HIGHER GENERATION FOR PURE BRAID GROUPS

MATTHEW C. B. ZAREMSKY

ABSTRACT. We exhibit some families of subgroups of the pure braid group that are highly generating, in the sense of Abels and Holz. In one class of examples, the relevant geometric object is a complex termed the restricted arc complex of a surface. Another arises by considering "dangling braiges," introduced by Bux, Fluch, Schwandt, Witzel and the author.

INTRODUCTION

The notion of a family of subgroups of a group being *highly generating* was introduced by Abels and Holz [AH93]. It is a very natural condition, with many strong consequences, but to date few examples have been explicitly constructed of highly generating subgroups for "interesting" groups. One prominent existing example, given by Abels and Holz, is standard parabolic subgroups of Coxeter groups, or standard parabolic subgroups of groups with a BN-pair. The relevant geometry is given by Coxeter complexes and buildings. Higher generation is also used in [MMV98] as a tool to calculate the Bieri-Neumann-Strebel-Renz invariants of right-angled Artin groups.

As an addition to the collection of interesting examples, we produce two classes of families of subgroups of the n-string pure braid group PB_n that we show to be highly generating. In the first case the geometry is given by certain complexes of arcs on a surface. In the second case the geometry is given by complexes of what are called "dangling braiges." In both cases, the definitions and techniques are heavily informed by the paper [BFS+12], in which arc complexes and braige complexes are used to prove that the braided Thompson's groups V_{br} and F_{br} are of type F_∞.

In Section 1 we recall some definitions and results from [AH93], and establish a criterion for detecting coset complexes in Proposition 1.5. In Section 2.1 we define the *restricted arc complex* on a surface, and in Section 2.2 we define the complex of *dangling pure braiges*. The relevant families of subgroups of PB_n are defined in the paragraphs before Lemma 2.6 and Corollary 2.11, and in Definition 2.12. Finally in Section 3 we calculate the connectivity of these complexes and deduce that the families of subgroups are highly generating. See Propositions 3.6 and 3.13 for the exact bounds.

Acknowledgments. This paper was inspired by observations made during the course of writing [BFS+12], and thanks are due to Kai-Uwe Bux, Martin Fluch, Marco Schwandt and Stefan Witzel for their part in the development of the present results. I should also thank Bux for the suggestion that I write these results up, after which they became much more precise and well formed. Finally, as in [BFS+12], I am grateful to Andy Putman for suggesting a new strategy to handle the complexes $\mathcal{MA}(\Gamma)$ in Section 2.1, after which Theorem 3.1 is possible.

1. HIGHER GENERATION

Higher generation is defined using nerves of coverings of groups by cosets. The relevant definitions are as follows.

Definition 1.1 (Nerve). Let X be a set and $\mathcal{U}$ a collection of subsets covering X. The *nerve* of the cover $\mathcal{U}$, denoted $\mathcal{N}(\mathcal{U})$, is a simplicial complex with vertex set $\mathcal{U}$, such that pairwise distinct vertices $U_0, \ldots, U_k$ span a k-simplex if and only if $U_0 \cap \cdots \cap U_k \neq \emptyset$.

The type of nerve we are interested in is the following *coset complex*.

Date: December 6, 2012.
2010 *Mathematics Subject Classification.* Primary 20F65; Secondary 20F36.
Key words and phrases. Braid group, higher generation, arc complex.
The author was supported by the SFB 701 of the DFG.

Definition 1.2 (Coset complex and higher generation)**.** Let G be a group and $\mathscr{F}$ a family of subgroups. Let $\mathcal{U} := \coprod_{H \in \mathscr{F}} G/H$ be the covering of G by cosets of subgroups in $\mathscr{F}$. We call $\mathcal{N}(\mathcal{U})$ the *coset complex* of G with respect to $\mathscr{F}$, and denote it $\mathrm{CC}(G, \mathscr{F})$. We say that $\mathscr{F}$ *n-generates* G if $\mathrm{CC}(G, \mathscr{F})$ is $(n-1)$-connected, and *∞-generates* G if $\mathrm{CC}(G, \mathscr{F})$ is contractible.

The following theorem indicates some ways higher generation can be used. The first part says that 1-generation equals generation, and the second part says that a 2-generating family yields a decomposition of G as an amalgamated product.

Theorem 1.3. [AH93, Theorem 2.4] *Let $\mathscr{F} = \{H_\alpha \mid \alpha \in \Lambda\}$ be a family of subgroups of G.*

(1) *$\mathscr{F}$ is 1-generating if and only if $\bigcup H_\alpha$ generates G.*
(2) *$\mathscr{F}$ is 2-generating if and only if the natural map $\coprod_{\cap} H_\alpha \to G$ is an isomorphism.*

Here by $\coprod_{\cap} H_\alpha$ we mean the amalgamated product of the H_α over their intersections. We remark that another equivalent condition in part (1) is that the map $\coprod_{\cap} H_\alpha \to G$ be surjective.

An important observation about coset complexes is that the action of the group on the complex has a very nice fundamental domain.

Observation 1.4 (Fundamental domain)**.** With the above notation, assume $\mathscr{F}$ is finite. Since $\bigcap_{H \in \mathscr{F}} H \neq \emptyset$, we see that $\mathscr{F}$ itself is the vertex set of a maximal simplex in $\mathrm{CC}(G, \mathscr{F})$. This maximal simplex, which we call C, is a fundamental domain for the action of G on $\mathrm{CC}(G, \mathscr{F})$ by left multiplication.

Proof. For any simplex σ in $\mathrm{CC}(G, \mathscr{F})$, there exist $H_0, \dots, H_k \in \mathscr{F}$ and $g \in G$ such that the vertices of σ are the cosets gH_i for $0 \le i \le k$. Then $g^{-1}\sigma$ is a face of C. This shows that every G-orbit intersects C, and indeed intersects C uniquely since if $gH_i = H_j$ then $g \in H_i = H_j$. □

A sort of converse of this observation is the following proposition, which allows us to detect highly generating families of subgroups as stabilizers of "nice" actions.

Proposition 1.5 (Detecting coset complexes)**.** *Let G be a group acting by simplicial automorphisms on a simplicial complex X, with a single maximal simplex C as fundamental domain. Let $\mathscr{F} := \{\mathrm{Stab}_G(v) \mid v$ is a vertex of $C\}$. Then $\mathrm{CC}(G, \mathscr{F})$ is isomorphic to X as a simplicial G-complex.*

Proof. Define a map $\phi\colon \mathrm{CC}(G, \mathscr{F}) \to X$ by sending the coset $g\,\mathrm{Stab}_G(v)$ to the vertex gv of X. This is a G-invariant map between the 0-skeleta, and it induces a simplicial map since the vertices of a simplex in $\mathrm{CC}(G, \mathscr{F})$ can be represented as cosets with a common left representative. Since C is a fundamental domain, ϕ is bijective. □

A good first example is when X is a tree, on which a group G acts edge transitively and without inversion. Then Theorem 1.3 and Proposition 1.5 imply that G decomposes as an amalgamated product. Namely, if e is a fundamental domain with endpoints v and w, then $G = G_v *_{G_e} G_w$ (this is standard Bass-Serre theory). Indeed, the vertex stabilizers are not just 2-generating, but ∞-generating.

This example is generalized by looking at groups acting on buildings.

Example 1.6 (Buildings)**.** Let G be a group acting chamber transitively on a building Δ, by type preserving automorphisms. See [AB08] for the relevant background. Let C be the fundamental chamber, and let $\mathscr{F} := \{\mathrm{Stab}_G(v) \mid v \text{ is a vertex of } C\}$. Then $\mathrm{CC}(G, \mathscr{F}) \cong \Delta$, and so $\mathscr{F}$ is highly generating for G. More precisely, if Δ is spherical of dimension n then $\mathscr{F}$ is n-generating, and if Δ is not spherical then $\mathscr{F}$ is ∞-generating. If the action is not just chamber transitive, but is even *Weyl transitive*, as in [AB08, Chapter 6], then the stabilizers $\mathrm{Stab}_G(v)$ are precisely the maximal *standard parabolic subgroups*. An even stronger condition is that the action is *strongly transitive*, in which case G has a BN-pair, and we recover the situation in [AH93, Section 3.2].

