Rev. Mat. Iberoam. 41 (2025), no. 5, 1711–1758
DOI 10.4171/RMI/1549

On the space of subgroups of Baumslag–Solitar groups I: Perfect kernel and phenotype

Alessandro Carderi, Damien Gaboriau, François Le Maître and Yves Stalder

Abstract. Given a Baumslag–Solitar group, we study its space of subgroups from a topological and dynamical perspective. We first determine its perfect kernel (the largest closed subset without isolated points). We then bring to light a natural partition of the space of subgroups into one closed subset and countably many open subsets that are invariant under the action by conjugation. One of our main results is that the restriction of the action to each piece is topologically transitive. This partition is described by an arithmetically defined function, that we call the phenotype, with values in the positive integers or infinity. We eventually study the closure of each open piece and also the closure of their union. We moreover identify in each phenotype a (the) maximal compact invariant subspace.

1. Introduction and presentation of the results

The Baumslag–Solitar group of non-zero integer parameters m and n is defined by the presentation

$$\mathrm{BS}(m,n) := \langle b, t \mid t b^m t^{-1} = b^n \rangle. \tag{1.1}$$

These one-relator two-generators groups were defined by Baumslag and Solitar [2] to provide examples of groups with surprising properties, depending on the arithmetic properties of the parameters.

It follows from the work of Baumslag and Solitar and of Meskin [24] that the group $\mathrm{BS}(m,n)$ is

- residually finite if and only if $|m| = 1$ or $|n| = 1$ or $|m| = |n|$,
- Hopfian if and only if it is residually finite or m and n have the same set of prime divisors.

The group $\mathrm{BS}(m,n)$ is amenable if and only if $|m| = 1$ or $|n| = 1$, and in this case, it is metabelian. All Baumslag–Solitar groups, however, share weak forms of amenability: they are inner-amenable [27] and a-T-menable [17].

Mathematics Subject Classification 2020: 20E06 (primary); 20E08, 20F65, 37B05 (secondary).
Keywords: Baumslag–Solitar groups, space of subgroups, perfect kernel, topologically transitive actions, Bass–Serre theory.

Over the years and despite the simplicity of their presentation, these groups have served as a standard source of examples and counter-examples, sometimes to published results (!). They have been considered from countless different perspectives in group theory and beyond.

Various aspects concerning the subgroups of the $\mathrm{BS}(m,n)$ have been considered such as the growth functions of their number of subgroups of finite index with various properties, or such as a description of the kind of fundamental group of graphs of groups that can be embedded as subgroups in some $\mathrm{BS}(m,n)$; see, for instance, [13, 18, 22].

In this article, we consider global aspects of the space $\mathrm{Sub}(\mathrm{BS}(m,n))$ of subgroups of the $\mathrm{BS}(m,n)$ and of the topological dynamics generated by the natural action by conjugation.

1.1. The perfect kernel

Let Γ be a countable group. We denote by $\mathrm{Sub}(\Gamma)$ the space of subgroups of Γ. If one identifies each subgroup with its indicator function, one can view the space $\mathrm{Sub}(\Gamma)$ as a closed subset of $\{0,1\}^\Gamma$. Thus, $\mathrm{Sub}(\Gamma)$ is a compact, metrizable space by giving it the restriction of the product topology. See Section 2.2 for the generalities about $\mathrm{Sub}(\Gamma)$.

By the Cantor–Bendixson theorem, $\mathrm{Sub}(\Gamma)$ admits a unique decomposition as a disjoint union of a perfect set, called the *perfect kernel* $\mathcal{K}(\Gamma)$ of Γ, and of a countable open subset. It is a challenging problem to determine the perfect kernel of a given countable group.

When Γ is finitely generated, the finite index subgroups are isolated in $\mathrm{Sub}(\Gamma)$. It is thus relevant to introduce the subspace $\mathrm{Sub}_{[\infty]}(\Gamma)$ consisting of all infinite index subgroups of Γ. It is a closed subspace of $\mathrm{Sub}(\Gamma)$ exactly when Γ is finitely generated (see Remark 2.3).

Our first main result completely describes the perfect kernel of the various Baumslag–Solitar groups. When $|m| = |n|$, the subgroup generated by b^m is normal; let us denote by π the corresponding quotient homomorphism

$$\mathrm{BS}(m,n) \xrightarrow{\pi} \mathrm{BS}(m,n)/\langle b^m \rangle.$$

We also denote by π the map it induces between the spaces of subgroups of $\mathrm{BS}(m,n)$ and $\mathrm{BS}(m,n)/\langle b^m \rangle$.

Theorem A (Perfect kernel of $\mathrm{BS}(m,n)$, Theorem 5.3). *Let $m,n \in \mathbb{Z} \setminus \{0\}$.*

(1) *If $|m| = 1$ or $|n| = 1$, then $\mathcal{K}(\mathrm{BS}(m,n))$ is empty.*

(2) *If $|m|, |n| > 1$, then*

  (a) *if $|m| \neq |n|$, then $\mathcal{K}(\mathrm{BS}(m,n)) = \mathrm{Sub}_{[\infty]}(\mathrm{BS}(m,n))$,*

  (b) *if $|m| = |n|$, then $\mathcal{K}(\mathrm{BS}(m,n)) = \pi^{-1}(\mathrm{Sub}_{[\infty]}(\mathrm{BS}(m,n)/\langle b^m \rangle))$.*

The fact that $\mathrm{Sub}(\mathrm{BS}(m,n))$ is countable when $|m| = 1$ or $|n| = 1$ (item (1)), i.e., for the Baumslag–Solitar groups that are metabelian, was already observed by Becker, Lubotzky and Thom, Corollary 8.4 in [3]. Fortuitously or not, it turns out that the equality $\mathcal{K}(\mathrm{BS}(m,n)) = \mathrm{Sub}_{[\infty]}(\mathrm{BS}(m,n))$ holds exactly when $\mathrm{BS}(m,n)$ is not residually finite.

There is a general correspondence between the transitive pointed Γ-actions and the subgroups of Γ. It sends an action α to the stabilizer of its base point. This Γ-equivariant map is a bijection when one considers the actions up to pointed isomorphisms (see Section 2.2). Item (2) of Theorem A has a unified reformulation in this setting:

(2′) *If* $|m|, |n| > 1$, *then* $\mathcal{K}(\mathrm{BS}(m,n))$ *is the space of subgroups* Λ *such that the right* $\mathrm{BS}(m,n)$*-action on* $\Lambda\backslash\mathrm{BS}(m,n)$ *has infinitely many* $\langle b\rangle$*-orbits.*

Note that this exactly means that the quotient of the Λ-action on the standard Bass–Serre tree (see Section 2.3) of $\mathrm{BS}(m,n)$ is infinite.

