemma 2.1. Next, fix an $s \in S$, and define an embedding $\iota_1^s \colon G \to SV_G$ by

$$\iota_1^s(\gamma) = x_s^{-1}(\mathrm{id}_1 \oplus \tau_\gamma)x_s$$

where id_1 denotes the identity map on $\mathfrak{C}^S$. Thus $\iota_1^s(\gamma)$ applies the twist τ_γ to the "right half" of $\mathfrak{C}^S$ determined by the direction s.

Proposition 3.1. SV_G *is generated by* $SV \cup \iota_1^s(G)$.

Proof. Let $H = \langle SV \cup \iota_1^s(G) \rangle$. Since the double coset $SV\,g\,SV$ of an element $g \in SV_G$ is determined by its set of germinal twists, it suffices to prove that there exists an element of H with any given set $\{\gamma_1, \ldots, \gamma_n\}$ of germinal twists.

Note first that elements of SV have germimal twist set $\{1\}$, and elements of $\iota_1^s(G)$ have germinal twist set $\{1, \gamma\}$ for some $\gamma \in G$, so H contains all elements having these sets of germinal twists. Next note that if $g \in H$ has germinal twist set $\{\gamma_1, \ldots, \gamma_n\}$ and $\gamma \in G$, then there exist $g_1, g_2 \in H$ with germinal twist set $\{1, \gamma\}$ so that $g_1 g$ has germinal twist set $\{\gamma\gamma_1\} \cup \{\gamma_2, \ldots, \gamma_n\}$ and $g_2 g$ has germinal twist set $\{\gamma\gamma_1\} \cup \{\gamma_1, \ldots, \gamma_n\}$. It follows easily by induction that H contains elements with any given set of germinal twists. $\square$

To obtain finite generation conditions for SV_G, we need some information about generating sets for SV. For $T \subseteq S$, recall that SV has a natural subgroup $TV^{(S)}$ that is isomorphic to TV. For the following proposition, we are particularly interested in the case where T is a two-element subset of S, so $TV^{(S)} \cong 2V$.

Proposition 3.2. *Let S be a set with $|S| \geq 2$, let Γ be a connected simple graph with $V(\Gamma) = S$, and for each edge $e \in E(\Gamma)$ let $eV^{(S)}$ be the associated copy of $2V$. Then the union $\bigcup_{e \in E(\Gamma)} eV^{(S)}$ generates SV.*

Proof. In [HM12, Theorem 25], Hennig and Matucci describe a set of generators for the group $nV = \{1, \ldots, n\}V$, each of which is contained in $\{1, d\}V^{(\{1, \ldots, n\})}$ for some $2 \leq d \leq n$. Thus the given proposition holds when $S = \{1, \ldots, n\}$ and $E(\Gamma) = \big\{\{1, 2\}, \{1, 3\}, \ldots, \{1, n\}\big\}$, i.e., whenever the graph Γ is a finite star.

For S finite, we can now prove the proposition by induction on $|S|$. The case where $|S| = 2$ is done, so suppose that $|S| \geq 3$. Let u be a vertex of Γ whose removal does not disconnect Γ, and let Γ' be the (connected) subgraph of Γ induced by $S \setminus u$. Let v be a vertex of Γ that is adjacent to u. By our induction hypothesis, the copies of $2V$ corresponding to the edges of Γ' generate $(S \setminus u)V^{(S)}$, so we can replace Γ' by the star $\big\{\{v, w\}\big\}_{w \in S \setminus \{u, v\}}$. Since $\{u, v\} \in E(\Gamma)$, the graph Γ now contains the entire star $\big\{\{v, w\}\big\}_{w \in S \setminus v}$, so by induction the proposition holds when S is finite.

The case where S is infinite follows easily, since any infinite connected graph is the union of its finite connected subgraphs. $\square$

Corollary 3.3. *Let Γ be a connected graph with vertex set S that is invariant under the action of G, and let $\{e_\alpha\}_{\alpha \in \mathcal{I}}$ be a system of representatives for the orbits of $E(\Gamma)$ under G. Then SV_G is generated by the subgroups $e_\alpha V^{(S)}$ (each of which is isomorphic to $2V$) together with $\iota_\varnothing(G)$ and $\iota_1^s(G)$.*

Proof. For $\gamma \in G$, conjugation by $\iota_\varnothing(\gamma)$ maps $e_\alpha V^{(S)}$ to $(\gamma e_\alpha)V^{(S)}$. So the subgroup generated by the $\iota_\varnothing(G) \cup \bigcup_{\alpha \in \mathcal{I}} e_\alpha V^{(S)}$ contains $eV^{(S)}$ for every $e \in E(\Gamma)$, and thus contains SV by Proposition 3.2. By Proposition 3.1, it follows that the given generators generate SV_G. $\qquad \square$

Theorem A. *The group SV_G is finitely generated if and only if G is finitely generated and the action of G on S has finitely many orbits.*

Proof. First suppose G is finitely generated and the action of G on S has finitely many orbits. Let $s_1, \ldots, s_m$ be representatives for the orbits of G in S, let $\gamma_1, \ldots, \gamma_n$ be generators for G, and let Γ be the graph with vertex set S and edge set

$$\big\{\{\gamma s_1, \gamma s_i\} \mid 2 \leq i \leq m, \gamma \in G\big\} \cup \big\{\{\gamma s_1, \gamma \gamma_j s_1\} \mid 1 \leq j \leq n, \ \gamma_j s_1 \neq s_1, \text{ and } \gamma \in G\big\}.$$

Then Γ is connected and G-invariant, and has finitely many orbits of edges under G. By Corollary 3.3, it follows that SV_G is generated by the subgroups $\{s_1, s_i\}V^{(S)}$ $(2 \leq i \leq m)$ and $\{s_1, \gamma_j s_1\}V^{(S)}$ $(1 \leq j \leq n$ and $\gamma_j s_1 \neq s_1)$ together with $\iota_\varnothing(G)$ and $\iota_1^s(G)$, all of which have finitely many generators.

Now suppose G is not finitely generated, so there is a filtration $G_1 < G_2 < G_3 < \cdots$ of G by proper subgroups. Then $SV_{G_1} \leq SV_{G_2} \leq SV_{G_3} \leq \cdots$ is a chain of subgroups whose union is SV_G. Moreover, for each i the set of all germinal twists of all elements of SV_{G_i} equals G_i, so each SV_{G_i} is a proper subgroup of $SV_{G_{i+1}}$. We conclude that SV_G is not finitely generated.

Lastly, suppose the action of G on S has infinitely many orbits. Then there is a filtration $S_1 \subset S_2 \subset S_3 \subset \cdots$ of S by proper G-invariant subsets. Fix an $s \in S_1$, and for each i let H_i be the subgroup of SV_G generated by $S_iV^{(S)}$ and $\iota_1^s(G)$. (This subgroup is isomorphic to S_iV_G if G acts faithfully on S_i, but in general S_iV_G is a quotient of H_i.) Since each S_i is G-invariant, we know that $H_i \cap SV = S_iV^{(S)}$, and therefore each H_i is a proper subgroup of H_{i+1}. By Proposition 3.1 the union of the H_i's is all of SV_G, and therefore SV_G is not finitely generated. $\qquad \square$

In the case where G is Thompson's group F acting on the dyadics S in $(0, 1)$, as in Corollary E, the proof of the theorem above tells us that SV_G is generated by $\{1/2, 3/4\}V^{(S)}$ (a copy of $2V$) together with $\iota_\varnothing(F)$ and $\iota_1^{1/2}(F)$ (two copies of F). If instead G is Houghton's group H_n and $S = \{1, \ldots, n\} \times \mathbb{N}$ as in Corollary G, a similar argument shows that SV_G is generated by one copy of $2V$ and two copies of H_n. (In general, SV_G is generated by one copy of $2V$ and two copies of G whenever the action of G on S is primitive, since in this case G has a generating set with only one generator that moves s_1.)