We also have examples from the world of Artin groups.

Example 1.7 (Deligne complexes)**.** Background for this example can be found in [CD95]. Let (A, S) be an Artin system with associated Coxeter system (W, S). For $T \subseteq S$ let A_T (respectively W_T) be the subgroup generated by T. Let $\widehat{\mathscr{F}} := \{A_T \mid T \subseteq S\}$ and $\mathscr{F} := \{A_T \mid T \subseteq S \text{ with } |W_T| < \infty\}$. The coset complexes $\mathrm{CC}(A, \widehat{\mathscr{F}})$ and $\mathrm{CC}(A, \mathscr{F})$ are, up to homotopy equivalence, the *Deligne complex* and *modified Deligne complex* of A. The connectivity of these complexes, and hence the higher generation properties of these families of subgroups, is tied to the $K(\pi, 1)$ Conjecture described in [CD95]. Namely, $\mathscr{F}$ is conjecturally ∞-generating; see [CD95, Conjecture 2]. This is known to hold for many Artin groups, including for braid groups.

2. Arc complexes and braige complexes

In this section we define and analyze some complexes on which the braid group and pure braid group act. After recalling some background on braid groups, in the first subsection we look at the *restricted arc complex* on a surface, and in the second subsection we look at the complex of *dangling (pure) braiges*. Arc complexes are well-studied, and the restricted arc complex here will provide a coset complex for PB_n using arc stabilizers as subgroups. Braige complexes are more esoteric (the name "braige" would only be familiar if the reader has seen [BFS+12]), but these will provide a coset complex for PB_n using subgroups obtained via the "strand cloning maps." These subgroups are smaller than the arc stabilizers, and more visualizable when using strand pictures for braids. We will save the connectivity calculations for Section 3, after which we will conclude that these families of subgroups are highly generating.

Before introducing the complexes, we quickly recall some background on braid groups. Let B_n be the braid group on n strands. This is the Artin group for the symmetric group Σ_n, so we have a projection $B_n \to \Sigma_n$. The kernel of this projection is the *pure braid group* on n strands, PB_n. Moreover, B_n is isomorphic to the mapping class group of the n-punctured disk D_n [Bir74], and PB_n is obtained by treating the punctures as distinguished points and only considering homeomorphisms taking each puncture to itself.

2.1. Arc complexes.

Let S be a connected surface, possibly with boundary ∂S. Let $P \subseteq S \setminus \partial S$ be a set of n points in the interior of the surface. An *arc* is a simple path in $S \setminus \partial S$ that intersects P precisely at its endpoints, and whose endpoints are distinct. Our reference for arc complexes is [Hat91]. This definition of arc is different from the definition in [Hat91], in that we do not allow the endpoints of a given arc to coincide. Also, in [Hat91], points in P may be contained in ∂S, which we do not allow.

Let $\{\alpha_0, \dots, \alpha_k\}$ be a collection of arcs. If the α_i are pairwise disjoint except possibly at endpoints, and if no distinct α_i and α_j are homotopic relative P, we call $\{\alpha_0, \dots, \alpha_k\}$ an *arc system*. The homotopy classes, relative P, of arc systems form the simplices of a simplicial complex, with the face relation given by passing to subcollections of arcs.

Definition 2.1 (Arc complex)**.** Let Γ be a simplicial graph with $|P|$ vertices, and identify P with the set of vertices of Γ. Call an arc *compatible* with Γ if its endpoints are connected by an edge in Γ. Let $\mathcal{HA}(\Gamma)$ be the *arc complex* on (S, P) corresponding to Γ, that is the simplicial complex with a k-simplex for each arc system $\{\alpha_0, \dots, \alpha_k\}$ such that all the α_i are compatible with Γ. When extra precision is required, we will write $\mathcal{HA}(S, P, \Gamma)$.

Given an arc system $\sigma = \{\alpha_0, \dots, \alpha_k\}$ in $\mathcal{HA}(\Gamma)$, denote by Γ_σ the following subgraph of Γ. Every vertex of Γ is a vertex of Γ_σ, and an edge e of Γ is in Γ_σ if and only if the endpoints of e are the endpoints of some α_i. Call Γ_σ *faithful* if it has precisely $(k+1)$ edges. Since we only consider simplicial graphs, i.e., there are no loops or multiple edges, this condition is equivalent to saying that no distinct α_i, α_j share both endpoints (they may share one).

Of particular interest is the complete graph K_n, which is the graph with n vertices and a single edge between any two distinct vertices. We remark that $\mathcal{HA}(K_n)$ is a proper subcomplex of the complex $\mathcal{A}(S, P)$ in [Hat91], since we only consider arcs with two distinct endpoints. We are also interested in the family of subgraphs of *linear graphs* L_n. The linear graph L_n has n vertices labeled 1 through n, and $n-1$ edges, one connecting the vertex labeled i to the vertex labeled $i+1$ for each $1 \leq i < n$. We will be interested in subgraphs of L_n having the same vertex set as L_n.

Terminological convention: From now on, a *subgraph* Γ' of a graph Γ always has the same vertex set as Γ.

An important subcomplex of $\mathcal{HA}(\Gamma)$ is the *matching complex* on the surface S with distinguished points P, introduced in [BFS+12]. First we define the matching complex of a graph, and then the matching complex on a surface.

Definition 2.2 (Matching complex of a graph)**.** Let Γ be a subgraph of K_n. The *matching complex* $\mathcal{M}(\Gamma)$ of Γ is the simplicial complex with a k-simplex for each collection of $k+1$ pairwise disjoint edges $\{e_0, \dots, e_k\}$ of Γ, and face relation given by passing to subcollections. Observe that if Γ' is a subgraph of Γ then $\mathcal{M}(\Gamma')$ is a subcomplex of $\mathcal{M}(\Gamma)$.

Definition 2.3 (Matching complex on a surface)**.** Let $\mathcal{MA}(K_n)$ be the subcomplex of $\mathcal{HA}(K_n)$ whose simplices are given by arc systems whose arcs are pairwise disjoint including at their endpoints. For a subgraph Γ of K_n let $\mathcal{MA}(\Gamma)$ be the preimage of $\mathcal{M}(\Gamma)$ under the map $\mathcal{MA}(K_n) \to \mathcal{M}(K_n)$ that sends an arc with endpoints labeled i and j to the edge of K_n with endpoints i and j. We call $\mathcal{MA}(\Gamma)$ the *matching complex* on (S, P) corresponding to Γ. We may also write $\mathcal{MA}(S, P, \Gamma)$.

The complex we are presently interested in is a complex $\mathcal{RA}(\Gamma)$, which we will call the *restricted arc complex.*

Definition 2.4 (Restricted arc complex)**.** The *restricted arc complex* $\mathcal{RA}(\Gamma)$ on (S, P) corresponding to Γ is the subcomplex of $\mathcal{HA}(\Gamma)$ consisting of arc systems σ for which Γ_σ is faithful. We may also write $\mathcal{RA}(S, P, \Gamma)$.

We could equivalently require that the subspace of S given by the union of the arcs is a simplicial graph, i.e., has no multiple edges. In this way we can view $\mathcal{RA}(\Gamma)$ as the complex of embeddings of subgraphs of Γ into S that send vertices in a prescribed way to the points of P. The $\Gamma = L_n$ case is especially nice, since all of L_n can be embedded into any connected surface. In fact, every simplex of $\mathcal{RA}(L_n)$ is a face of a maximal simplex of dimension $n-2$. See Figure 1 for some examples of arc systems.

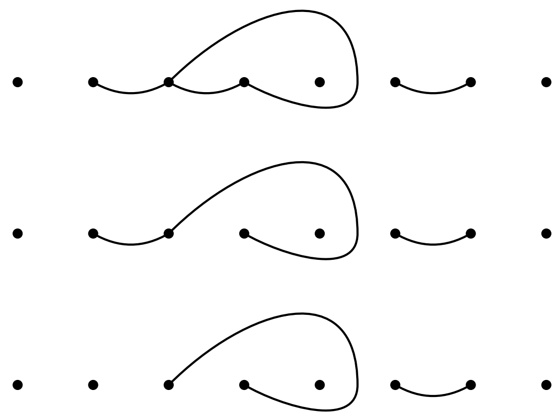

FIGURE 1. From top to bottom, an arc system in $\mathcal{HA}(L_8) \setminus \mathcal{RA}(L_8)$, one in $\mathcal{RA}(L_8) \setminus \mathcal{MA}(L_8)$ and one in $\mathcal{MA}(L_8)$.