Let us now give some more context for Theorem A. By Brouwer's characterization of Cantor spaces, the space $\mathcal{K}(\Gamma)$ is either empty or a Cantor space. It is empty exactly when $\mathrm{Sub}(\Gamma)$ is countable. This happens, for example, for groups all whose subgroups are finitely generated, also known as Noetherian groups. For instance, all finitely generated nilpotent groups and more generally all polycyclic groups have a countable space of subgroups.

On the opposite side, for the free group with a countably infinite number of generators, no subgroup is isolated, thus $\mathcal{K}(\mathbf{F}_\infty) = \mathrm{Sub}(\mathbf{F}_\infty)$ (see Proposition 2.1 in [11]).

There are some classical groups for which we know that $\mathcal{K}(\Gamma) = \mathrm{Sub}_{[\infty]}(\Gamma)$. This is the case for the free groups $\mathbf{F}_n$ (for $1 < n < \infty$), see, for instance, Proposition 2.1 of [11]. This is also the case for the groups with infinitely many ends, for the fundamental groups of the closed surfaces of genus ≥ 2, and for the finitely generated LERF groups with non-zero first ℓ^2-Betti number (see [1]). Recall that a group Γ is LERF when its set of finite index subgroups is dense in $\mathrm{Sub}(\Gamma)$ (see, for instance, Theorem 3.1 in [19]).

Bowen, Grigorchuk and Kravchenko established that the perfect kernel of the lamplighter group $(\mathbb{Z}/p\mathbb{Z})^n \wr \mathbb{Z} = (\bigoplus_{\mathbb{Z}}(\mathbb{Z}/p\mathbb{Z})^n) \rtimes \mathbb{Z}$ (for a prime number p) is exactly the space $\mathrm{Sub}(\bigoplus_{\mathbb{Z}}(\mathbb{Z}/p\mathbb{Z})^n)$ of subgroups of the normal subgroup, Theorem 1.1 in [7]. Skipper and Wesolek uncovered the perfect kernel for a class of branch groups containing the Grigorchuk group and the Gupta–Sidki 3 group [26].

The perfect kernel can be obtained by successively, and transfinitely, removing the isolated points, thus obtaining for every ordinal α the α-th Cantor–Bendixson derivative $\mathrm{Sub}(\Gamma)^{(\alpha)} := \mathrm{Sub}(\Gamma)^{(\beta)} \setminus \{\text{isolated points of } \mathrm{Sub}(\Gamma)^{(\beta)}\}$ if $\alpha = \beta + 1$, and $\mathrm{Sub}(\Gamma)^{(\alpha)} := \bigcap_{\beta<\alpha} \mathrm{Sub}(\Gamma)^{(\beta)}$ if α is a limit ordinal. The *Cantor–Bendixson rank* $\mathrm{rk}_{\mathrm{CB}}(\Gamma)$ of Γ is the first ordinal ζ for which the derived space $\mathrm{Sub}(\Gamma)^{(\zeta)}$ has no more isolated points, and is thus equal to the perfect kernel (see, for instance, Section 6.C of [20] for details). When $|m|, |n| > 1$ and $|m| \neq |n|$, then Theorem A implies that $\mathrm{rk}_{\mathrm{CB}}(\mathrm{BS}(m,n)) = 1$. The determination of the Cantor–Bendixson ranks $\mathrm{rk}_{\mathrm{CB}}(\mathrm{BS}(m,n))$ for the other cases is postponed to the sequel [6].

1.2. Dynamical partition of the perfect kernel

The compact space of subgroups $\mathrm{Sub}(\Gamma)$ is equipped with the continuous action of Γ by conjugation: $\Lambda \cdot \gamma := \gamma^{-1}\Lambda\gamma$. The perfect kernel is Γ-invariant. This action is of course not minimal in general, even when restricted to the perfect kernel: the latter may contain normal subgroups and these subgroups are fixed points! However, the first three named authors observed a particularly nice feature in the case of the free group $\mathbf{F}_n$ (for $1<n<\infty$): the action $\mathcal{K}(\mathbf{F}_n) \curvearrowleft \mathbf{F}_n$ is topologically transitive (which means that the space admits a dense G_δ subset of points whose individual orbits are dense). These $\mathbf{F}_n$-actions are called totipotent, see [11].

To our surprise, we uncovered a dramatically different and very rich situation for the Baumslag–Solitar groups.

Theorem B. *Whenever* $|m|, |n| \neq 1$*, the perfect kernel* $\mathcal{K}(\mathrm{BS}(m,n))$ *admits a countably infinite partition into* $\mathrm{BS}(m,n)$*-invariant and topologically transitive subspaces. For the induced topology on* $\mathcal{K}(\mathrm{BS}(m,n))$*), one of the subspaces is closed; all the other ones are open.*

Theorem B follows from Proposition 5.8 and Theorem 5.14. In a further work [16], we show that topological transitivity can be upgraded to high topological transitivity.

From now on in this introduction, we stick to the case $|m| \neq 1$ and $|n| \neq 1$. In order to describe the partition in Theorem B, we introduce a new invariant: the *phenotype*.

The relation $tb^m b^{-1} = b^n$ imposes some arithmetic conditions between the cardinalities of the b-orbit of a point x and the b-orbit of xt. For instance, the b-orbit of x is infinite if and only if the b-orbit of xt is infinite.

In Definition 4.1, we introduce a function $\mathrm{Ph}_{m,n}\colon \mathbb{Z}_{\geq 1} \cup \{\infty\} \to \mathbb{Z}_{\geq 1} \cup \{\infty\}$, called the (m,n)-phenotype, with the following property, which directly follows from Proposition 4.6, Theorem 4.13 and Proposition 3.22.

Theorem C. *Whenever* $|m|, |n| \neq 1$*, there is a transitive* $\mathrm{BS}(m,n)$*-action with two* b*-orbits of cardinal* k *and* ℓ*, respectively, if and only if* $\mathrm{Ph}_{m,n}(k) = \mathrm{Ph}_{m,n}(\ell)$.