Theorem 3.4. *The twisted Brin–Thompson group SV_G is simple. Moreover, if SV_G is finitely generated then it can be generated by two elements of finite order.*

Proof. In [Hyd17, BEH20], Bleak, Elliott, and Hyde prove that a group K of homeomorphisms of a Cantor space X is simple provided that it satisfies the following three conditions:

(i) K is *vigorous*. That is, for every clopen set $A \subseteq X$ and proper clopen subsets B and C of A there exists a $k \in K$ supported on A such that $k(B) \subseteq C$.

(ii) K is generated by elements of *small support*, i.e., those elements which are supported on proper clopen subsets of X.

(iii) K is *perfect*, i.e., $[K, K] = K$.

Moreover, they prove that every simple, vigorous, finitely generated group of homeomorphisms of the Cantor set is generated by two elements of finite order. Thus it suffices to verify that SV_G satisfies the three conditions above.

Clearly SV is vigorous, so SV_G is as well. Moreover, since SV is generated by elements of small support and each $\iota_1^s(\gamma)$ has small support, it follows from Proposition 3.1 that SV_G is generated by elements of small support. To prove that SV_G is perfect, recall first from [Bri10] that SV is perfect, and therefore $[SV_G, SV_G]$ contains SV. By Proposition 3.1, it suffices to prove that $[SV_G, SV_G]$ contains $\iota_1^s(\gamma)$ for some fixed $s \in S$ and all $\gamma \in G$. Let $B = x_s^{-1}(\mathfrak{C}_2^S)$ be the support of $\iota_1^s(\gamma)$, and let f be any element of SV that maps B to a proper subset of itself. Then the commutator $k = [\iota_1^s(\gamma), f] = \iota_1^s(\gamma) f \iota_1^s(\gamma)^{-1} f^{-1}$ has germinal twist γ on $B \setminus f(B)$ and the identity everywhere else. Since $\iota_1^s(\gamma)$ also has germinal twists γ and the identity, there exist elements $h_1, h_2 \in SV$ so that $h_1 k h_2 = \iota_1^s(\gamma)$, and therefore $\iota_s^1(\gamma) \in [SV_G, SV_G]$. $\square$

4. Quasi-retracts

If X and Y are metric spaces, a function $f \colon X \to Y$ is said to be *coarse Lipschitz* if there exist constants $C, D > 0$ so that

$$d(f(x), f(x')) \leq C \, d(x, x') + D$$

for all $x, x' \in X$. For example, any homomorphism between finitely generated groups is coarse Lipschitz with respect to the word metrics.

A function $\iota \colon X \to Y$ is said to be a *quasi-isometric embedding* if there exist constants $C, D > 0$ so that

$$\frac{1}{C} \, d(x, x') - D \leq d(f(x), f(x')) \leq C \, d(x, x') + D$$

for all $x, x' \in X$. Clearly any quasi-isometric embedding is coarse Lipschitz, but the converse need not hold.

A function $\rho \colon X \to Y$ is said to be a *quasi-retraction* if it is coarse Lipschitz and there exists a coarse Lipschitz function $\iota \colon Y \to X$ and a constant $E > 0$ so that $d(\rho \circ \iota(y), y) \leq E$ for all $y \in Y$. If such a function ρ exists, then Y is said to be a *quasi-retract* of X. In this case, it is not hard to show that ι must be a quasi-isometric embedding of Y into X.

Now suppose G and SV_G are finitely generated, with word metrics corresponding to some finite generating sets. Our next result is that G is a quasi-retract of SV_G. In order to prove this we need a way of canonically extracting an element of G from an element of SV_G. We will do this using the notion of germinal twists $\mathrm{gtwist}_\kappa(h)$ from Section 2.

Theorem B. *Suppose G and SV_G are finitely generated. Then there exists a coarse Lipschitz map $\rho\colon SV_G \to G$ such that $\rho \circ \iota_\varnothing$ is the identity. In particular, the group SV_G quasi-retracts onto G, and $\iota_\varnothing$ is a quasi-isometric embedding of G into SV_G.*

Proof. Fix a point $\kappa \in \mathfrak{C}^S$, and let $\rho\colon SV_G \to G$ be the function

$$\rho(h) = \mathrm{gtwist}_\kappa(h).$$

Note then that $\rho \circ \iota_\varnothing$ is the identity function on G. We claim that ρ is coarse Lipschitz with respect to the *left* word metrics on SV_G and G.

Fix an $s \in S$ and a generating set A_G for G, and recall from Corollary 3.3 that SV_G is generated by $\iota_\varnothing(A_G) \cup \iota_1^s(A_G)$ together with finitely many elements of SV. But:

(i) $\rho\big(\iota_\varnothing(a)h\big) = a\rho(h)$ for all $a \in A_G$ and $h \in SV_G$,

(ii) $\rho\big(\iota_1^s(a)h\big) \in \{\rho(h), a\rho(h)\}$ for all $a \in A_G$ and $h \in SV_G$, and

(iii) $\rho(h'h) = \rho(h)$ for all $h' \in SV$ and $h \in SV_G$.

It follows that ρ is nonexpanding and hence coarse Lipschitz. Since $\iota_\varnothing$ is a homomorphism, it must be coarse Lipschitz as well. We conclude that ρ is a quasi-retraction and $\iota_\varnothing$ is a quasi-isometric embedding. (In fact for these choices of generating sets, $\iota_\varnothing$ is even an isometric embedding.) $\qquad\square$

5. COMPLEXES

In this section we construct complexes on which the twisted Brin–Thompson groups SV_G act. First we define a directed poset P_1 on which SV_G acts and then build the so called Stein complex X. Most of the ideas here are directly inspired by the construction of the analogous complexes for sV in [FMWZ13]. However, unlike the complex in [FMWZ13], the Stein complex here will be locally infinite if S is infinite, which makes a number of arguments more complicated.

The poset P. If $h\colon \mathfrak{C}^S(m') \to \mathfrak{C}^S(m)$ is an element of SV_G write $[h]$ for the coset $\mathcal{G}(m)h$, where $\mathcal{G}(m)$ is the group of twisted permutations of $\mathfrak{C}^S(m)$. Let P be the set of equivalence classes $[h]$ of elements h of SV_G.

If $h\colon \mathfrak{C}^S(m') \to \mathfrak{C}^S(m)$ is an element of SV_G, we will refer to m as the *rank* of h and m' as the *corank* of h. Note that rank and corank are invariant under multiplication by twisted permutations, so it makes sense to refer to the *rank* and *corank* of an element of P.

We can place a partial order on P as follows. Given $v, w \in P$, we say that w is an *expansion* of v, denoted $v \le w$, if there exists an $h \in SV_G$ and a multicolored forest f so that $v = [h]$ and $w = [fh]$.

Lemma 5.1. *The relation $\le$ is a partial order on P.*

Proof. Clearly $\le$ is reflexive. For antisymmetry, suppose that $v \le w$ and $w \le v$ for some $v, w \in P$. Then there exist $h_1, h_2 \in SV_G$ and multicolored forests f_1, f_2 so that $v = [h_1] = [f_2 h_2]$ and $w = [h_2] = [f_1 h_1]$. Then $h_1 = g_2 f_2 h_2$ and $h_2 = g_1 f_1 h_1$ for some twisted permutations g_1 and g_2. It follows that $g_1 f_1 g_2 f_2$ is the identity, so f_1 and f_2 must be permutations. Then $g_1 f_1$ is a twisted permutation, so $[h_1] = [h_2]$ and hence $v = w$.