Remark 2.5. Embedding graphs into surfaces is an interesting enterprise in its own right, so the complex $\mathcal{RA}(\Gamma)$ may be of further general interest. For instance, the dimension of $\mathcal{RA}(S, P, \Gamma)$ is one less than the number of edges in a maximal subgraph of Γ embeddable into (S, P).

Being the mapping class group of D_n, B_n acts on $\mathcal{HA}(K_n)$, and this action stabilizes $\mathcal{MA}(K_n)$ and $\mathcal{RA}(K_n)$. For general Γ, B_n will not necessarily stabilize $\mathcal{HA}(\Gamma)$, since general braids may not stabilize P pointwise. However, pure braids do stabilize P pointwise, and so PB_n stabilizes $\mathcal{HA}(\Gamma)$, $\mathcal{MA}(\Gamma)$ and $\mathcal{RA}(\Gamma)$ for any Γ.

Denote by $[m]$ the set $\{1, \dots, m\}$ for $m \in \mathbb{N}$. Let S be the unit disk, and fix an embedding $L_n \hookrightarrow S$ of the linear graph with n vertices into S. Let P be the image of the vertex set, so P is a set of n points in S, labeled 1 through n. Under this embedding, the edges of L_n yield a maximal simplex of $\mathcal{RA}(L_n)$, which we will denote C. For each $\emptyset \neq J \subseteq [n-1]$ define σ_J to be the face of C consisting only of those arcs with endpoints $j, j+1$ for $j \in J$. In particular, σ_J is a $(|J|-1)$-simplex in $\mathcal{RA}(L_n)$.

For each $\emptyset \neq J \subseteq [n-1]$ define

$$PB_n^J := \mathrm{Stab}_{PB_n}(\sigma_J)$$

and set $\mathscr{AF}_n := \{PB_n^J \mid \emptyset \neq J \subseteq [n-1] \text{ with } |J| = 1\}$.

Lemma 2.6. *The coset complex* $\mathrm{CC}(PB_n, \mathscr{AF}_n)$ *and the restricted arc complex* $\mathcal{RA}(L_n)$ *are isomorphic as simplicial* PB_n*-complexes.*

Proof. It suffices by Proposition 1.5 to show that C is a fundamental domain for the action of PB_n on $\mathcal{RA}(L_n)$. A maximal simplex of $\mathcal{RA}(L_n)$ is an embedding of L_n into S, such that the vertex labeled i maps to the point in P labeled i, for each $1 \leq i \leq n$. Any such simplex is in the PB_n-orbit of C. Moreover, if $p\sigma_J = \sigma_K$ for $p \in PB_n$ and σ_J, σ_K are faces of C, then since p is pure we know that $J = K$. We conclude that C is a fundamental domain. □

In Section 3 we will calculate the connectivity of $\mathcal{RA}(L_n)$, and deduce that $\mathscr{AF}_n$ is highly generating for PB_n. Before doing that, we describe another complex with a nice PB_n action.

2.2. **Braige complexes.** The terminology "braige" is short for *braid-merge.* For our purposes, a braige consists of a braid whose strands may merge together at the bottom, in some sense. A more formal definition is as follows.

Definition 2.7 (Braiges)**.** A *braige* on n strands is a pair (b, Γ), consisting of a braid $b \in B_n$ and a subgraph Γ of L_n. If the edges of Γ are disjoint, we call (b, Γ) *elementary.* If the braid is pure, then the braige is a *pure braige.* See Figure 2 for some examples.

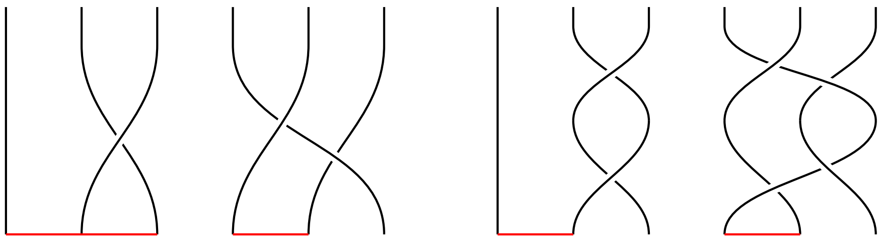

FIGURE 2. A braige on 6 strands and an elementary pure braige on 6 strands.

As a remark, this definition is *not* the same as the one in [BFS+12]. Here a "merge" amounts to choosing some pairs of adjacent strands that should be stuck together at the bottom with edges. In [BFS+12], the merging is more subtle; strands merge two at a time, not in a square shape but in more of a triangle, and a new strand continues down out of the merge. This new strand may merge further with other strands, but one must keep track of the order of merging. However, the notions of elementary braiges are the same here and in [BFS+12], since it does not matter in which order the merges occur.

Let $\mathscr{B}_n(L_n)$ be the set of all braiges on n strands. There is a left action of B_n on $\mathscr{B}_n(L_n)$, via $b(c, \Gamma) := (bc, \Gamma)$. We can think of $\mathscr{B}_n(L_n)$ as a simplicial complex, where (b, Γ) is a face of (b', Γ') if $b = b'$ and Γ' is a subgraph of Γ. Restricting to pure braids, we get the set $\mathscr{PB}_n(L_n)$ of pure braiges, with an action of PB_n. A nice feature of this is that (id, L_n) is a fundamental domain for the action of B_n on $\mathscr{B}_n(L_n)$, or PB_n on $\mathscr{PB}_n(L_n)$. However, it is easy to see that $\mathscr{B}_n(L_n)$ and $\mathscr{PB}_n(L_n)$ stand little chance of being connected, since we can only "move" by changing the merges, and not the braid. To get a highly connected complex, we will consider an equivalence relation on these complexes via the notion of "dangling." First we need to define what it means for a strand in a braid to be a *clone.*

Definition 2.8 (Clones)**.** Let $b \in B_n$. Number the strands of b from left to right at their tops by $1, \ldots, n$. Let ρ_b be the permutation induced by b under $B_n \to \Sigma_n$. Think of b as living in 3-space $\mathbb{R}^3$, with the top of the i^{th} strand at the point $(i, 1, 0)$ and the bottom at $(\rho_b(i), 0, 0)$, for each $i \in [n]$. In particular all the tops and bottoms of the strands are in the xy-plane. Note that for any given strand, b has a representation wherein that strand is entirely contained in the xy-plane. Now suppose that for some $i \in [n-1]$, b can be represented in such a way that the i^{th} and $(i+1)^{\text{st}}$ strands are *simultaneously* in the xy-plane, and moreover, no strands of the braid other than those two intersect the closed region of the xy-plane bounded by the two strands and the line segments from $(i,1,0)$ to $(i+1,1,0)$ and from $(\rho_b(i),0,0)$ to $(\rho_b(i+1),0,0)$. In this case we will refer to the $(i+1)^{\text{st}}$ strand as a *clone*, specifically a *clone of the i^{th} strand.* Note that necessarily $\rho_b(i+1) = \rho_b(i)+1$.

Our convention is to always consider the strand on the right to be the clone of the strand on the left, as opposed to the other way around. See Figure 3 for an example.

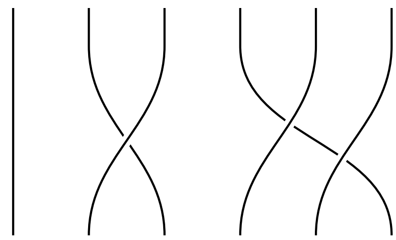

FIGURE 3. The sixth strand is a clone of the fifth.