If, for instance, m and n are coprime, the phenotype $\mathrm{Ph}_{m,n}(k)$ of any natural number $k \in \mathbb{Z}_{\geq 1}$ is obtained as k expunged of all its prime divisors that appear in either m or n. The general form is more complicated, see Definition 4.1 and Example 4.3, but it follows readily from Definition 4.1 that $\mathrm{Ph}_{m,n}(q) = q$ for every $q \geq 1$ that is coprime with m and n. Hence, the set of possible (m,n)-phenotypes

$$\mathcal{Q}_{m,n} := \{\mathrm{Ph}_{m,n}(k) : k \in \mathbb{Z}_{\geq 1}\} \cup \{\infty\}.$$

is always infinite.

Theorem C allows us to define the *phenotype* of a transitive $\mathrm{BS}(m,n)$-action as the common (m,n)-phenotype of the cardinalities of its b-orbits. Then we define the *phenotype* $\mathbf{Ph}(\Lambda)$ of a subgroup $\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n))$ as the phenotype of the (right) $\mathrm{BS}(m,n)$-action on the homogeneous space $\Lambda\backslash\mathrm{BS}(m,n)$.

Notice that the $\mathrm{BS}(m,n)$-actions on $\Lambda\backslash\mathrm{BS}(m,n)$ and $(g^{-1}\Lambda g)\backslash\mathrm{BS}(m,n)$ are isomorphic (both are transitive with some point stabilizer equal to Λ), so they have the same phenotype. Hence, the partition

$$\mathrm{Sub}(\mathrm{BS}(m,n)) = \bigsqcup_{q\in\mathcal{Q}_{m,n}} \mathbf{Ph}^{-1}(q) \tag{1.2}$$

is invariant under the $\mathrm{BS}(m,n)$-action (recall this is the action by conjugation). Let us mention from Proposition 5.8 that

- for each finite $q \in \mathcal{Q}_{m,n}$, the piece $\mathbf{Ph}^{-1}(q)$ is open; it is also closed if and only if $|m| = |n|$;
- the piece $\mathbf{Ph}^{-1}(\infty)$ is closed but not open.

In particular, the function $\mathbf{Ph}\colon \mathrm{Sub}(\mathrm{BS}(m,n)) \to \mathbb{Z}_{\geq 1} \cup \{+\infty\}$ is Borel. It is continuous if and only if $|m| = |n|$.

It now follows from Theorem 5.14 that the restriction of the partition (1.2) to the perfect kernel

$$\mathcal{K}(\mathrm{BS}(m,n)) = \bigsqcup_{q \in \mathcal{Q}_{m,n}} \mathcal{K}_q(\mathrm{BS}(m,n)), \tag{1.3}$$

where $\mathcal{K}_q(\mathrm{BS}(m,n)) := \mathcal{K}(\mathrm{BS}(m,n)) \cap \mathbf{Ph}^{-1}(q)$, satisfies all the conclusions of Theorem B. The pieces $\mathcal{K}_q(\mathrm{BS}(m,n))$ are indeed non-empty, see Remark 5.12.

1.3. Approximations by subgroups of other phenotypes

We still stick to the case $|m| \neq 1$ and $|n| \neq 1$. Since the only non-open piece in partition (1.2) is $\mathbf{Ph}^{-1}(\infty)$, the subgroups of infinite phenotype are the only ones which can be approximated in $\mathrm{Sub}(\mathrm{BS}(m,n))$ by subgroups of other (that is, finite) phenotypes.

The set of limits of subgroups of finite phenotype depends on whether we fix the phenotype or we let it vary. About approximations by subgroups with a constant phenotype, we have the following result (see Proposition 5.8 and Theorem 6.2).

Theorem D. *Assume $|m|, |n| \neq 1$ and let us fix a finite (m,n)-phenotype q.*

1. *If $|m| = |n|$, then $\mathbf{Ph}^{-1}(q)$ is closed, hence no infinite phenotype subgroup can be approximated by subgroups of phenotype q.*
2. *If $|m| \neq |n|$, then an infinite phenotype subgroup Λ can be approximated by subgroups of phenotype q if and only if $\Lambda \leq \langle\langle b \rangle\rangle$, where $\langle\langle b \rangle\rangle$ is the normal subgroup generated by b.*

It is remarkable that the set $\overline{\mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty)$ is independent of q in the previous result.

Allowing the finite phenotype to vary yields new limit points. Our result is the following (see Proposition 6.7 and Corollary 6.11).

Theorem E. *Assume $|m|, |n| \neq 1$.*

1. *If $|m| = |n|$, then every infinite phenotype subgroup is a limit of finite (and varying) phenotype subgroups.*
2. *On the contrary, if $|m| \neq |n|$, then the set of subgroups in $\mathbf{Ph}^{-1}(\infty)$ which are limits of finite (and varying) phenotypes subgroups has empty interior in $\mathbf{Ph}^{-1}(\infty)$.*

Therefore, in the case $|m| = |n|$, all subgroups of infinite phenotype are limits of subgroups of finite phenotype, but none of them is a limit of subgroups of fixed finite phenotype.

The case $|m| \neq |n|$ is more complex. We do not have a nice description of the limit set from the above theorem. We can show however that this limit set is strictly larger than its fixed phenotype counterpart, see Proposition 6.12 and Theorem 6.14.

1.4. Closures of orbits in finite phenotype

We still stick to the case $|m| \neq 1$, $|n| \neq 1$, and assume moreover $|m| \neq |n|$. The previous subsection shows that for any finite phenotype q, we have

$$\mathbf{Ph}^{-1}(q) \subsetneq \overline{\mathbf{Ph}^{-1}(q)} \subsetneq \mathbf{Ph}^{-1}(q) \cup \mathbf{Ph}^{-1}(\infty).$$

Theorem B further yields that $\mathbf{Ph}^{-1}(q)$ contains dense orbits. For such an orbit $\mathcal{O}$, one has $\overline{\mathcal{O}} = \overline{\mathbf{Ph}^{-1}(q)}$, thus $\overline{\mathcal{O}}$ intersects $\mathbf{Ph}^{-1}(\infty)$. In fact, Theorem D completely described $\overline{\mathcal{O}}$. We now turn our attention to the orbits whose closure is contained in $\mathbf{Ph}^{-1}(q)$. Quite remarkably, they form a compact set.

Theorem F (See Theorem 5.20). *Suppose $|m|, |n| \neq 1$ and $|m| \neq |n|$. For every finite phenotype q, there is a positive integer $s = s(q, m, n)$ such that the subset*

$$\mathcal{MC}_q := \mathbf{Ph}^{-1}(q) \cap \{\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n)) \colon \Lambda \geq \langle\!\langle b^s \rangle\!\rangle\}$$

is compact and contains all the invariant compact subsets of $\mathbf{Ph}^{-1}(q)$.