For transitivity suppose that $u \leq v$ and $v \leq w$ for some $u, v, w \in P$. Then there exist $h_1, h_2 \in SV_G$ and multicolored forests f_1, f_2 so that $u = [h_1]$, $v = [h_2] = [f_1 h_1]$, and $w = [f_2 h_2]$. Let g be a twisted permutation so that $h_2 = g f_1 h_1$, and observe that $f_2 g = g' f_2'$ for some twisted permutation g' and multicolored forest f_2' by Lemma 2.3. Then $w = [f_2 h_2] = [f_2 g f_1 h_1] = [g' f_2' f_1 h_1] = [f_2' f_1 h_1]$, and since $f_2' f_1$ is a multicolored forest it follows that $u \leq w$. $\square$

Let P_1 be the subposet of P consisting of corank-1 elements. We are especially interested in the restriction of $\leq$ to P_1. Recall that a poset is *directed* if any two elements have a common upper bound. It is a standard fact that directed posets have contractible geometric realizations.

Proposition 5.2. *The poset P_1 is directed, so its geometric realization $|P_1|$ is contractible.*

Proof. Let $[h_1], [h_2] \in P_1$. By Lemma 2.2, there exist multicolored forests f_1, f_2 and a G-twist t so that $h_1 h_2^{-1} = f_2^{-1} t f_1$, and rearranging gives $f_2 h_1 = t f_1 h_2$. Then $[h_1] \leq [f_2 h_1]$ and $[h_2] \leq [f_1 h_2] = [t f_1 h_2] = [f_2 h_1]$, so $[f_2 h_1]$ is a common upper bound. $\square$

Elementary bricks and elementary multicolored forests. Our next goal is to define the Stein complex for SV_G, which is a certain contractible subcomplex of $|P_1|$. This will be based on the notion of elementary bricks and elementary multicolored forests.

We say that a dyadic brick B in $\mathfrak{C}^S$ is *elementary* if its defining function $\psi \colon S \to \{0,1\}^*$ has the property that $\psi(s) \in \{0, 1, \varnothing\}$ for all $s \in S$. Such a brick can be obtained from $\mathfrak{C}^S$ by cutting at most once in each direction.

A dyadic partition $\{B_1, \ldots, B_n\}$ of $\mathfrak{C}^S$ is called *elementary* if each of the bricks B_i is an elementary brick. Similarly, a dyadic partition of $\mathfrak{C}^S(n)$ is elementary if it restricts to an elementary partition of each of the Cantor cubes $\mathfrak{C}_1^S, \ldots, \mathfrak{C}_n^S$.

A multicolored forest $f \colon \mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ is called *elementary* if there exists an elementary dyadic partition $B_1, \ldots, B_n$ of $\mathfrak{C}^S(m)$ such that f maps each B_i to $\mathfrak{C}_i^S$ by a canonical homeomorphism.

The lattice of forests. Let F_m be the subset of P consisting of all vertices of corank m that can be represented by multicolored forests. A vertex of F_m is called *elementary* if it can be represented by an elementary multicolored forest.

Proposition 5.3. *Every pair of elements $[f_1], [f_2] \in F_m$ have a least upper bound $[f_1] \vee [f_2]$ in F_m. If $[f_1]$ and $[f_2]$ are elementary then so is $[f_1] \vee [f_2]$.*

Proof. Recall that the multicolored forests $f \colon \mathfrak{C}^S(m) \to \mathfrak{C}^S(n)$ are in one-to-one correspondence with ordered dyadic partitions $(B_1, \ldots, B_n)$ of $\mathfrak{C}^S(m)$, where f maps each B_i to $\mathfrak{C}_i^S$ by a canonical homeomorphism. If f and f' are multicolored forests corresponding to dyadic partitions $(B_1, \ldots, B_r)$ and $(B_1', \ldots, B_s')$ of $\mathfrak{C}^S(m)$, then it is easy to check that:

   (i) $[f] = [f']$ if and only if $\{B_1, \ldots, B_r\} = \{B_1', \ldots, B_s'\}$, and

   (ii) $[f] \leq [f']$ if and only if $\{B_1', \ldots, B_s'\}$ is a refinement of $\{B_1, \ldots, B_r\}$.
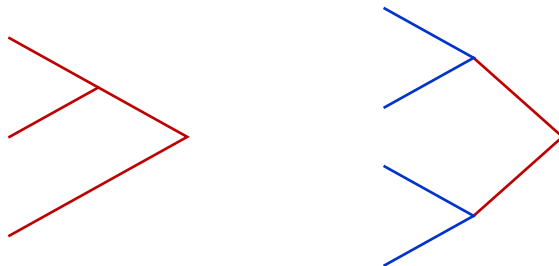
FIGURE 3. Let $S = \{r, b\}$. The first picture is the multicolored forest $f_1 = (x_r \oplus \mathrm{id}_1)x_r$, which is not elementary. The second picture is the multicolored forest $f_2 = (x_b \oplus x_b)x_r$, which is elementary.

That is, the vertices in F_m are in one-to-one correspondence with (unordered) dyadic partitions of $\mathfrak{C}^S(m)$, with the partial order $\leq$ corresponding to refinement. But any two dyadic partitions $\{B_1, \ldots, B_r\}$ and $\{B'_1, \ldots, B'_s\}$ of $\mathfrak{C}^S(m)$ have a least common refinement, namely the dyadic partition consisting of all nonempty intersections $B_i \cap B'_j$. Moreover, since the nonempty intersection of elementary bricks is an elementary brick, the least common refinement of two elementary dyadic partitions is again elementary. □

Recall that a partially ordered set is called a *lattice* if every pair of elements v, w have a least upper bound $v \vee w$, called the *join*, and a greatest lower bound $v \wedge w$, called the *meet*.

Corollary 5.4. *Each poset F_m is a lattice, as is the subset of F_m consisting of elementary vertices.*

Proof. Note that F_m has a unique minimal element, namely the equivalence class of the trivial corank-m forest. Moreover, every element of F_m has finitely many predecessors with respect to $\leq$. Since finite subsets of F_m have least upper bounds by Proposition 5.3, it follows easily that any finite subset of F_m has a greatest lower bound as well. A similar argument applies to the elementary vertices. □

The Stein complex. Given vertices $v, w \in P$, we say that w is an *elementary expansion* of v, denoted $v \preceq w$, if there exists an $h \in SV_G$ and an elementary multicolored forest f so that $v = [h]$ and $w = [fh]$. See Figure 3 for an example.

Now, recall that a simplex of the geometric realization $|P_1|$ is a finite chain $v_0 < \cdots < v_k$ of elements of P_1. Such a simplex is *elementary* if $v_0 \preceq v_k$ (and hence $v_i \preceq v_j$ for all $i \leq j$). The elementary simplices form a subcomplex X of $|P_1|$, which we refer to as the *Stein complex* for SV_G. The term "Stein complex" is a reference to the analogous complex constructed by Stein for Thompson's group F in [Ste92].

We will use interval notation in P, for instance $[v, w] = \{u \mid v \leq u \leq w\}$, with similar definitions for open and half-open intervals. Call an interval $[v, w]$ *elementary* if $v \preceq w$.

Proposition 5.5 (Elementary core). *For any $v < w$ in P the set of elementary expansions of v that lie in $[v, w]$ has a maximum element, which is not equal to v.*

Proof. Without loss of generality $v = [\text{id}_m]$ and $w = [f]$ for some multicolored forest f of corank m. Then the elementary expansions of v are precisely the elementary multicolored forests of corank m, of which only finitely many lie in $[v, w]$. By Proposition 5.3, these elementary multicolored forests have a least upper bound u. This is again an elementary multicolored forest and it lies in $[v, w]$, so u is the desired maximum.

Moreover, since $v \neq w$ the multicolored forest f is not a permutation, so we can write

$$f = p_\sigma(f_1 \oplus f_2)x_s$$

for some permutation homeomorphism p_σ, some multicolored forests f_1 and f_2, and some simple split x_s. Then $[x_s]$ is an elementary expansion of v that lies in $[v, w]$ and is strictly greater than v, so the maximum elementary expansion of v in $[v, w]$ is not equal to v. □

Given an interval $[v, w]$ in P, we refer to the maximum elementary expansion of v that it contains as the *elementary core* of w relative to v, denoted $\text{core}_v(w)$.