For each $i \in [n-1]$ there is a *cloning map* $\kappa_i \colon B_{n-1} \to B_n$ given by cloning the i^{th} strand. This is not a homomorphism, but becomes one when restricted to $\kappa_i \colon PB_{n-1} \to PB_n$. For $I = \{i_1, \dots, i_r\} \subseteq [n-r]$, with $i_1 < \cdots < i_r$, define the cloning map $\kappa_I := \kappa_{i_1} \circ \cdots \circ \kappa_{i_r} \colon B_{n-r} \to B_n$. The restriction $\kappa_I \colon PB_{n-r} \to PB_n$ is again a homomorphism. Now for $J = \{j_1, \dots, j_r\} \subseteq [n-1]$, with $j_1 < \cdots < j_r$, let $I_J \subseteq [n-r]$ be the set $\{j_i - (i-1) \mid 1 \le i \le r\}$. The point is that a braid $b \in B_n$ is in the image of κ_{I_J} if and only if for each $j \in J$, the $(j+1)^{\text{st}}$ strand is a clone of the j^{th} strand. Denote the subset of such braids by $B_n^{(J)}$, and the sub*group* of such pure braids by $PB_n^{(J)}$. (The parentheses distinguish $PB_n^{(J)}$ from the arc system stabilizer PB_n^J from the previous section.)

We can now define the equivalence relation between braiges, given by "dangling."

Definition 2.9 (Dangling)**.** Let (b,Γ) be a braige on n strands, and number the vertices of Γ by $1,\dots,n$ from left to right. Let $J_\Gamma \subseteq [n-1]$ be the set of left endpoints of edges of Γ. Now consider any braid c from the set $B_n^{(J_\Gamma)}$. For each $j \in J_\Gamma$, we know that $\rho_c(j+1) = \rho_c(j)+1$, so there is a subgraph of L_n whose edges are precisely those connecting $\rho_c(j)$ and $\rho_c(j+1)$ for $j \in J_\Gamma$. Call this graph Γ^c. The point is that, if we draw c below the braige, and "pull" the merges through c, we get the braige (bc,Γ^c). Now declare that (b,Γ) is equivalent to (bc,Γ^c) for each $c \in B_n^{(J_\Gamma)}$. One checks that this is an equivalence relation, called equivalence *under dangling.* Denote the equivalence class of (b,Γ) by $[(b,\Gamma)]$, and call it a *dangling braige.* The idea is that the top of a braige is static, but the strands at the bottom are free to "dangle," modulo the restriction that the merges remain rigid (and oriented) during the dangling. We analogously get the notion of a *dangling pure braige*, where we only consider c as above coming from $PB_n^{(J_\Gamma)}$, so in particular Γ^c always equals Γ in the pure case. An example of dangling can be seen in Figure 4.

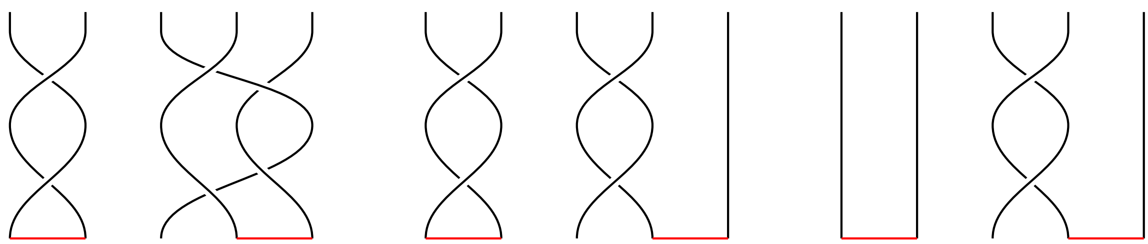

FIGURE 4. The first and second braiges are equivalent under pure dangling, but are *not* equivalent to the third one.

Let $\mathcal{B}_n(L_n)$ be the set of equivalence classes under dangling of braiges in $\mathscr{B}_n(L_n)$. The simplicial structure of the latter induces a simplicial structure on the former, for example the faces of $[(b, \Gamma)]$ are precisely of the form $[(bc, \Gamma')]$, for $c \in B_n^{(J_\Gamma)}$ and Γ' a subgraph of Γ^c. Also let $\mathcal{PB}_n(L_n)$ be the set of dangling pure braiges. The faces of a dangling pure braige $[(p, \Gamma)]$ are the dangling pure braiges of the form $[(pc, \Gamma')]$ for $c \in PB_n^{(J_\Gamma)}$ and Γ' a subgraph of Γ. Heuristically, in $\mathcal{B}_n(L_n)$ we can move around not only by changing the merges, but now also by changing the braid in certain controlled ways, so $\mathcal{B}_n(L_n)$ and $\mathcal{PB}_n(L_n)$ stand a chance of being connected (for large enough n), and even highly connected. In the pure case we can also define $\mathcal{PB}_n(\Gamma)$ for any subgraph Γ of L_n, by only considering braiges from $\mathscr{PB}_n(\Gamma)$. We also have the subcomplexes of dangling *elementary* braiges or dangling elementary pure braiges, denoted $\mathcal{EB}_n(L_n)$ and $\mathcal{EPB}_n(L_n)$ respectively. In the pure case we can use any Γ instead of L_n, and get the complex $\mathcal{EPB}_n(\Gamma)$. This will be an important subcomplex for proving that $\mathcal{PB}_n(L_n)$ is highly connected.

The left action of B_n on $\mathscr{B}_n(L_n)$ induces an action of B_n on $\mathcal{B}_n(L_n)$; for $c \in B_n$ we have $c[(b, \Gamma)] := [(cb, \Gamma)]$. Similarly, PB_n acts from the left on $\mathcal{PB}_n(L_n)$, and indeed stabilizes $\mathcal{PB}_n(\Gamma)$ for any subgraph Γ of L_n. The action of PB_n on $\mathcal{PB}_n(L_n)$ is of particular interest, since there is a fundamental domain consisting of a single maximal simplex, namely $[(\mathrm{id}, L_n)]$. This tells us that $\mathcal{PB}_n(L_n)$ is a coset complex, using the family of stabilizers of faces of $[(\mathrm{id}, L_n)]$.

Lemma 2.10 (Stabilizers of dangling braiges)**.** *Let Γ be a subgraph of L_n. Then the stabilizer $\mathrm{Stab}_{PB_n}([(\mathrm{id}, \Gamma)])$ is precisely the subgroup $PB_n^{(J_\Gamma)}$.*

Proof. First let $p \in PB_n^{(J_\Gamma)}$. Then $p[(\mathrm{id}, \Gamma)] = [(p, \Gamma)] = [(\mathrm{id}, \Gamma)]$. Now suppose $p[(\mathrm{id}, \Gamma)] = [(\mathrm{id}, \Gamma)]$, so $[(p, \Gamma)] = [(\mathrm{id}, \Gamma)]$. Then there exists $c \in PB_n^{(J_\Gamma)}$ such that $(p, \Gamma) = (c, \Gamma)$. But this implies that $p = c$, so we are done. □

Let $\mathscr{BF}_n := \{PB_n^{(J_\Gamma)} \mid \Gamma \text{ is a subgraph of } L_n \text{ with one edge}\}$.

Corollary 2.11. *$\mathrm{CC}(PB_n, \mathscr{BF}_n)$ is isomorphic to $\mathcal{PB}_n(L_n)$ as a simplicial PB_n-complex.*

Proof. This is immediate from Proposition 1.5, since $[(\mathrm{id}, L_n)]$ is a fundamental domain. □

In the next section we will calculate the connectivity of $\mathcal{RA}(L_n)$ and $\mathcal{PB}_n(L_n)$, and hence of $\mathrm{CC}(PB_n, \mathscr{AF}_n)$ and $\mathrm{CC}(PB_n, \mathscr{BF}_n)$, from which we deduce higher generation.

We close this section by setting up a generalization of the complexes we have constructed. Note that in the definition of $\mathscr{AF}_n$ we require $|J| = 1$, and in the definition of $\mathscr{BF}_n$ we require Γ to have only one edge (this is the same as saying $|J_\Gamma| = 1$). The subgroups in these families consist of braids that, respectively, stabilize some arc, or feature at least one cloned strand. Of course, as n grows, it becomes increasingly "easy" for a braid to be very complicated while still featuring a cloned strand, or stabilizing an arc. Hence, higher generation becomes an even more interesting question if we consider requirements like, e.g., *all but 5 strands are clones.* (Observe that any of the standard generators of PB_n satisfy this very requirement.)