In particular, every normal subgroup of phenotype q, and hence every finite index subgroup, contains $\langle\!\langle b^s \rangle\!\rangle$. Moreover, $\mathcal{MC}_q \cap \mathcal{K}_q(\mathrm{BS}(m,n))$ has empty interior in the corresponding piece of the perfect kernel $\mathcal{K}_q(\mathrm{BS}(m,n))$ (Theorem 5.20 (4)).

When $\gcd(m,n) = 1$, the above theorem takes an easier form, that is, $s = q$ and $\mathcal{MC}_q \cap \mathcal{K}(\mathrm{BS}(m,n)) = \{\langle\!\langle b^q \rangle\!\rangle\}$. In particular, $\langle\!\langle b^q \rangle\!\rangle$ is the unique normal subgroup of phenotype q and infinite index, see Theorem 5.20 (5). On the other hand, if $\gcd(m,n) \neq 1$, then the perfect kernel contains continuum many normal subgroups of phenotype q, see Theorem 5.26.

1.5. An example: The case of BS(2, 3)

Let us specialize our theorems to the case of BS(2, 3). An illustrative picture is given in Figure 1.

Since $2 \neq 3$, Theorem A tells us that $\mathcal{K}(\mathrm{BS}(2,3)) = \mathrm{Sub}_{[\infty]}(\mathrm{BS}(2,3))$. In this case the phenotype is given by the following simple formula:

$$\mathbf{Ph}(\Lambda) = \frac{I}{2^{|I|_2}\, 3^{|I|_3}},$$

where I is the index $I := [\langle b\rangle : \Lambda \cap \langle b\rangle]$, and where $|I|_p$ denotes the p-adic valuation of I subject to the convention that $|\infty|_p = 0$.

Therefore, the possible phenotypes for subgroups of BS(2, 3) are given by all the positive integers not divisible by 2 and 3, and infinity. Denoting $\mathcal{K}_q = \{\Lambda \leq \mathrm{BS}(2,3) : \mathbf{Ph}(\Lambda) = q\}$, the partition (1.3) becomes

$$\mathcal{K}(\mathrm{BS}(2,3)) = \mathcal{K}_\infty \sqcup \bigsqcup_{q \,:\, \gcd(q,2)=\gcd(q,3)=1} \mathcal{K}_q.$$

By Theorem B, the action on each $\mathcal{K}_q$ is topologically transitive. Note that all finite index subgroups have finite phenotype. The set $\mathcal{K}_\infty$ is closed and colored in black in Figure 1; the subsets $\mathcal{K}_q$ are open and colored in gray in the figure. Finally, the finite index subgroups are denoted by the dotted lines. Note that there are infinitely many finite index subgroups and they accumulate on the sets $\mathcal{K}_q$.

Note that for every finite q, the set $\overline{\mathcal{K}_q} \cap \mathcal{K}_\infty$ is non-empty and independent of q; indeed, by Theorem D, this is the set of subgroups of infinite phenotype contained in $\langle\!\langle b \rangle\!\rangle$. This set is illustrated as the black central disk in the figure. As one can guess in the

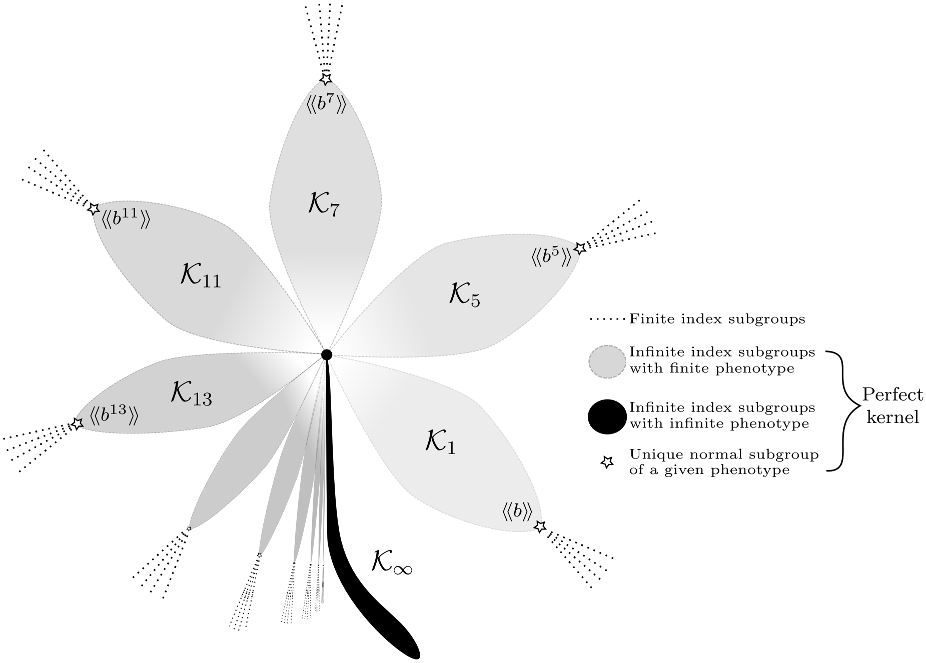

Figure 1. The space of subgroups of $\mathrm{BS}(2, 3)$.

figure, $\overline{\bigcup_{q \text{ finite}} \mathcal{K}_q} \cap \mathcal{K}_\infty$ is strictly bigger than this set, and yet not the entirety of $\mathcal{K}_\infty$, as prescribed by Theorem E.

We finally apply Theorem F. Since $\gcd(2, 3) = 1$, for every finite phenotype q, the largest compact invariant subset of $\mathcal{K}_q$ consists only of one point: the unique normal subgroup contained in $\mathcal{K}_q$, namely, $\langle\langle b^q \rangle\rangle$, pictured with a small star in the figure. Moreover, $\mathcal{MC}_q$ consists of the finite index subgroups of phenotype q represented by the dotted lines emanating from the star together with the single accumulation point $\langle\langle b^q \rangle\rangle$ of $\mathcal{MC}_q$.

Remark. Figure 1 is actually quite general. As soon as $|m| \neq |n|$, we have the exact same picture except that the possible phenotypes are different, and the stars turn into bigger compact maximal invariant subsets. Moreover, the phenotype is given by a more complicated formula.