Proposition 5.6. *The Stein complex X is homotopy equivalent to $|P_1|$, hence contractible.*

Proof. We proceed by a now-standard strategy (see for example [Bro92, FMWZ13, BFM+16, BM16]). Observe that each interval $[v, w]$ in P_1 is itself a directed set, whose geometric realization $|[v, w]|$ is a contractible subcomplex of $|P_1|$. Since every finite chain in P_1 is contained in an interval, every simplex in $|P_1|$ is contained in some $|[v, w]|$, so $|P_1|$ is the union of these subcomplexes. Note that the Stein complex $X \subseteq |P_1|$ is precisely the union of the subcomplexes $|[v, w]|$ for which $[v, w]$ is an elementary interval.

Now, define the *length* of an interval $[v, w]$ to be $m' - m$, where m is the rank of v and m' is the rank of w. Then we can write $|P_1|$ as an ascending union of subcomplexes

$$X = X_0 \subseteq X_1 \subseteq X_2 \subseteq \cdots$$

where each X_n is the union of X_{n-1} with the geometric realizations of all non-elementary intervals $[v, w]$ of length n. It suffices to prove that each inclusion $X_{n-1} \to X_n$ is a homotopy equivalence.

Note first that each simplex in X_n which is not in X_{n-1} lies in the geometric realization of a unique non-elementary interval $[v, w]$ of length n. Thus it suffices to prove that the intersection $|[v, w]| \cap X_{n-1}$ is contractible for each non-elementary interval $[v, w]$ of length n. This intersection is precisely the union $|[v, w)| \cup |(v, w]|$, which is isomorphic to the suspension of $|(v, w)|$ with v and w serving as suspension points, so it suffices to prove that each $|(v, w)|$ is contractible.

Let $r \colon |(v, w)| \to |(v, \text{core}_v(w)]|$ be the retraction induced by the order-preserving function $\text{core}_v \colon (v, w) \to (v, \text{core}_v(w)]$. Quillen observed in [Qui78, Section 1.3] that if $f, g \colon A \to B$ are order-preserving functions between posets and $f(a) \leq g(a)$ for all $a \in A$ then the induced maps from $|A|$ to $|B|$ are homotopic. Since $\text{core}_v(x) \leq x$ for all $x \in (v, w)$, it follows that r is homotopic to the identity map on $|(v, w)|$. But $|(v, \text{core}_v(w)]|$ is clearly contractible, being the cone of $|(v, \text{core}_v(w))|$ with cone point $\text{core}_v(w)$, and therefore $|(v, w)|$ is contractible as well. □

6. Actions on the complexes

In this section we prove that if G is oligomorphic and if every stabilizer in G of a finite subset of S is of type F_n (including $G = \text{Stab}_G(\emptyset)$ itself), then the only remaining step toward proving that the twisted Brin–Thompson group SV_G is of type F_n (Theorem D) is a descending link analysis. This analysis will then happen in the next section. It seems worth remarking that these conditions on orbits and stabilizers of finite subsets are very similar to those used in [BdCK15] to characterize finiteness properties of wreath products.

The strategy is similar in spirit to that used in [FMWZ13] for the Brin–Thompson groups sV. For those familiar with the approach, the problem here is that elementary intervals can be arbitrarily long if we have an infinite set of colors S. This makes it difficult to apply any existing techniques that have been used for related groups. For instance one could try and use the "cloning system" tools from [WZ18] and [Wit19], but these are quickly stymied by the fact that elementary intervals can be arbitrarily long, so the analog of the crucial Observation 4.6 in [WZ18] fails (but see Remark 2.4 for further thoughts on the connection to cloning systems). For the same reason, the "groundedness" of Belk and Forrest (see [BM16, Theorem 6.2] and [BF19, Theorem 4.9]) does not apply directly.

We will use a combination of Bestvina–Brady Morse theory [BB97] and Brown's Criterion [Bro87] that by now is quite standard, which we set up now.

Definition 6.1 (Morse function, sublevel set, descending link)**.** Let Y be an affine cell complex and $\phi\colon Y \to \mathbb{R}$ a map. We call ϕ a *Morse function* if it is affine and non-constant on positive dimensional cells and the image $\phi(Y^{(0)}) \subseteq \mathbb{R}$ is closed and discrete. For $m \in \mathbb{R}$, the *sublevel complex* Y_m is the induced subcomplex of Y with vertex set $Y^{(0)} \cap \phi^{-1}((-\infty, m])$. For a vertex $y \in Y$ the *descending link* $\text{lk}^{\downarrow}y$ of y with respect to ϕ is the link of y in $Y_{\phi(y)}$.

Theorem 6.2. *Let H be a group acting cellularly on a contractible affine cell complex Y and let $\phi\colon Y \to \mathbb{R}$ be an H-invariant Morse function. Suppose that each sublevel set Y_m is H-cocompact, that the stabilizer in H of any cell of Y is of type F_n, and that for each $k \in \mathbb{N}$ there exists $m \in \mathbb{N}$ such that for all vertices $y \in Y$ with $\phi(y) \geq m$, the descending link $\text{lk}^{\downarrow}y$ is k-connected. Then H is of type F_n.*

Note that SV_G acts by right multiplication on the set of corank-1 elements of $S\mathcal{V}_G$, which is compatible with left multiplication by $\mathcal{G}$, and so induces an action of SV_G on P_1. This action is compatible with $\leq$ and so extends to an action on $|P_1|$, under which X is invariant.

Lemma 6.3. *The stabilizer in SV_G of a vertex of X with rank m is isomorphic to $G \wr \Sigma_m$.*

Proof. We have that $a \in SV_G$ stabilizes $[h]$ if and only if $hah^{-1} \in \mathcal{G}$ (recall that $\mathcal{G}$ is the set of twisted permutations). If h has rank m this is equivalent to $hah^{-1} \in \mathcal{G}(m)$, so we have $\text{Stab}_{SV_G}([h]) = h^{-1}(\mathcal{G}(m))h$. Since $\mathcal{G}(m) \cong G \wr \Sigma_m$, we conclude that $\text{Stab}_{SV_G}([h]) \cong G \wr \Sigma_m$. $\qquad\square$

Lemma 6.4. *Let $\gamma \in G$ and let f be a multicolored tree. If γ fixes $\text{Spec}(f)$ pointwise then $[f\tau_\gamma] = [f]$, and if $[f\tau_\gamma] = [f]$ then γ stabilizes $\text{Spec}(f)$.*

Proof. If x_s is a simple split and $\gamma s = s$ then $x_s \tau_\gamma = (\tau_\gamma \oplus \tau_\gamma) x_s$ by Lemma 2.1(7). It follows easily by induction that if f is a multicolored tree and γ fixes $\mathrm{Spec}(f)$ pointwise then $f \tau_\gamma = (\tau_\gamma \oplus \cdots \oplus \tau_\gamma) f$, and therefore $[f \tau_\gamma] = [f]$.

For the second statement, let f be any multicolored tree and suppose that $[f \tau_\gamma] = [f]$. By Lemma 2.3, there exists a G-twist g and a multicolored forest f' so that $f \tau_\gamma = g f'$. Note then that $\mathrm{Spec}(f) = \gamma \, \mathrm{Spec}(f')$. But $[f'] = [g f'] = [f \tau_\gamma] = [f]$ and hence $\mathrm{Spec}(f') = \mathrm{Spec}(f)$, so γ must stabilize $\mathrm{Spec}(f)$. $\qquad\square$

For the following proposition, recall that a group H_1 is *commensurate with* a group H_2 if there exist finite index subgroups of H_1 and H_2 that are isomorphic.

Proposition 6.5. *Let Δ be a simplex in $|P_1|$ corresponding to a chain $v_0 < \cdots < v_k$, where $v_0 = [h]$ and $v_k = [(f_1 \oplus \cdots \oplus f_n)h]$ for some $h \in SV_G$ and some multicolored trees $f_1, \ldots, f_n$. Then the stabilizer of Δ in SV_G is commensurate with $\prod_{i=1}^{n} \mathrm{Stab}_G(\mathrm{Spec}(f_i))$.*

Proof. Since the action of SV_G is order-preserving, any element of SV_G that stabilizes this simplex must fix v_0 and v_k. Moreover, since the interval $[v_0, v_k]$ is finite, the stabilizer of Δ has finite index in the stabilizer of $[v_0, v_k]$.