Definition 2.12 (More restrictive families)**.** Let $s \in \mathbb{N}$. Define $\mathscr{AF}_n^s := \{PB_n^J \mid J \subseteq [n-1] \text{ with } |J| = s\}$. Also define $\mathscr{BF}_n^s := \{PB_n^{(J_\Gamma)} \mid \Gamma \text{ is a subgraph of } L_n \text{ with } s \text{ edges}\}$. Hence $\mathscr{AF}_n^1 = \mathscr{AF}_n$ and $\mathscr{AF}_n^{n-1} = \{Z(PB_n)\}$, and also $\mathscr{BF}_n^1 = \mathscr{BF}_n$ and $\mathscr{BF}_n^{n-1} = \{\{1\}\}$.

3. Connectivity of the complexes

For $\ell \in \mathbb{Z}$ define $\eta(\ell) := \lfloor \frac{\ell-2}{4} \rfloor$. The main goal of this section is to prove that $\mathcal{RA}(L_n)$ and $\mathcal{PB}_n(L_n)$ are $(\eta(n) - 1)$-connected.

The only blackbox we will use in this section is the following result from [BFS+12]. The proof there is informed by techniques from the proof of Proposition 5.2 in [Put12].

Theorem 3.1. [BFS+12, Theorem 3.9] *Let Γ_m be a subgraph of L_n with m edges. Then $\mathcal{MA}(\Gamma_m)$ is $(\eta(m + 1) - 1)$-connected.*

The function η is defined slightly differently in [BFS+12], hence the unusual-looking connectivity bound here. The point it that L_n has $n - 1$ edges, so in particular $\mathcal{MA}(L_n)$ is $(\eta(n) - 1)$-connected.

3.1. **Connectivity of arc complexes.** Our first goal is to deduce the connectivity of $\mathcal{RA}(L_n)$ from Theorem 3.1. Let Γ_m be a subgraph of L_n with m edges. For a k-simplex $\sigma = \{\alpha_0, \dots, \alpha_k\}$ in $\mathcal{RA}(\Gamma_m)$, define $r(\sigma)$ to be the number of points in P that are used as endpoints of arcs in σ. Then define the *defect* $d(\sigma)$ to be $2(k+1) - r(\sigma)$. Let h be the function on the barycentric subdivision $\mathcal{RA}(\Gamma_m)'$ of $\mathcal{RA}(\Gamma_m)$ given by $h(\sigma) = (d(\sigma), -\dim(\sigma))$, ordered lexicographically. Note that $d(\sigma) = 0$ if and only if the arcs are all disjoint, even at their endpoints. Hence, thinking of h as a height function on the vertices of $\mathcal{RA}(\Gamma_m)'$, in the sense of [BB97], we observe that the sublevel set $(\mathcal{RA}(\Gamma_m)')^{d=0}$ is precisely $\mathcal{MA}(\Gamma_m)'$. Hence we can compare the homotopy types of the two complexes using discrete Morse theory, with [BB97, Corollary 2.6] as the guide. The key is to inspect the descending links with respect to h. This is very similar to the procedure used in [BFS+12] to deduce connectivity of $\mathcal{MA}(K_n)$ from connectivity of $\mathcal{HA}(K_n)$.

Proposition 3.2. *$\mathcal{RA}(\Gamma_m)$ is $(\eta(m+1)-1)$-connected.*

Proof. We know that $\mathcal{MA}(\Gamma_m)$ is $(\eta(m+1)-1)$-connected by Theorem 3.1. We claim that the inclusion $\mathcal{MA}(\Gamma_m) \to \mathcal{RA}(\Gamma_m)$ induces a surjection in homotopy π_k for $k \le \eta(m+1)-1$, from which the proposition follows. To prove the claim, it suffices by [BB97, Corollary 2.6] to prove that for $\sigma \in \mathcal{RA}(\Gamma_m) \setminus \mathcal{MA}(\Gamma_m)$, i.e., $h(\sigma) > 0$, the descending link $\mathrm{lk}{\downarrow}(\sigma)$ is $(\eta(m+1)-2)$-connected. We suppose that σ is a k-simplex, with $\sigma = \{\alpha_0, \dots, \alpha_k\}$.

There are two types of arc systems in $\mathrm{lk}{\downarrow}(\sigma)$. First, we could have $\sigma' < \sigma$ and $h(\sigma') < h(\sigma)$. Then σ' is obtained from σ by removing arcs and strictly decreasing the defect. Call the full subcomplex of $\mathrm{lk}{\downarrow}(\sigma)$ spanned by these σ' the *down-link*. Second, we could have $\widetilde{\sigma} > \sigma$ and $h(\widetilde{\sigma}) < h(\sigma)$. Here $\widetilde{\sigma}$ is obtained by adding new arcs to σ, so that the new arcs are all disjoint from each other and from any existing arcs, even at endpoints. Call the full subcomplex of $\mathrm{lk}{\downarrow}(\sigma)$ spanned by such $\widetilde{\sigma}$ the *up-link*. Any simplex in the down-link is a face of every simplex in the up-link, so $\mathrm{lk}{\downarrow}(\sigma)$ is the join of the down-link and up-link.

First consider the down-link. A face σ' of σ fails to be in the down-link if and only if each arc in $\sigma \setminus \sigma'$ is disjoint from every other arc of σ, since then and only then do σ and σ' have the same defect. Let σ_0 be the face of σ consisting precisely of all such arcs, if any exist. Since $d(\sigma) > 0$, we know $\sigma_0 \neq \sigma$. The boundary of σ is a $(k-1)$-sphere, and the complement in the boundary of the down-link is either empty, or is a cone with cone point σ_0. Hence the down-link is either a $(k-1)$-sphere or is contractible, so in particular is $(k-2)$-connected. At this point we may assume without loss of generality that the down-link is a $(k-1)$-sphere, and so every arc in σ shares an endpoint with some other arc in σ. This means that every edge of Γ_σ shares an endpoint with some other edge of Γ_σ. In particular $k \ge 1$.

Now consider the up-link. The simplices in the up-link are given by adding arcs to σ that are all disjoint from each other and from the arcs in σ. Consider the connected surface $S' := S \setminus \{\alpha_0, \dots, \alpha_k\}$, obtained by cutting out the arcs α_i. If $P' := S' \cap P$, then $|P'| = n - r(\sigma)$. Also let Γ'_{m-2k-2} be the subgraph of Γ_m obtained by removing the edges of Γ_σ, and all edges sharing a vertex with any of these, so Γ'_{m-2k-2} has at most $m-2k-2$ edges (here we use the fact that every edge of Γ_σ shares an endpoint with some other edge of Γ_σ). The up-link of σ is isomorphic to the matching complex $\mathcal{MA}(S', P', \Gamma'_{m-2k-2})$, which is $(\eta(m-2k-1)-1)$-connected. Since $\mathrm{lk}{\downarrow}(\sigma)$ is the join of the down- and up-links, we conclude that $\mathrm{lk}{\downarrow}(\sigma)$ is $(\eta(m-2k-1)+k-1)$-connected.

We have $\eta(m-2k-1)+k-1 \ge \dfrac{m-2k-3}{4} + k - 2 \ge \eta(m+1) + \dfrac{k}{2} - \dfrac{5}{2} \ge \eta(m+1) - 2$ since $k \ge 1$, and so we are done. ☐

The next corollary is immediate, keeping in mind that with our notation L_n has $n-1$ edges.

Corollary 3.3. *$\mathcal{RA}(L_n)$ is $(\eta(n)-1)$-connected.* ☐

Corollary 3.4. *$\mathrm{CC}(PB_n, \mathscr{AF}_n)$ is $(\eta(n)-1)$-connected, and hence $\mathscr{AF}_n$ is $\eta(n)$-generating for PB_n.*

Proof. This is immediate from Lemma 2.6 and Corollary 3.3. ☐

We also want to show that the families $\mathscr{AF}_n^s$ from Definition 2.12 are highly generating. For $s > 1$, the coset complex $\mathrm{CC}(PB_n, \mathscr{AF}_n^s)$ is obtained up to homotopy equivalence from $\mathrm{CC}(PB_n, \mathscr{AF}_n^{s-1})$ by removing the open stars of the vertices, i.e., the cosets pPB_n^J for $|J| = s-1$. Hence the problem amounts to showing high connectivity of links. This is more or less the procedure done in the proof of Theorem 3.3 in [AH93], in the context of buildings. It is a bit harder here though; links in buildings are themselves buildings, but links in restricted arc complexes are not themselves restricted arc complexes. Nonetheless, we can get the right connectivity without too much extra work.