1.6. Some ideas on the techniques of proofs

The definition of the topology on the space of subgroups leads us to look at the restriction of transitive actions to some part of their Schreier graph and then on assembling such parts from different actions (to form new actions). This leads us to the notion of pre-action, as considered in [14], where to facilitate the verification of the group relation, we impose that b is defined everywhere, i.e., on the whole domain of the pre-action (see Section 3.1). These pre-actions are more malleable but the algebraic conditions underlying them still make them difficult to manipulate.

This is why we further downgrade the data and move on to purely combinatorial objects associated with actions and pre-actions: the (m,n)-graphs (Section 3.3). These are oriented graphs which carry labels on the vertices and on the edges and which satisfy simple arithmetic conditions linking degrees and labels (Definition 3.12, equalities (3.13) and inequalities (3.14)). They generalize the Bass–Serre graphs of pre-actions used in [14] by adding their labels which record the size of the orbits of b, b^m or b^n according to the graph element considered. Notice that in [14] the b-orbits were assumed to be infinite.

All the vertex labels of a connected (m,n)-graph have the same (m,n)-phenotype (Proposition 4.6), which is thus defined to be the phenotype of the graph (Definition 4.8).

At this level, we can consider assembling together different parts (originating from different actions). Consider two connected (m,n)-graphs that are non-saturated (at least one of the inequalities (3.14) is strict). Then they can appear as subgraphs of the same (m,n)-graph as soon as they have the same phenotype (Theorem 4.13). This relies on two basic constructions, the Welding Lemma 4.16 and the connecting Theorem 4.17.

We then proceed by upgrading from (m,n)-graphs to pre-actions and actions (Proposition 3.23). These upgrades are not uniquely determined, however, if an (m,n)-graph $\mathcal{G}_2$ contains the (m,n)-graph $\mathcal{G}_1$ of a pre-action α_1, then the upgraded pre-action α_2 can be chosen to extend α_1 (Proposition 3.23).

To summarize, we will use several times the same construction scheme: Considering two actions, we restrict them to a large but proper part of their domain (pre-actions). We downgrade the resulting pre-actions to (m,n)-graphs and glue them together. We saturate the resulting (m,n)-graph and upgrade it into one action that "contains" the chosen parts of both original actions as sub-pre-actions (Theorem 4.12).

1.7. Subsequent work

Since the first version of the present paper appeared, two preprints have enriched the picture as follows.

On the one hand, the three last-named authors proved in [16] that the dynamics on the pieces $\mathcal{K}_q$ is in fact *highly topologically transitive*. They also studied the property of high transitivity for transitive actions of $\mathrm{BS}(m,n)$: They characterized the pieces containing subgroups Λ such that the action $\Lambda\backslash\mathrm{BS}(m,n) \curvearrowleft \mathrm{BS}(m,n)$ is highly transitive and they established that this property is generic in these pieces.

On the other hand, Sasha Bontemps has extended Theorems A, B and C to generalized Baumslag–Solitar groups, where the right notion of phenotype is more subtle [5]. She also obtained high topological transitivity results generalizing Theorem C from the aforementioned preprint [16].

2. Preliminaries and notations

In this text, we denote by $\mathbb{Z}_{\geq 0} := \{0,1,2,\dots\}$ the set of non-negative integers and by $\mathbb{Z}_{\geq 1} := \{1,2,3,\dots\}$ the set of positive integers. Given two integers $k,l \in \mathbb{Z}\setminus\{0\}$, we denote by $\gcd(k,l) \in \mathbb{Z}_{\geq 1}$ the *greatest common divisor* of k and l. We use the convention that $\gcd(k,\infty) = k$ and $\infty/k = k\infty = \infty$.

Let $\mathcal{P}$ be the set of prime numbers. Given an integer $k \in \mathbb{Z} \setminus \{0\}$ and a prime $p \in \mathcal{P}$, we denote by $|k|_p$ the *p-adic valuation* of k, that is, $|k|_p$ is the largest positive integer such that $p^{|k|_p}$ divides k.

2.1. Graphs and Schreier graphs

All our graphs are defined as in [25]. That is, a graph $\mathcal{G}$ is a couple $(V(\mathcal{G}), E(\mathcal{G}))$, where $V(\mathcal{G})$ is the *vertex set* and $E(\mathcal{G})$ is the *edge set*, endowed with:

- two maps $\mathtt{s}, \mathtt{t}\colon E(\mathcal{G}) \to V(\mathcal{G})$ called *source* and *target*, respectively,
- a fixed-point-free involution $E(\mathcal{G}) \to E(\mathcal{G}), e \mapsto \bar{e}$,

such that $\mathtt{s}(\bar{e}) = \mathtt{t}(e)$ and $\mathtt{t}(\bar{e}) = \mathtt{s}(e)$.

An *orientation* of the graph $\mathcal{G}$ is a partition $E(\mathcal{G}) = E^+(\mathcal{G}) \sqcup E^-(\mathcal{G})$ whose pieces are exchanged by the involution $e \mapsto \bar{e}$. Edges in $E^+(\mathcal{G})$ are called *positive* edges and edges in $E^-(\mathcal{G})$ are *negative*.

Remark 2.1. In order to define an oriented graph $\mathcal{G}$, it is enough to define the set of vertices $V(\mathcal{G})$, the set of positive edges $E^+(\mathcal{G})$, and the restrictions of the source and target maps $\mathtt{s}, \mathtt{t}$ to $E^+(\mathcal{G})$. Indeed, we can define $E^-(\mathcal{G})$ to be a copy of $E^+(\mathcal{G})$ and the involution $e \mapsto \bar{e}$ to be the natural identification of $E^+(\mathcal{G})$ with $E^-(\mathcal{G})$. We extend the source and target map by setting $\mathtt{s}(\bar{e}) := \mathtt{t}(e)$ and $\mathtt{t}(\bar{e}) := \mathtt{s}(e)$.

The *degree* of a vertex v in a graph $\mathcal{G}$, is the cardinal

$$\deg(v) := |\{e \in E(\mathcal{G}) : \mathtt{s}(e) = v\}| = |\{e \in E(\mathcal{G}) : \mathtt{t}(e) = v\}|.$$

If $\mathcal{G}$ is oriented, we say that an edge e is:

- a v*-outgoing* edge if it is positive and $\mathtt{s}(e) = v$,
- a v*-incoming* edge if it is positive and $\mathtt{t}(e) = v$.

The *outgoing degree* $\deg_{\text{out}}(v)$ of v is the number of v-outgoing edges, while its *incoming degree* $\deg_{\text{in}}(v)$ is the number of v-incoming edges. We clearly have $\deg_{\text{out}}(v) + \deg_{\text{in}}(v) = \deg(v)$.