Now, the stabilizer of v_0 in SV_G is $h^{-1} \mathcal{G}(n) h$, which is isomorphic to $G \wr \Sigma_n$ by Lemma 6.3. This has a finite index subgroup isomorphic to G^n, consisting of all $h^{-1}(\tau_{\gamma_1} \oplus \cdots \oplus \tau_{\gamma_n}) h$ for $\gamma_1, \ldots, \gamma_n \in G$. Such an element fixes v_k if and only if $[f_i \tau_{\gamma_i}] = [f_i]$ for all i. By Lemma 6.4, the group of all such tuples $(\gamma_1, \ldots, \gamma_n)$ is commensurate with $\prod_{i=1}^{n} \mathrm{Stab}_G(\mathrm{Spec}(f_i))$. $\quad\square$

Corollary 6.6. *If G is of type F_n and the stabilizer in G of any finite subset of S is of type F_n, then the stabilizer of any simplex in $|P_1|$ is of type F_n.*

Proof. Finite products of groups of type F_n are of type F_n, and any group commensurate to a group of type F_n is of type F_n, so this is immediate from Proposition 6.5. $\qquad\square$

Let $\phi \colon SV_G \to \mathbb{N}$ be the function "rank" that sends $h \colon \mathfrak{C}^S(m') \to \mathfrak{C}^S(m)$ to m. Denote also by ϕ the restriction of ϕ to $P_1 = X^{(0)}$, and the extension of this to a map $X \to \mathbb{R}$ given by affinely extending to simplices. Note that ϕ is SV_G-invariant. It is a Morse function since it is non-constant on edges, by virtue of a non-trivial multicolored forest having different rank and corank.

For each $m \in \mathbb{N}$ consider the sublevel set $X_m := \phi^{-1}([1, m])$. Since ϕ is invariant under the action of SV_G, each X_m is SV_G-invariant.

Proposition 6.7 (Cocompact). *If the action of G on S is oligomorphic then each X_m is SV_G-cocompact.*

Proof. First note that for any $m \in \mathbb{N}$, SV_G acts transitively on the set of vertices v with $\phi(v) = m$. This is because if $\phi([h]) = \phi([h'])$ then the product $h^{-1}h'$ makes sense, lies in SV_G, and takes $[h]$ to $[h']$. In particular there are finitely many SV_G-orbits of vertices in X_m. It now suffices to show that for a vertex $v \in X_m$, say with $\phi(v) = r \le m$, there are finitely many $\mathrm{Stab}_{SV_G}(v)$-orbits of simplices in X_m containing v as their vertex of minimum rank. Without loss of generality $v = [\mathrm{id}_r]$, so we want to show that there are finitely many $\mathcal{G}(r)$-orbits of simplices in X_m containing $[\mathrm{id}_r]$ as their vertex of minimum rank.

For each $i \in \mathbb{N}$ choose a subset $T_i \subseteq S$ with $|T_i| = i$. Since G is oligomorphic, the proof of Lemma 2.3 implies that there exists $i \in \mathbb{N}$ such that every such $\mathcal{G}(r)$-orbit contains a simplex of the form $[\mathrm{id}_r] < [f_1] < \cdots < [f_k]$ for multicolored forests $f_1, \ldots, f_k$ such that $\mathrm{Spec}(f_1), \ldots, \mathrm{Spec}(f_k) \subseteq T_i$. The result now follows from the fact that for each $i, m \in \mathbb{N}$ there are only finitely many multicolored forests with spectrum lying in T_i and with rank at most m. $\square$

At this point if G is oligomorphic and if every stabilizer in G of a finite subset of S is of type F_n, then the twisted Brin–Thompson group SV_G acts cellularly on the contractible complex X, with stabilizers of type F_n, and we have an SV_G-cocompact filtration $(X_m)_{m \in \mathbb{N}}$ for the Morse function ϕ. In order to show that SV_G is of type F_n it now suffices by Theorem 6.2 to prove that as $\phi([f])$ goes to infinity, the descending link $\mathrm{lk}^{\downarrow}[f]$ of $[f]$ in X becomes arbitrarily highly connected. This is the content of the next section.

7. Descending links

For $[h] \in P_1$ the descending link $\mathrm{lk}^{\downarrow}[h]$ of $[h]$ in X is the subcomplex of $X_{\phi([h])}$ supported on vertices of the form $[f^{-1}gh]$ for g a twisted permutation and f a non-trivial elementary multicolored forest. If $\phi([h]) = m$ then $\mathrm{lk}^{\downarrow}[h] \cong \mathrm{lk}^{\downarrow}[\mathrm{id}_m]$. Let us denote this complex by E_m. A vertex of E_m is an element of P of the form $v = [f^{-1}g]$ for f a non-trivial elementary multicolored forest and $g \in \mathcal{G}(m)$.

Nerve lemma. In the course of analyzing E_m we will need to use the Nerve Lemma, which we record here for reference. Recall that the *nerve* of a covering of a complex Z by subcomplexes $Z_1, \ldots, Z_m$ is the simplicial complex with a vertex for each Z_i and such that $Z_{i_1}, \ldots, Z_{i_k}$ span a simplex whenever $Z_{i_1} \cap \cdots \cap Z_{i_k} \neq \emptyset$.

Lemma 7.1 (Nerve Lemma). [BLVŽ94, Lemma 1.2] *Let Z be a simplicial complex covered by subcomplexes $Z_1, \ldots, Z_m$. Suppose for each $1 \leq r \leq m$ that every non-empty intersection of r of these subcomplexes is $(n-r)$-connected. Then Z is $(n-1)$-connected if and only if the nerve of the covering is $(n-1)$-connected.*

The subcomplex VE_m. We say that a multicolored forest f is *very elementary* if it can be written as a direct sum

$$f = f_1 \oplus \cdots \oplus f_m$$

where each f_i is either the identity homeomorphism on $\mathfrak{C}^S$ or a simple split. Clearly any very elementary multicolored forest is elementary.

Let VE_m be the subcomplex of E_m spanned by vertices of the form $v = [f^{-1}g]$ for f a non-trivial very elementary multicolored forest, and for each $k \in \mathbb{N}$ let

$$\nu(k) = \left\lfloor \frac{k-2}{3} \right\rfloor.$$

Lemma 7.2. *The complex VE_m is $(\nu(m)-1)$-connected.*

Proof. We induct on m (the base case that VE_m is non-empty for $m \geq 2$ is trivially satisfied). Write $\mathfrak{C}^S(m) = \mathfrak{C}_1^S \sqcup \cdots \sqcup \mathfrak{C}_m^S$. For $1 \leq i < j \leq m$ say that a vertex $[f^{-1}g]$ of VE_m *matches* i and j if $f^{-1}g$ sends $\mathfrak{C}_i^S$ and $\mathfrak{C}_j^S$ into the same copy of $\mathfrak{C}^S$ in its range. Note that this is independent of the choice of representative, since postcomposition by a twisted permutation just maps the copies of $\mathfrak{C}^S$ in the range of $f^{-1}g$ homeomorphically to each other. For each $1 \leq i < j \leq m$ let $VE_m^{i,j}$ be the subcomplex of VE_m spanned by vertices v that either match i and j, or else do not match i or j to any k. By construction VE_m is covered by the subcomplexes $VE_m^{i,j}$.