Lemma 3.5 (Links in $\mathcal{RA}(\Gamma_m)$)**.** *Let $\sigma = \{\alpha_0, \dots, \alpha_k\}$ be a k-simplex in $\mathcal{RA}(\Gamma_m)$ for Γ_m as above (with m edges). Then the link $\mathrm{lk}_{\mathcal{RA}(\Gamma_m)}(\sigma)$ is $(\eta(m-k)-1)$-connected.*

To make precise the terminology, here by "link" we mean the subcomplex of simplices τ disjoint from σ for which there exists a simplex with τ and σ as faces.

Proof. Set $L := \mathrm{lk}_{\mathcal{RA}(\Gamma_m)}(\sigma)$. An arc system τ is in L if and only if each arc of τ is distinct from, but compatible with, every α_i. For such a τ, by retracting each arc α_i to a point, τ maps to an arc system in $\mathcal{RA}(\Gamma_{m-(k+1)})$. Here $\Gamma_{m-(k+1)}$ is a subgraph of Γ_m with $m-(k+1)$ edges. More formally, for $0 \le d \le k$ consider the homotopy equivalence of surfaces $r_d \colon S \to S_d$, obtained by collapsing α_i to a point, for each $0 \le i \le d$. Recall $S = D_n$, and here S_d is just our name for the copy of $D_{n-(d+1)}$ obtained by collapsing these arcs. Here we do not think of D_n as a punctured disk, but rather as a disk with n distinguished points; hence r_d is really a homotopy equivalence. Also let P_d be the image of P under r_d. We have induced maps of complexes $R_d \colon L \to \mathcal{RA}(\Gamma_{m-(d+1)})$. Note that these maps are surjective, but not injective; see Figure 5 for an example of the non-injectivity. Note however that the connectivity of $\mathcal{RA}(\Gamma_{m-(k+1)})$ is precisely the connectivity we are trying to verify for L.

The r_d also induce epimorphisms $\phi_d \colon PB_n \to PB_{n-(d+1)}$, with kernels $K_d := \ker(\phi_d)$. Also declare K_{-1} to be the trivial subgroup. Note that $K_{-1} \le K_0 \le \cdots \le K_k$. Colloquially, the pure braids p in $K_d \setminus K_{d-1}$ are precisely those that do "twist" α_d but don't twist any α_i for $i > d$. For $p \in K_k$, define $D(p) := \min\{d+1 \mid p \in K_d\}$. We will call $D(p)$ the *deviation* of p; note that $D(p) = 0$ if and only if $p = \mathrm{id}$. Now fix a map $s_{\mathrm{id}} \colon S_k \to S$ with $s_{\mathrm{id}} \circ r_k$ homotopic to the identity. This essentially amounts to fixing a choice of how to "blow up" each arc α_i to get from S_k back to S. We get an induced map $\iota_{\mathrm{id}} \colon \mathcal{RA}(\Gamma_{m-(k+1)}) \to L$, with $R_k \circ \iota_{\mathrm{id}}$ equal to the identity on $\mathcal{RA}(\Gamma_{m-(k+1)})$. For each $p \in K_k$, set $\iota_p := p \circ \iota_{\mathrm{id}}$. These maps are all injective simplicial maps that can be thought of as different choices of how to blow up each α_i, and we see that $R_k \circ \iota_p$ is the identity for all p. Every arc system in L is the image of an arc system in $\mathcal{RA}(\Gamma_{m-(k+1)})$ under some ι_p, so $L = \bigcup_{p \in PB_n} \mathrm{Im}(\iota_p)$. Also, each $\mathrm{Im}(\iota_p)$ is isomorphic to $\mathcal{RA}(\Gamma_{m-(k+1)})$, and hence is an $(\eta(m-k)-1)$-connected subcomplex of L. We now need to glue these $\mathrm{Im}(\iota_p)$ together in a clever order, always along $(\eta(m-k)-2)$-connected relative links, from which we will deduce that L is $(\eta(m-k)-1)$-connected.

The measurement $D(p)$ provides such an order. For $0 \le d \le k$ let $L^d := \bigcup_{D(p) \le d} \mathrm{Im}(\iota_p)$. We claim that L^d is $(\eta(m-k)-1)$-connected for all d. The base case $d = 0$ is clear. For a given d, the intersection $\mathrm{Im}(\iota_p) \cap \mathrm{Im}(\iota_q)$ with $p \neq q$ and $D(p) = D(q) = d+1$ is contained in L^d. This is because p and q must twist the arc α_d differently, and so if β is an arc in $\mathrm{Im}(\iota_p) \cap \mathrm{Im}(\iota_q)$ then β cannot share endpoints with α_d. For this reason, we can build up from L^d to L^{d+1} by attaching the $\mathrm{Im}(\iota_p)$ with deviation $d+1$, in any order, and the relative links will always be in L^d. Now, for p with $D(p) = d+1$, we attach $\mathrm{Im}(\iota_p)$ to L^d along the intersection $\mathrm{Im}(\iota_p) \cap L^d$. This intersection consists precisely of those arc systems in $\mathrm{Im}(\iota_p)$ that do not use arcs sharing endpoints with α_d. Applying R_k (so retracting each α_i to a point), this gives us the subcomplex of $\mathcal{RA}(\Gamma_{m-(k+1)})$ whose arcs are disjoint from the endpoint obtained by collapsing α_d. But this is just $\mathcal{RA}(\Gamma')$ for Γ' a subgraph of $\Gamma_{m-(k+1)}$ with at most two fewer edges. This is $(\eta(m-k)-2)$-connected, and so we are done. □

FIGURE 5. Distinct arcs in the link of σ that map to the same arc under R_d.

Proposition 3.6. *For $s \in \mathbb{N}$, $\mathrm{CC}(PB_n, \mathscr{AF}_n^s)$ is $(\eta(n-(s-1))-1)$-connected, and hence $\mathscr{AF}_n^s$ is $(\eta(n-(s-1)))$-generating for PB_n.*

Proof. It suffices to show that for $|J| = s-1$, the link of PB_n^J in $\mathrm{CC}(PB_n, \mathscr{AF}_n^{s-1})$ is $(\eta(n-(s-1))-1)$-connected. Equivalently, we need the link of σ_J in $\mathcal{RA}(L_n)$ to be $(\eta(n-(s-1))-1)$-connected. Since σ_J is a $(|J|-1)$-simplex, its link is $(\eta(n-|J|)-1)$-connected by Lemma 3.5 (since L_n has $n-1$ edges), and since $|J| = s-1$, we conclude that indeed the link is $(\eta(n-(s-1))-1)$-connected. ☐

3.2. **Connectivity of braige complexes.** Now we inspect $\mathrm{CC}(PB_n, \mathscr{BF}_n)$, or more accurately $\mathcal{PB}_n(L_n)$. To pass from the world of arcs to the world of braiges, we will project the braiges onto arcs in the following way. For each $J \subseteq [n-1]$, let σ_J be the simplex of $\mathcal{MA}(L_n)$ defined before Lemma 2.6. Consider the action of PB_n on $\mathcal{RA}(L_n)$ as a right action, and define a map

$$\pi\colon \mathcal{PB}_n(L_n) \to \mathcal{RA}(L_n)$$
$$[(p,\Gamma)] \mapsto (\sigma_{J_\Gamma})p^{-1}$$

where J_Γ is as in Definition 2.9. We will use π to also denote the restrictions $\mathcal{EPB}_n(L_n) \to \mathcal{MA}(L_n)$, $\mathcal{PB}_n(\Gamma) \to \mathcal{RA}(\Gamma)$ and $\mathcal{EPB}_n(\Gamma) \to \mathcal{MA}(\Gamma)$ for Γ a subgraph of L_n. As in [BFS+12], think of π as the procedure of combing the braid straight and watching where the arcs get moved.