A *subgraph* $\mathcal{G}'$ of a graph $\mathcal{G}$ is a graph such that $V(\mathcal{G}') \subseteq V(\mathcal{G})$, $E(\mathcal{G}') \subseteq E(\mathcal{G})$ and the structural maps of $\mathcal{G}'$ are restrictions of those of $\mathcal{G}$.

Still following [25], we call *circuit* a subgraph isomorphic to a circular graph (of length $l \geq 1$) and *loop* a circuit of length 1. The edge of a loop is also called a loop.

A *path* in a graph $\mathcal{G}$ is a finite sequence of edges $(e_1, \dots, e_n)$ such that $\mathtt{t}(e_k) = \mathtt{s}(e_{k+1})$ for all $1 \leq k \leq n-1$. Similarly, an *infinite path* is a sequence of edges $(e_k)_{k \geq 1}$ such that $\mathtt{t}(e_k) = \mathtt{s}(e_{k+1})$ for all $k \geq 1$. Finally, a (possibly infinite) path is called *simple* when the induced sequence of vertices is injective.

The *ball* $B(v, R)$ of radius R centered at a vertex v in a graph $\mathcal{G}$ is the subgraph induced by the set of vertices of $\mathcal{G}$ at distance $\leq R$ from v in the path metric.

Schreier graphs. Let Γ be a group and let S be a generating set of Γ. Consider a (right) action $\alpha\colon X \curvearrowleft \Gamma$. The *Schreier graph* of α relatively to S is the oriented graph $\mathbf{Sch}(\alpha) = \mathbf{Sch}(\alpha, S)$ defined by

$$V(\mathbf{Sch}(\alpha)) := X \quad \text{and} \quad E^+(\mathbf{Sch}(\alpha)) := \{(x, s) : x \in X, s \in S\},$$

where $\mathtt{s}(x, s) = x$ and $\mathtt{t}(x, s) = xs$, together with the following labeling: the edge (x, s) is labeled s and its opposite $\overline{(x, s)}$ is labeled by s^{-1}.

Given a subgroup $\Lambda \leq \Gamma$, we denote by $\mathbf{Sch}(\Lambda, S)$ the Schreier graph of the natural action $\Lambda\backslash\Gamma \curvearrowleft \Gamma$.

The *Cayley graph* of Γ relatively to S is the Schreier graph $\mathbf{Sch}(\alpha, S)$ of the action $\alpha\colon \Gamma \curvearrowleft \Gamma$ by (right) translations. This graph is denoted by $\mathbf{Cay}(\Gamma, S)$ and we clearly have $\mathbf{Cay}(\Gamma, S) = \mathbf{Sch}(\{\mathrm{id}\}, S)$. The Γ-action by left translations extends to the standard left action of Γ on $\mathbf{Cay}(\Gamma, S)$ by graph automorphisms.[1] In particular, $\Lambda\backslash\mathbf{Cay}(\Gamma, S) = \mathbf{Sch}(\Lambda, S)$.

Let $\varphi\colon X \to Y$ be a Γ-equivariant map from $\alpha\colon X \curvearrowleft \Gamma$ to $\beta\colon Y \curvearrowleft \Gamma$ and let S be a generating set of Γ. The map φ extends to a graph morphism from $\mathbf{Sch}(\alpha, S)$ to $\mathbf{Sch}(\beta, S)$ which respects the labelings. In particular, given subgroups $\Lambda_1 \leq \Lambda_2 \leq \Gamma$, the equivariant map $\Lambda_1\backslash\Gamma \to \Lambda_2\backslash\Gamma$ defines a surjective morphism $\mathbf{Sch}(\Lambda_1, S) \to \mathbf{Sch}(\Lambda_2, S)$.

2.2. Space of subgroups

Let Γ be a countable group. We identify its set of subsets with $\{0, 1\}^\Gamma$ and we endow it with the product topology, thus turning it into a Polish compact space. The *space of subgroups* of Γ is the closed, hence compact Polish, subspace

$$\mathrm{Sub}(\Gamma) := \{\Lambda \in \{0, 1\}^\Gamma : \Lambda \text{ is a subgroup}\},$$

which is also totally disconnected. The clopen subsets

$$\mathcal{V}(I, O) := \{\Lambda \in \mathrm{Sub}(\Gamma) : I \subseteq \Lambda \text{ and } \Lambda \cap O = \varnothing\}$$

of $\mathrm{Sub}(\Gamma)$, where I, O run over finite subsets of Γ, form a basis of the topology. Note that a sequence $(\Lambda_n)_{n\geq 0}$ of subgroups converges to Λ if and only if for all $\gamma \in \Gamma$,

$$(\gamma \in \Lambda) \iff (\gamma \in \Lambda_i \text{ for } i \text{ large enough}).$$

By the Cantor–Bendixson theorem [4, 10] (see, e.g., Theorem 6.4 in [20]), there is a unique decomposition

$$\mathrm{Sub}(\Gamma) = \mathcal{C}(\Gamma) \sqcup \mathcal{K}(\Gamma),$$

where $\mathcal{C}(\Gamma)$ is a countable open subset and $\mathcal{K}(\Gamma)$ is a closed perfect[2] subspace called the *perfect kernel* of Γ. The set $\mathcal{K}(\Gamma)$ is the largest subset $\mathcal{K} \subseteq \mathrm{Sub}(\Gamma)$ without isolated points for the induced topology. In fact, $\mathcal{K}(\Gamma)$ is exactly the set of *condensation points*, that is, the points whose neighborhoods in $\mathrm{Sub}(\Gamma)$ are all uncountable.

Remark 2.2. By a theorem of Brouwer, the space $\mathcal{K}(\Gamma)$ is either empty or a Cantor space, see Theorem. 7.4 in [20].

Remark 2.3. The subset $\mathrm{Sub}_{[\infty]}(\Gamma)$ of infinite index subgroups of Γ is closed in Γ if and only if Γ is finitely generated. Indeed if Γ is finitely generated, then its finite index subgroups are isolated. If Γ is not finitely generated, its finite index subgroups are not

[1] This is why Schreier graphs were defined with respect to right actions.
[2] A topological space is called *perfect* if it has no isolated points.

finitely generated, but they are limit points of finitely generated (thus of infinite index) subgroups; so $\mathrm{Sub}_{[\infty]}(\Gamma)$ is dense in $\mathrm{Sub}(\Gamma)$.