We will apply the Nerve Lemma 7.1 to the covering of VE_m by the $VE_m^{i,j}$. Note that $VE_m^{i,j}$ is the union of the (contractible) stars of its rank-$(m-1)$ vertices that match i and j (with one for each element of S). The intersection of two or more of these stars retracts onto the subcomplex spanned by those v that do not match i or j to any k. This is isomorphic to VE_{m-2}, which is $(\nu(m) - 2)$-connected by induction. This shows that each $VE_m^{i,j}$ is $(\nu(m) - 1)$-connected. It is straightforward to see that $VE_m^{i_1,j_1}, \ldots, VE_m^{i_r,j_r}$ have non-empty intersection if and only if the sets $\{i_1, j_1\}, \ldots, \{i_r, j_r\}$ are pairwise disjoint. Any non-empty such $VE_m^{i_1,j_1} \cap \cdots \cap VE_m^{i_r,j_r}$ is therefore the union of (contractible) stars of its rank-$(m-r)$ vertices that match i_ℓ to j_ℓ for each $1 \leq \ell \leq r$ (with one for each element of S^r). The intersection of two or more of these stars retracts onto the subcomplex spanned by those v that do not match any i_ℓ or j_ℓ to any k. This is isomorphic to VE_{m-2r}, which is $(\nu(m) - r - 1)$-connected by induction. This shows that $VE_m^{i_1,j_1} \cap \cdots \cap VE_m^{i_r,j_r}$ is $(\nu(m) - r)$-connected. In particular the hypotheses of the Nerve Lemma are all satisfied.

The nerve of this covering consists of a vertex for each $VE_m^{i,j}$ and a simplex spanned by any collection of these with non-empty intersection. As we said, $VE_m^{i_1,j_1} \cap \cdots \cap VE_m^{i_r,j_r}$ is non-empty if and only if the sets $\{i_1, j_1\}, \ldots, \{i_r, j_r\}$ are pairwise disjoint. Hence the nerve is precisely the matching complex of the complete graph K_m, which is $(\nu(m) - 1)$-connected ([BLVŽ94, Theorem 4.1]). □

The Morse function μ. Now we want to build up from VE_m to E_m and get high connectivity for E_m. We will use a Morse function similar to the one used in [FMWZ13] for (what here would be called) the $|S| < \infty$ and $G = \{1\}$ case. Let f be an elementary multicolored forest with rank m and corank q. Write the domain of f as

$$\mathfrak{C}^S(q) = R_1 \sqcup \cdots \sqcup R_q$$

and the range as

$$\mathfrak{C}^S(m) = L_1 \sqcup \cdots \sqcup L_m,$$

with each R_i and L_i a copy of $\mathfrak{C}^S$, and call the R_i *roots* and the L_i *leaves*. The image of each root under f is a union of some number of leaves. Define the *weight* of a root to be the number of leaves in its image under f. For each $k \in \mathbb{N}$ define the *kth weight multiplicity* $\mu_k(f)$ of f to be the number of roots with weight k. Note that the μ_k are invariant under left and right multiplication by permutations and twists, so it makes sense to define the kth weight multiplicity of a vertex $v = [f^{-1}g] \in E_m$ to be

$$\mu_k(v) := \mu_k(f).$$

Let $\mu\colon E_m^{(0)} \to \mathbb{N}_0^m$ (where $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$) be the function

$$\mu := (\mu_m, \ldots, \mu_3, \phi).$$

Note that μ_2 and μ_1 are not included in this tuple. Consider the lexicographic ordering on $\mathbb{N}_0^m$, and choose any order-preserving embedding $\mathbb{N}_0^m \to \mathbb{R}$, to view μ as a map from $E_m^{(0)}$ to $\mathbb{R}$. Extend this affinely to the simplices of E_m to get a map $\mu\colon E_m \to \mathbb{R}$.

Observation 7.3. The map $\mu\colon E_m \to \mathbb{R}$ is a Morse function.

Proof. Since $\mu_k \leq m$ for all k and ϕ restricted to E_m is bounded by $m-1$, the image of $E_m^{(0)}$ under μ is finite, hence closed and discrete. By construction, μ restricts to an affine function on each simplex. On any edge of E_m, ϕ is non-constant, hence μ is non-constant. $\qquad\square$

As a remark, the Morse function μ is defined slightly differently than the corresponding Morse function in [FMWZ13], though it dictates the same behavior for descending links.

The μ-descending link. Note that VE_m is precisely the sublevel set $E_m^{\mu \leq (0,\ldots,0,m-1)}$ since a vertex $v \in E_m$ lies in VE_m if and only if $\mu_k(v) = 0$ for all $k \geq 3$ (this is exactly why we did not use μ_2 or μ_1 in the definition of μ). In particular, to inspect the higher connectivity of E_m it suffices to analyze μ-descending links $\mathrm{lk}^{\downarrow}_{\mu} v$ for vertices $v \in E_m$ satisfying $\mu(v) > (0,\ldots,0,m-1)$, i.e., satisfying $\mu_k(v) > 0$ for some $k \geq 3$. Every vertex of $\mathrm{lk}^{\downarrow}_{\mu} v$ is either an upper bound or a lower bound of v (as elements of P), and any upper bound of v is an upper bound of any lower bound of v, so $\mathrm{lk}^{\downarrow}_{\mu} v$ decomposes as the join of two subcomplexes, the μ-*descending split link* and μ-*descending merge link*.

Lemma 7.4 (Descending merge link). *The μ-descending merge link of v is isomorphic to $VE_{\mu_1(v)}$, hence is $(\nu(\mu_1(v)) - 1)$-connected.*

Proof. A vertex w of E_m with $w < v$ is in the μ-descending merge link of v if and only $\mu_k(w) = \mu_k(v)$ for all $k \geq 3$. This is because if $\mu_k(w) < \mu_k(v)$ for $w < v$ then necessarily $\mu_\ell(w) > \mu_\ell(v)$ for some $\ell > k$ (and on the other hand $\phi(w) < \phi(v)$). Suppose $v = [f^{-1}g]$ and write the domain of $f^{-1}g$ as the disjoint union $L_1 \sqcup \cdots \sqcup L_m$ of leaves L_i. Without loss of generality (that is, up to the action of $\mathcal{G}(m)$), we can assume that for each $1 \leq i \leq \mu_1(v)$, L_i does not get mapped by $f^{-1}g$ into the same root as any other L_j. In other words, the first $\mu_1(v)$ of the leaves are the ones revealing $\mu_1(v)$, and note that this property does not depend on the choice of representative of v. Now suppose $w = [f_1^{-1}g_1] < v$ is in the μ-descending merge link of v. For a homeomorphism h with domain $\mathfrak{C}^S(m)$ write $\mathrm{res}_\alpha(h)$ for the restriction of h to $\mathfrak{C}_1^S \sqcup \cdots \sqcup \mathfrak{C}_{\mu_1(v)}^S$ and $\mathrm{res}_\omega(h)$ for the restriction of h to $\mathfrak{C}_{\mu_1(v)+1}^S \sqcup \cdots \sqcup \mathfrak{C}_m^S$. Since $\mu_k(w) = \mu_k(v)$ for all $k \geq 3$ we know that $[\mathrm{res}_\omega(f_1^{-1}g_1)] = [\mathrm{res}_\omega(f^{-1}g)]$. We thus have a well defined injective simplicial map from the μ-descending merge link of v to $VE_{\mu_1(v)}$ given by sending $[f_1^{-1}g_1]$ to $[\mathrm{res}_\alpha(f_1^{-1}g_1)]$. It is straightforward to see that it is also surjective. The connectivity statement follows from Lemma 7.2. $\qquad\square$

The μ-descending split link. Now we focus on the μ-descending split link of $v = [f^{-1}g]$. The μ-descending split link of v lies in the realization of the open interval $(v, [\mathrm{id}_m])$. An element w of $(v, [\mathrm{id}_m])$ is in the μ-descending split link of v if and only $\mu_k(v) > \mu_k(w)$ for some $k \geq 3$. This is because if $\mu_\ell(v) < \mu_\ell(w)$ for $v < w$ then necessarily $\mu_k(v) > \mu_k(w)$ for some $k > \ell$ (and on the other hand $\phi(v) > \phi(w)$ is impossible). Call a root R_i in the range of $f^{-1}g$ *fat* if it has weight at least 3. Any element of $(v, [\mathrm{id}_m])$ is of the form $[ef^{-1}g] = [(f')^{-1}g]$ for e and f' multicolored forests such that $f = f'e$. Note that the roots R_i comprise the domain of e. The above characterization implies that such an element is μ-descending if and only if e is non-trivial on at least one fat root. Consider the subposet

$$D(f) := \{[e] \in ([\mathrm{id}_q], [f]) \mid e \text{ is non-trivial on at least one fat root}\}$$

of $([\mathrm{id}_q], [f])$. The μ-descending split link of $v = [f^{-1}g]$ is isomorphic to $|D(f)|$ via the map $[(f')^{-1}g] \mapsto [e]$ induced by right multiplication by $g^{-1}f$.