Proposition 3.7 (Braige connectivity from arc connectivity)**.** *For Γ_m be a subgraph of L_n with m edges, $\mathcal{EPB}_n(\Gamma_m)$ is $(\eta(m+1)-1)$-connected.*

Proof. By Theorem 3.1, $\mathcal{MA}(\Gamma_m)$ is $(\eta(m+1)-1)$-connected. Let $\sigma = \{\alpha_0, \dots, \alpha_k\}$ be a k-simplex in $\mathcal{MA}(\Gamma_m)$. The link $\mathrm{lk}(\sigma)$ of σ in $\mathcal{MA}(\Gamma_m)$ is isomorphic to $\mathcal{MA}(\Gamma')$ for Γ' a subgraph of Γ_m with at least $m-3(k+1)$ edges, so $\mathrm{lk}(\sigma)$ is $(\eta(m-3(k+1)+1)-1)$-connected, and hence $(\eta(m+1)-k-2)$-connected. It now suffices by [Qui78, Theorem 9.1] to prove that the fiber $\pi^{-1}(\sigma)$ is $(k-1)$-connected (here we treat a simplex as a closed cell). Indeed, we will prove that $\pi^{-1}(\sigma)$ is the join of the fibers $\pi^{-1}(\alpha_i)$ of the vertices α_i of σ. See also Proposition 4.3 in [BFS+12].

Let $\mathscr{JVF} := *_{i=0}^{k}\pi^{-1}(\alpha_i)$ be the join of the vertex fibers. Clearly $\pi^{-1}(\sigma) \subseteq \mathscr{JVF}$. Also, the 0-skeleton of $\mathscr{JVF}$ is contained in $\pi^{-1}(\sigma)$. Now suppose that the same is true of the r-skeleton for $r > 0$. An $(r+1)$-simplex in $\mathscr{JVF}$ is the join of a 0-simplex and an r-simplex, both of which are contained in $\pi^{-1}(\sigma)$. It now suffices to prove the following claim.

Claim: Let $[(p,E)]$ be a vertex in $\mathcal{EPB}_n(\Gamma_m)$, so $p \in PB_n$ and E is a one-edge subgraph of Γ_m. Let $[(q,\Gamma)]$ be a simplex in $\mathcal{EPB}_n(\Gamma_m)$ such that $\pi([(q,\Gamma)])$ does not contain $\pi([(p,E)])$ but does share a simplex with $\pi([(p,E)])$ in $\mathcal{MA}(\Gamma_m)$. Then $[(q,\Gamma)]$ shares a simplex with $[(p,E)]$ in $\mathcal{EPB}_n(\Gamma_m)$.

This hypothesis is rephrased in terms of arcs as: $(\Gamma)q^{-1}$ shares a simplex with $(E)p^{-1}$. By acting from the left with PB_n, we can assume without loss of generality that $p = \mathrm{id}$, so we have $\pi([(p,E)]) = E$. Let $\{\beta_0, \dots, \beta_\ell\} := (\Gamma)q^{-1}$, chosen so that E is disjoint from the β_i, even at endpoints (remember we are in $\mathcal{MA}(\Gamma_m)$, not just $\mathcal{RA}(\Gamma_m)$). This is possible by the hypothesis, and implies that the dangling equivalence class $[(q,\Gamma)]$ contains a representative in which the $(j+1)^{\text{st}}$ strand is a clone of the j^{th} strand, where j and $j+1$ are the endpoints of the edge of E. We can assume (q,Γ) itself is such a representative, in which case the dangling braige $[(q, \Gamma \cup E)]$ is a simplex of $\mathcal{EPB}_n(\Gamma_m)$ containing $[(q,\Gamma)]$ and $[(p,E)]$, proving the claim. ☐

It might be possible to mimic this proof using $\pi\colon \mathcal{PB}_n(\Gamma) \to \mathcal{RA}(\Gamma)$ instead, and get the connectivity of $\mathcal{PB}_n(L_n)$ right away, but the downside is that the fibers are not joins of vertex fibers. Hence one would have to do extra work to show that fibers have the right connectivity.

To calculate the connectivity of $\mathcal{PB}_n(\Gamma_m)$, we will use a similar procedure as for $\mathcal{RA}(\Gamma_m)$. Namely, we will build up from $\mathcal{EPB}_n(\Gamma_m)$ to $\mathcal{PB}_n(\Gamma_m)$ using discrete Morse theory. A k-simplex in $\mathcal{PB}_n(\Gamma_m)$ is a dangling equivalence class of a pair (p,Γ), for $p \in PB_n$ and Γ a subgraph of Γ_m with $k+1$ edges. Let $r(\Gamma)$ be the number of vertices that are endpoints of an edge in Γ. Then define the *defect* $d(p,\Gamma)$ to be $2(k+1) - r(\Gamma)$. Extend these definitions to the dangling equivalence classes, and observe that $\mathcal{EPB}_n(\Gamma_m)$ is the $d = 0$ sublevel set of $\mathcal{PB}_n(\Gamma_m)$. We now apply Morse theory, as before.

Proposition 3.8. *$\mathcal{PB}_n(\Gamma_m)$ is $(\eta(m+1)-1)$-connected.*

Proof. By Proposition 3.7, $\mathcal{EPB}_n(\Gamma_m)$ is $(\eta(m+1)-1)$-connected. Mimicking the proof of Proposition 3.2, it suffices to prove that for $\sigma \in \mathcal{PB}_n(\Gamma_m) \setminus \mathcal{EPB}_n(\Gamma_m)$, the descending link $\mathrm{lk}{\downarrow}(\sigma)$ is $(\eta(m+1)-2)$-connected. Let σ be such a k-simplex, say $\sigma = [(p,\Gamma)]$. The down-link is either S^{k-1}, or contractible if Γ has an isolated edge. Suppose there is no such isolated edge, so the down-link is S^{k-1}. Now, the up-link is obtained by dangling and then adding extra edges to the graph, such that the new edges are disjoint from Γ and from each other. Since Γ has no isolated edges, there are at most $2(k+1)$ edges of Γ_m that share an endpoint with an edge of Γ. Hence the up-link of σ is isomorphic to $\mathcal{EPB}_\ell(\Gamma_{m-2k-2})$ for some ℓ, which is $(\eta(m-2k-1)-1)$-connected. The calculation from the proof of Proposition 3.2 now tells us that $\mathrm{lk}{\downarrow}(\sigma)$ is $(\eta(m+1)-2)$-connected. □

Corollary 3.9. *$\mathcal{PB}_n(L_n)$ is $(\eta(n)-1)$-connected.* □

Corollary 3.10. *$\mathrm{CC}(PB_n, \mathscr{BF}_n)$ is $(\eta(n)-1)$-connected, and hence $\mathscr{BF}_n$ is $\eta(n)$-generating for PB_n.* □

Example 3.11. For $n \geq 6$, $\mathrm{CC}(PB_n, \mathscr{BF}_n)$ is connected, so PB_n has a generating set in which each generator features at least one cloned strand. Indeed, the standard generating set from Section 1.3.1 of [KT08] satisfies this property for $n \geq 6$, and fails for $n < 6$. For $n \geq 10$, $\mathrm{CC}(PB_n, \mathscr{BF}_n)$ is simply connected, so PB_n is 2-generated by $\mathscr{BF}_n$. Hence there exists a presentation for PB_n in which every generator features a cloned strand, and the relations all arise from relations in the subgroups of braids with a cloned strand. Again we note that the standard presentation works precisely in this range.

We conclude by showing that the families $\mathscr{BF}_n^s$ for $s \in \mathbb{N}$, defined in Definition 2.12, are highly generating as well. For $s > 1$, the coset complex $\mathrm{CC}(PB_n, \mathscr{BF}_n^s)$ is obtained up to homotopy equivalence from $\mathrm{CC}(PB_n, \mathscr{BF}_n^{s-1})$ by removing the open stars of vertices, i.e, cosets $pPB_n^{(J)}$ for $|J| = s-1$, just like in the arc case.