The group Γ acts (on the right) by conjugation via $\Lambda \cdot \gamma := \gamma^{-1}\Lambda\gamma$ on the space of its subgroups $\mathrm{Sub}(\Gamma)$. This action is continuous and the Cantor–Bendixson decomposition $\mathrm{Sub}(\Gamma) = \mathcal{C}(\Gamma) \sqcup \mathcal{K}(\Gamma)$ is Γ-invariant.

By the Baire category theorem, any countable closed subset of $\mathrm{Sub}(\Gamma)$ contains an isolated point, so $\mathrm{Sub}(\Gamma)$ has trivial perfect kernel if and only if it is countable. The following well-known proposition is useful for showing the latter property.

Proposition 2.4. *Let Γ be a countable group, let N be a normal subgroup of Γ such that Γ/N is Noetherian (all its subgroups are finitely generated), and assume that* $\mathrm{Sub}(N)$ *is countable. Then* $\mathrm{Sub}(\Gamma)$ *is countable.*

Proof. Let $\Lambda \leq \Gamma$ and denote by $\pi\colon \Gamma \to \Gamma/N$ the quotient map. Since Γ/N is Noetherian, we have $\pi(\Lambda) = \langle S\rangle$ for some finite set S. Fix a finite set $S' \subseteq \Lambda$ such that $\pi(S') = S$. Then we can recover Λ from S' and its intersection with N as

$$\Lambda = \langle S' \cup (\Lambda \cap N)\rangle.$$

In other words, the map $(S', N') \mapsto \langle S' \cup N'\rangle$ surjects $\mathcal{P}_f(\Gamma) \times \mathrm{Sub}(N)$ onto $\mathrm{Sub}(\Gamma)$, where $\mathcal{P}_f(\Gamma)$ is the set of finite subsets of Γ, which is countable. Since $\mathrm{Sub}(N)$ is countable as well, we conclude that $\mathrm{Sub}(\Gamma)$ is countable. ∎

Corollary 2.5. *If $|m| = 1$ or $|n| = 1$, then* $\mathrm{Sub}(\mathrm{BS}(m,n))$ *is countable.*

Sketch of proof. We sketch the proof contained in Corollary 8.4 of [3]. By symmetry, we may as well assume that $m = 1$. Then $\mathrm{BS}(m,n)$ is isomorphic to the semi-direct product $\mathbb{Z}[1/n] \rtimes \mathbb{Z}$ where $\mathbb{Z}$ acts by multiplication by n. As explained in the proof of Corollary 8.4 in [3], $\mathrm{Sub}(\mathbb{Z}[1/n])$ is countable, so the result follows from the previous proposition. ∎

Space of pointed actions. Let us now interpret the topological space $\mathrm{Sub}(\Gamma)$ in terms of pointed transitive group actions and their pointed Schreier graphs. Given any pointed transitive group action (α, v), where $\alpha\colon V \curvearrowleft \Gamma$ and $v \in V$, we associate to it the stabilizer $\mathrm{Stab}_\alpha(v) \in \mathrm{Sub}(\Gamma)$, and we notice that $\mathrm{Stab}_{\alpha_1}(v_1) = \mathrm{Stab}_{\alpha_2}(v_2)$ if and only if (α_1, v_1) and (α_2, v_2) are isomorphic as pointed transitive actions.

Notation 2.6. We shall denote by $[\alpha, v]$ the isomorphism class of any pointed transitive action (α, v).

We therefore have a canonical bijection $[\alpha, v] \mapsto \mathrm{Stab}_\alpha(v)$ between the collection of isomorphism classes of pointed transitive actions and $\mathrm{Sub}(\Gamma)$. Its inverse is given by $\Lambda \mapsto [\Lambda\backslash\Gamma \curvearrowleft \Gamma, \Lambda]$. Through this bijection, the action by conjugation of Γ on $\mathrm{Sub}(\Gamma)$ becomes $[\alpha, v] \cdot \gamma = [\alpha, v\alpha(\gamma)]$, i.e., it moves the basepoint.

Via the above identification, we obtain a topology on the set of isomorphism classes of pointed actions $[\alpha, v]$.

It is clear that two pointed actions are isomorphic if and only if their Schreier graphs are isomorphic as pointed labeled graphs. Given two pointed labeled oriented graphs $(\mathcal{G}, v)$ and $(\mathcal{H}, w)$, and a positive integer R, we write $(\mathcal{G}, v) \simeq_R (\mathcal{H}, w)$ to mean that the R-balls around v in $\mathcal{G}$ and around w in $\mathcal{H}$ are isomorphic as pointed oriented labeled graphs. It is an exercise to check that if Γ is generated by a finite set S, then the sets of the form

$$\mathcal{N}([\alpha, v], R) := \{[\alpha', v'] : (\mathbf{Sch}(\alpha, S), v) \simeq_R (\mathbf{Sch}(\alpha', S), v')\}, \tag{2.7}$$

constitute a basis of clopen neighborhoods of $[\alpha, v]$.

2.3. Bass–Serre theory

Associated with the standard HNN-presentation of

$$\mathrm{BS}(m,n) = \langle b, t \mid t b^m t^{-1} = b^n \rangle,$$

we have the $\mathrm{BS}(m,n)$-action on its Bass–Serre tree $\mathcal{T}$. Recall that $\mathcal{T}$ is the oriented tree with $V(\mathcal{T}) = \mathrm{BS}(m,n)/\langle b\rangle$, $E^+(\mathcal{T}) = \mathrm{BS}(m,n)/\langle b^n\rangle$,

$$\mathtt{s}(\gamma\langle b^n\rangle) = \gamma\langle b\rangle \quad \text{and} \quad \mathtt{t}(\gamma\langle b^n\rangle) = \gamma t\langle b\rangle$$

and given a subgroup $\Lambda \leq \mathrm{BS}(m,n)$, the quotient $\Lambda\backslash\mathcal{T}$ has the structure of a graph of groups whose fundamental group is Λ, see [25].

Remark 2.8. Let $\Lambda \leq \mathrm{BS}(m,n)$ be a subgroup. If $\Lambda \cap \langle b\rangle = \{\mathrm{id}\}$, then Λ acts freely on $\mathcal{T}$; thus, it is the fundamental group of the quotient graph $\Lambda\backslash\mathcal{T}$, hence Λ is a free group.

Let us now concentrate on a subgroup $\Lambda \leq \mathrm{BS}(m,n)$ such that $\Lambda \cap \langle b\rangle \neq \{\mathrm{id}\}$. Then for the induced action $\Lambda \curvearrowright \mathcal{T}$, each edge and vertex stabilizer is infinite cyclic: the tree $\mathcal{T}$ is a GBS-tree (for Generalized Baumslag–Solitar), in the sense of [15, 21]. One can use this point of view to understand the graph of groups description of Λ. However, taking advantage of the transitivity of the $\mathrm{BS}(m,n)$-action on the edges and the vertices, we provide a slightly more precise description.