Lemma 7.5 (Descending split link, first case). *If $\mu_2(v) \neq 0$ then $|D(f)|$, and hence the μ-descending split link of v, is contractible.*

Proof. There is a poset map π from $[[\mathrm{id}_q], [f]]$ to itself given by sending $[e]$ to $[\bar{e}]$ where $\bar{e}$ is the homeomorphism coinciding with e on each fat root in the domain and acting as the identity on each non-fat root in the domain (picture π as "pruning" away any parts of f on non-fat roots). Since $[\bar{e}] \leq [e]$ and $[\bar{e}] = [e]$ whenever $[e]$ is in the image of π, this yields a retraction of $|[[\mathrm{id}_q], [f]]|$ onto the realization of the image of π. Moreover, if e is non-trivial on at least one fat root then so is $\bar{e}$, so π restricts to a retraction of $|D(f)|$ onto its image under π. Since $\mu_2(v) \neq 0$ we have $\pi([f]) \in ([\mathrm{id}_q], [f])$, so $\pi(D(f))$ is conically contractible with cone point $\pi([f])$ (see [Qui78, Section 1.5]). $\qquad\square$

If $\mu_2(v) = 0$ then $D(f) = ([\mathrm{id}_q], [f])$ and π is the identity, so we need to just inspect $([\mathrm{id}_q], [f])$ directly. Say $f = f_1 \oplus \cdots \oplus f_q$ for f_i of corank 1. For each f_i, recall that $s \in S$ lies in $\mathrm{Spec}(f_i)$ if there exists $c \in \mathfrak{C}^S$ with $f_i(c)(s) \neq c(s)$. Now define the *special spectrum* $\mathrm{SSpec}(f_i)$ to be the set of $s \in \mathrm{Spec}(f_i)$ such that $f_i(c)(s) \neq c(s)$ holds for all $c \in \mathfrak{C}^S$ (rather than just some c). For example the spectrum of $(x_t \oplus \mathrm{id})x_s$ is $\{s, t\}$ and the special spectrum is $\{s\}$, whereas the special spectrum of $(x_t \oplus x_t)x_s$ is the whole spectrum $\{s, t\}$.

In the following note that for $v \in E_m$ we have $\mu_1(v) \leq m - 1$, so $\log_2(m - \mu_1(v))$ is defined (and non-negative).

Lemma 7.6 (Descending split link, second case). *Let $v \in E_m \setminus VE_m$. If $\mu_2(v) = 0$ then $|([\mathrm{id}_q], [f])|$, and hence the μ-descending split link of v, is $(\log_2(m - \mu_1(v)) - 3)$-connected.*

Proof. Our strategy is to cover $([\mathrm{id}_q], [f])$ by the stars of its minimal elements (i.e., atoms in the lattice $[[\mathrm{id}_q], [f]]$). Each minimal element is of the form $z_i^s := [\mathrm{id}_{i-1} \oplus x_s \oplus \mathrm{id}_{q-i}]$ for some i such that R_i is a fat root and some $s \in \mathrm{SSpec}(f_i)$. Write Z_i^s for the star of z_i^s in $|([\mathrm{id}_q], [f])|$. By construction $|([\mathrm{id}_q], [f])|$ is covered by the collection of those Z_i^s such that R_i a fat root and $s \in \mathrm{SSpec}(f_i)$. Any intersection of complexes Z_i^s equals the star in $|([\mathrm{id}_q], [f])|$ of the join of the corresponding elements z_i^s, which is contractible as soon as it is non-empty, and so by the Nerve Lemma 7.1 it suffices to prove that the nerve of this covering is $(\log_2(m - \mu_1(v)) - 3)$-connected. If $\mathrm{SSpec}(f_i) \subsetneq \mathrm{Spec}(f_i)$ for some i

then the z_i^s all have a common upper bound strictly smaller than $[f]$, and so the nerve is a (contractible) simplex. Now assume $\mathrm{SSpec}(f_i) = \mathrm{Spec}(f_i)$ for all i. Then the join of a collection of z_i^s equals $[f]$ if we use every z_i^s and otherwise is a strict lower bound of $[f]$. Hence the nerve of the covering is the boundary of a simplex whose vertices are indexed by pairs (i, s) such that R_i is a fat root and $s \in \mathrm{Spec}(f_i)$. Since $\mathrm{Spec}(f_i)$ is empty if R_i is not fat, the number of vertices of this simplex is $\sum_{i=1}^{q} |\mathrm{Spec}(f_i)|$, and so the nerve is $((\sum_{i=1}^{q} |\mathrm{Spec}(f_i)|) - 3)$-connected. Each f_i has $2^{|\mathrm{Spec}(f_i)|}$ leaves, and f has m leaves, so $m = \sum_{i=1}^{q} 2^{|\mathrm{Spec}(f_i)|}$. Since $\log_2(2^a + 2^b) \leq a + b$ whenever $a, b \geq 1$, we conclude that $\log_2(m - \mu_1(v)) \leq \sum_{i=1}^{q} |\mathrm{Spec}(f_i)|$, so we are done. $\qquad\square$

Remark 7.7. In the case when S is finite and G is trivial, the proof of Lemma 7.6 also serves to correct a mistake in the proof of [FMWZ13, Lemma 3.8]. There the complex called the "up-link" is implied to be a join of some number of non-empty complexes, but in fact it is not quite a join, since we are not allowed to completely split apart every block, so there is one maximal simplex missing. It is easy to show that removing this maximal simplex does not affect the claimed connectivity, and alternatively the nerve argument given here directly proves the claimed connectivity.

Corollary 7.8. *For any vertex* $v \in E_m \setminus VE_m$, *the* μ-*descending link of* v *is* $(\nu(\mu_1(v)) + \log_2(m - \mu_1(v)) - 2)$-*connected. In particular its connectivity is at least the minimum of* $\nu(m/2) - 2$ *and* $\log_2(m/2) - 2$.

Proof. By Lemma 7.4 the μ-descending merge link of v is $(\nu(\mu_1(v)) - 1)$-connected. By Lemmas 7.5 and 7.6, the μ-descending split link of v is $(\log_2(m - \mu_1(v)) - 3)$-connected, so the first claim follows from the fact that the join of an a-connected complex and a b-connected complex is $(a + b + 2)$-connected. To see the second claim, note that if $\mu_1(v) \geq m/2$ then $\nu(\mu_1(v)) + \log_2(m - \mu_1(v)) - 2 \geq \nu(m/2) - 2$ and if $\mu_1(v) \leq m/2$ then $\nu(\mu_1(v)) + \log_2(m - \mu_1(v)) - 2 \geq \log_2(m/2) - 2$. $\qquad\square$

Proof of finiteness theorem. Now that we have analyzed the connectivity of the μ-descending links we are ready to prove our main finiteness theorem. This depends on the following proposition.

Proposition 7.9. *The connectivity of* E_m *is at least the minimum of* $\nu(m/2) - 2$ *and* $\log_2(m/2) - 2$.