Lemma 3.12 (Links in $\mathcal{PB}_n(\Gamma_m)$)**.** *Let σ be a k-simplex in $\mathcal{PB}_n(\Gamma_m)$ for Γ_m as above (with m edges). Then the link $\mathrm{lk}_{\mathcal{PB}_n(\Gamma_m)}(\sigma)$ is $(\eta(m-k)-1)$-connected.*

Proof. Links in the braige case are nicer than links in the arc case, since they are actually isomorphic to smaller braige complexes. In the arc case, namely in the proof of Lemma 3.5, we related a given link to a smaller arc complex, via a map that was not an isomorphism. In the present case, we claim that $\mathrm{lk}_{\mathcal{PB}_n(\Gamma_m)}(\sigma)$ is just isomorphic to $\mathcal{PB}_{n-(k+1)}(\Gamma_{m-(k+1)})$, for $\Gamma_{m-(k+1)}$ a graph with $m-(k+1)$ edges, and then the connectivity result is immediate. Say $\sigma = [(p, \Gamma_{k+1})]$ for Γ_{k+1} a subgraph of Γ_m with $k+1$ edges. Let $L := \mathrm{lk}_{\mathcal{PB}_n(\Gamma_m)}(\sigma)$. The simplices in L are dangling braiges of the form $\tau = [(pq, \Gamma)]$, where $q \in PB_n^{(J_{\Gamma_{k+1}})}$ and Γ is a subgraph of Γ_m having no edges in common with Γ_{k+1}. The first condition ensures that τ and σ share a simplex, namely $[(pq, \Gamma \cup \Gamma_{k+1})]$, and the second condition ensures that τ and σ are disjoint. Acting from the left with PB_n, we can assume $p = \mathrm{id}$. We have a map $\phi \colon L \to \mathcal{PB}_{n-(k+1)}(\Gamma_{m-(k+1)})$, where $\Gamma_{m-(k+1)}$ is the graph with $n-(k+1)$ vertices that is obtained from Γ_m by retracting each edge of Γ_{k+1} to a point. The map ϕ sends $\tau = [(q, \Gamma)]$ to $[(q', \Gamma')]$, where Γ' is the image of Γ under the retraction $\Gamma_m \to \Gamma_{m-(k+1)}$, and q' is the preimage of q under the cloning map $\kappa_{J_{\Gamma_{k+1}}}$. See Figure 6 for an example. Since q' is uniquely determined by q, we have an inverse ϕ^{-1}, induced by the cloning map. (This is the essential difference from the arc case, that there is only one way to "blow up" a braige via cloning.) Since ϕ and ϕ^{-1} are of course simplicial maps, we conclude that ϕ is a simplicial isomorphism, and the result follows. □

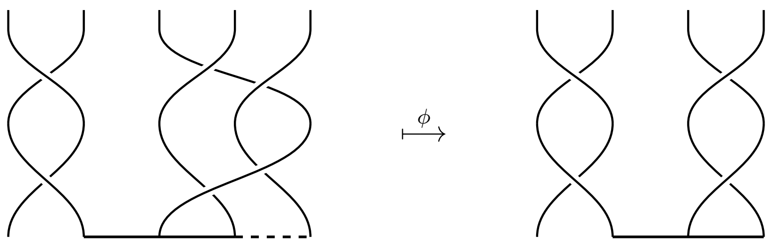

FIGURE 6. The map ϕ takes an element of $\mathrm{lk}_{\mathcal{PB}_5(L_5)}(\sigma)$ to an element of $\mathcal{PB}_4(L_4)$. Here σ is $[(\mathrm{id}, E_4)]$, for E_4 the subgraph with a single edge indicated by the dashed line.

Proposition 3.13. *For $s \in \mathbb{N}$, $\mathrm{CC}(PB_n, \mathscr{BF}_n^s)$ is $(\eta(n-(s-1))-1)$-connected, and hence $\mathscr{BF}_n^s$ is $\eta(n-(s-1))$-generating for PB_n.*

Proof. As in the proof of Proposition 3.6, it suffices to prove that for Γ with $s-1$ edges, the link of the $(s-2)$-simplex $[(\mathrm{id}, \Gamma)]$ in $\mathcal{PB}_n(L_n)$ is $(\eta(n-(s-1))-1)$-connected. Since L_n has $n-1$ edges, this follows from Lemma 3.12. □

Example 3.14. To generalize the previous example, we have that for any $n \geq 6$, $\mathscr{BF}_n^{n-5}$ is 1-generating for PB_n. This means that PB_n has a set of generators such that in each generator, all but 5 strands are clones (indeed the standard generators have this property). Similarly for $n \geq 10$, $\mathscr{BF}_n^{n-9}$ is 2-generating for PB_n, so PB_n has a presentation in which each relation can be realized by using only 9 non-clone strands. Again, the standard presentation fits the bill.

Example 3.15. In the situation of arcs, the *swing presentation* for PB_n, described in Section 4 of [MM09], provides an explicit example of $\mathscr{AF}_n^{n-5}$ being 1-generating for $n \geq 6$ and $\mathscr{AF}_n^{n-9}$ being 2-generating for $n \geq 10$. In this presentation the generators are Dehn twists, each of which must stabilize at least one arc of the form σ_j, as soon as $n \geq 6$. Each relation in [MM09, Theorem 4.10] (specifically the second presentation) is a product of Dehn twists, and for $n \geq 10$ this product stabilizes at least one arc of the form σ_j. See Figure 7 for an example.

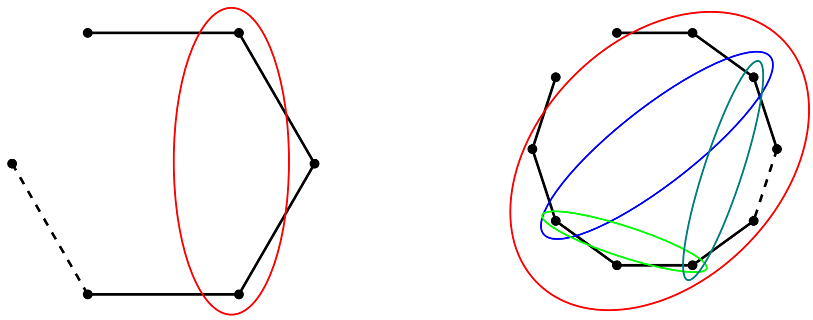

FIGURE 7. With 6 points, each generator must stabilize an arc. With 10 points, each relation must stabilize an arc. The dashed lines indicate the arcs stabilized in the examples. The relation pictured here is a *lantern relation*, as in Figure 12 of [MM09].

REFERENCES

[AB08] P. Abramenko and K. S. Brown. *Buildings*, volume 248 of *Graduate Texts in Mathematics*. Springer, New York, 2008. Theory and applications.

[AH93] H. Abels and S. Holz. Higher generation by subgroups. *J. Algebra*, 160(2):310–341, 1993.

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

[BFS+12] K.-U. Bux, M. Fluch, M. Schwandt, S. Witzel, and M. C. B. Zaremsky. The braided Thompson's groups are of type F_∞. arXiv:1210.2931, 2012.

[Bir74] J. S. Birman. *Braids, links, and mapping class groups*. Princeton University Press, Princeton, N. J., 1974. Annals of Mathematics Studies, No. 82.

[CD95] R. Charney and M. W. Davis. The $K(\pi, 1)$-problem for hyperplane complements associated to infinite reflection groups. *J. Amer. Math. Soc.*, 8(3):597–627, 1995.

[Hat91] A. Hatcher. On triangulations of surfaces. *Topology Appl.*, 40(2):189–194, 1991.

[KT08] C. Kassel and V. Turaev. *Braid groups*, volume 247 of *Graduate Texts in Mathematics*. Springer, New York, 2008. With the graphical assistance of Olivier Dodane.

[MM09] D. Margalit and J. McCammond. Geometric presentations for the pure braid group. *J. Knot Theory Ramifications*, 18(1):1–20, 2009.

[MMV98] J. Meier, H. Meinert, and L. VanWyk. Higher generation subgroup sets and the Σ-invariants of graph groups. *Comment. Math. Helv.*, 73(1):22–44, 1998.

[Put12] A. Putman. Representation stability, congruence subgroups, and mapping class groups. 2012. Version 2 of arXiv:1201.4876.

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

DEPARTMENT OF MATHEMATICS, BIELEFELD UNIVERSITY, PO BOX 100131, 33501 BIELEFELD, GERMANY
E-mail address: zaremsky@math.uni-bielefeld.de