Proof. By Corollary 7.8 the connectivity of the μ-descending link of any vertex v in $E_m \setminus VE_m$ is at least the minimum of $\nu(m/2) - 2$ and $\log_2(m/2) - 2$. By Lemma 7.2 VE_m is $(\nu(m) - 1)$-connected, hence $(\nu(m/2) - 2)$-connected, so Morse theory (e.g., [BB97, Corollary 2.6]) tells us that the inclusion $VE_m \to E_m$ induces isomorphisms in π_k for all k up to the minimum of $\nu(m/2) - 2$ and $\log_2(m/2) - 2$ and a surjection in π_k for k one more than this. Hence the connectivity of E_m is at least the minimum of $\nu(m/2) - 2$ and $\log_2(m/2) - 2$. $\qquad\square$

Since both of these functions go to ∞ as m goes to ∞, we have shown that the connectivity of E_m goes to ∞ as m goes to ∞, i.e., the connectivity of $\mathrm{lk}^{\downarrow}[h]$ goes to ∞ as $\phi([h])$ goes to ∞.

Theorem D. *Let G be an oligomorphic group of permutations of a countable set S. Let $n \in \mathbb{N} \cup \{\infty\}$, and suppose that*

    (i) *G is of type F_n, and*

    (ii) *The stabilizer in G of every finite subset of S is of type F_n.*

Then SV_G is of type F_n as well.

Proof. We verify the requirements of Theorem 6.2. The complex X is contractible (Proposition 5.6), the action of SV_G on X has stabilizers of type F_n (Corollary 6.6), and each sublevel set X_m is cocompact (Proposition 6.7). By Proposition 7.9 the descending link of a vertex gets arbitrarily highly connected as the ϕ value of the vertex goes to infinity. Now Theorem 6.2 says SV_G is of type F_n. $\square$

References

[Alo94]   Juan M. Alonso. Finiteness conditions on groups and quasi-isometries. *J. Pure Appl. Algebra*, 95(2):121–129, 1994.

[BB97]   Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[BBM20]   James Belk, Collin Bleak, and Francesco Matucci. Embedding right-angled Artin groups into Brin-Thompson groups. *Math. Proc. Cambridge Philos. Soc.*, 169(2):225–229, 2020.

[BdCK15]   Laurent Bartholdi, Yves de Cornulier, and Dessislava H. Kochloukova. Homological finiteness properties of wreath products. *Q. J. Math.*, 66(2):437–457, 2015.

[BEH20]   Collin Bleak, Luke Elliott, and James Hyde. Sufficient conditions for a group of homeomorphisms of the Cantor set to be two generated. arXiv:2008.04791.

[BF19]   James Belk and Bradley Forrest. Rearrangement groups of fractals. *Trans. Amer. Math. Soc.*, 372(7):4509–4552, 2019.

[BFM+16]   Kai-Uwe Bux, Martin G. Fluch, Marco Marschler, Stefan Witzel, and Matthew C. B. Zaremsky. The braided Thompson's groups are of type F_∞. *J. Reine Angew. Math.*, 718:59–101, 2016. With an appendix by Zaremsky.

[BG84]   Kenneth S. Brown and Ross Geoghegan. An infinite-dimensional torsion-free FP_∞ group. *Inventiones mathematicae*, 77(2):367–381, 1984.

[BL10]   Collin Bleak and Daniel Lanoue. A family of non-isomorphism results. *Geom. Dedicata*, 146:21–26, 2010.

[BLVŽ94]   A. Björner, L. Lovász, S. T. Vrećica, and R. T. Živaljević. Chessboard complexes and matching complexes. *J. London Math. Soc. (2)*, 49(1):25–39, 1994.

[BM16]   James Belk and Francesco Matucci. Röver's simple group is of type F_∞. *Publ. Mat.*, 60(2):501–524, 2016.

[Bri98]   Martin R Bridson. Controlled embeddings into groups that have no non-trivial finite quotients. *Geometry and Topology Monographs*, 1:99–116, 1998.

[Bri04]   Matthew G. Brin. Higher dimensional Thompson groups. *Geom. Dedicata*, 108:163–192, 2004.

[Bri05]   Matthew G. Brin. Presentations of higher dimensional Thompson groups. *J. Algebra*, 284(2):520–558, 2005.

[Bri10]   Matthew G. Brin. On the baker's map and the simplicity of the higher dimensional Thompson groups nV. *Publ. Mat.*, 54(2):433–439, 2010.

[Bro87]   Kenneth S. Brown. Finiteness properties of groups. In *Proceedings of the Northwestern conference on cohomology of groups (Evanston, Ill., 1985)*, volume 44, pages 45–75, 1987.

[Bro92]   Kenneth S. Brown. The geometry of finitely presented infinite simple groups. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 121–136. Springer, New York, 1992.

[Cam90]   Peter J. Cameron. *Oligomorphic permutation groups*, volume 152 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1990.

[CFP96]   J. W. Cannon, W. J. Floyd, and W. R. Parry. Introductory notes on Richard Thompson's groups. *Enseign. Math. (2)*, 42(3-4):215–256, 1996.

[CR09]   Pierre-Emmanuel Caprace and Bertrand Rémy. Simplicity and superrigidity of twin building lattices. *Invent. Math.*, 176(1):169–221, 2009.

[CR10]   Pierre-Emmanuel Caprace and Bertrand Rémy. Non-distortion of twin building lattices. *Geom. Dedicata*, 147:397–408, 2010.

[DS20]   Arman Darbinyan and Markus Steenbock. Embeddings into left-orderable simple groups. arXiv:2005.06183.

[FMWZ13]   Martin G. Fluch, Marco Marschler, Stefan Witzel, and Matthew C. B. Zaremsky. The Brin–Thompson groups sV are of type F_∞. *Pacific J. Math.*, 266(2):283–295, 2013.

[Gor74]   A. P. Goryushkin. Imbedding of countable groups in 2-generated simple groups. *Mathematical Notes*, 16(2):725–727, 1974.

[Hal74]   P. Hall. On the embedding of a group in a join of given groups. *Journal of the Australian Mathematical Society*, 17(4):434–495, 1974.

[HM12]   Johanna Hennig and Francesco Matucci. Presentations for the higher-dimensional Thompson groups nV. *Pacific J. Math.*, 257(1):53–74, 2012.

[Hyd17]   James Thomas Hyde. *Constructing 2-generated subgroups of the group of homeomorphisms of Cantor space*. PhD thesis, University of St Andrews, 2017.

[KMPN13]   Dessislava H. Kochloukova, Conchita Martínez-Pérez, and Brita E. A. Nucinkis. Cohomological finiteness properties of the Brin-Thompson-Higman groups $2V$ and $3V$. *Proc. Edinb. Math. Soc. (2)*, 56(3):777–804, 2013.

[Qui78]   Daniel Quillen. Homotopy properties of the poset of nontrivial p-subgroups of a group. *Adv. in Math.*, 28(2):101–128, 1978.

[Sch76]   Paul E. Schupp. Embeddings into simple groups. *J. London Math. Soc. (2)*, 13(1):90–94, 1976.

[Ste92]   Melanie Stein. Groups of piecewise linear homeomorphisms. *Trans. Amer. Math. Soc.*, 332(2):477–514, 1992.

[SWZ19]   Rachel Skipper, Stefan Witzel, and Matthew C. B. Zaremsky. Simple groups separated by finiteness properties. *Invent. Math.*, 215(2):713–740, 2019.

[Wit19]   Stefan Witzel. Classifying spaces from Ore categories with Garside families. *Algebr. Geom. Topol.*, 19(3):1477–1524, 2019.

[WZ18]   Stefan Witzel and Matthew Zaremsky. Thompson groups for systems of groups, and their finiteness properties. *Groups Geom. Dyn.*, 12(1):289–358, 2018.

SCHOOL OF MATHEMATICS AND STATISTICS, UNIVERSITY OF ST ANDREWS, ST ANDREWS, KY16 9SS, SCOTLAND.

Email address: jmb42@st-andrews.ac.uk

DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY AT ALBANY (SUNY), ALBANY, NY 12222.

Email address: mzaremsky@albany.edu