Proposition 2.9. *Let m and n be non-zero integers. Let $\Lambda \leq \mathrm{BS}(m,n)$ be a subgroup such that $\Lambda \cap \langle b\rangle \neq \{\mathrm{id}\}$. The quotient graph of groups arising from the action $\Lambda \curvearrowright \mathcal{T}$ is isomorphic to the graph of groups obtained by attaching a copy of $\mathbb{Z}$ to every vertex and every edge of the quotient graph $\Lambda\backslash\mathcal{T}$, with structural maps of positive edges*

$$\mathbb{Z}_e \hookrightarrow \mathbb{Z}_{\mathtt{s}(e)}, \quad k \mapsto \frac{n}{\deg_{\mathrm{out}}(\mathtt{s}(e))}\cdot k \quad \textit{and} \quad \mathbb{Z}_e \hookrightarrow \mathbb{Z}_{\mathtt{t}(e)}, \quad k \mapsto \frac{m}{\deg_{\mathrm{in}}(\mathtt{t}(e))}\cdot k.$$

Proof. In this proof we set $\Gamma := \mathrm{BS}(m,n)$. Let us consider the action of Λ on the tree $\mathcal{T}$. Since $\mathcal{T}$ is locally finite, any edge adjacent to a vertex with infinite stabilizer has itself infinite stabilizer. It follows that all vertex and edge Λ-stabilizers are infinite. Being subgroups of the Γ-stabilizers, they are all isomorphic to $\mathbb{Z}$.

Observe that since Γ acts transitively and the Γ-stabilizers are abelian, the Γ-stabilizers are canonically pairwise isomorphic: given any vertex $u \in V(\mathcal{T})$ and $a \in \mathrm{Stab}_\Gamma(u)$, one has

$$gag^{-1} = hah^{-1} \quad \text{for any } g, h \in \Gamma \text{ such that } gu = hu. \tag{2.10}$$

Indeed, since $h^{-1}g \in \mathrm{Stab}_\Gamma(u)$, we get that $h^{-1}gag^{-1}h = a$.

We now focus on the quotient graph of groups arising from the action $\Lambda \curvearrowright \mathcal{T}$. Let us recall from [25] that its vertex groups are $G_v := \mathrm{Stab}_\Lambda(\tilde{v})$ and its edge groups are $G_e := \mathrm{Stab}_\Lambda(\tilde{e})$, where $\tilde{v}$ and $\tilde{e}$ are some lifts of v and e in $\mathcal{T}$. Given any $e \in E^+(\Lambda\backslash\mathcal{T})$, the structural map $G_e \hookrightarrow G_{\mathtt{t}(e)}$ is

$$G_e = \mathrm{Stab}_\Lambda(\tilde{e}) \hookrightarrow \mathrm{Stab}_\Lambda(\mathtt{t}(\tilde{e})) \to \mathrm{Stab}_\Lambda(\widetilde{\mathtt{t}(e)}) = G_{\mathtt{t}(e)}, \quad a \mapsto a \mapsto gag^{-1}, \tag{2.11}$$

where $g \in \Lambda$ is any element such that $g \cdot \mathtt{t}(\tilde{e}) = \widetilde{\mathtt{t}(e)}$ and the map $G_e \hookrightarrow G_{\mathtt{s}(e)}$ is similar. This is unambiguous by (2.10).

Let us call *orientation* of an infinite cyclic group the choice of one generator (over two). This provides an identification to $\mathbb{Z}$. Once every stabilizer is oriented, the inclusions $G_e \hookrightarrow G_{\mathtt{s}(e)}$ and $G_e \hookrightarrow G_{\mathtt{t}(e)}$ become multiplications by non-zero integers $\lambda_\Lambda^-(e)$ and $\lambda_\Lambda^+(e)$, respectively. It now suffices to prove that, for well-chosen orientations, one has

$$\lambda_\Lambda^-(e) = \frac{n}{\deg_{\mathrm{out}}(\mathtt{s}(e))} \quad \text{and} \quad \lambda_\Lambda^+(e) = \frac{m}{\deg_{\mathrm{in}}(\mathtt{t}(e))} \tag{2.12}$$

for every positive edge $e \in E^+(\Lambda\backslash\mathcal{T})$.

Let us first observe that the absolute value of $\lambda_\Lambda^\pm(e)$ does not depend on the orientations: it is equal to $[G_v : G_e]$. In other words, if $\tilde{e}$ is a lift of e, $\tilde{v} := \mathtt{s}(\tilde{e}) =$ and $\tilde{w} := \mathtt{t}(\tilde{e})$, we have

$$|\lambda_\Lambda^-(e)| = [\mathrm{Stab}_\Lambda(\tilde{v}) : \mathrm{Stab}_\Lambda(\tilde{e})] = |\mathrm{Stab}_\Lambda(\tilde{v}) \cdot \tilde{e}|, \tag{2.13}$$

$$|\lambda_\Lambda^+(e)| = [\mathrm{Stab}_\Lambda(\tilde{w}) : \mathrm{Stab}_\Lambda(\tilde{e})] = |\mathrm{Stab}_\Lambda(\tilde{w}) \cdot \tilde{e}|. \tag{2.14}$$

Let $E_{\mathrm{out}}(\tilde{v})$ be the set of $\tilde{v}$-outgoing edges. Its cardinal is $|E_{\mathrm{out}}(\tilde{v})| = |n|$. Any generator of $\mathrm{Stab}_\Gamma(\tilde{v})$ acts as a single $|n|$-cycle on $E_{\mathrm{out}}(\tilde{v})$. Hence, $E_{\mathrm{out}}(\tilde{v})$ splits into $\mathrm{Stab}_\Lambda(\tilde{v})$-orbits of equal size, that is, $|\lambda_\Lambda^-(e)|$ according to (2.13). The number of these $\mathrm{Stab}_\Lambda(\tilde{v})$-orbits is $\deg_{\mathrm{out}}(v)$, thus $|n| = |\lambda_\Lambda^-(e)| \cdot \deg_{\mathrm{out}}(v)$. We obtain similarly $|m| = |\lambda_\Lambda^+(e)| \cdot \deg_{\mathrm{in}}(w)$, using incoming edges and (2.14). We have established that (2.12) holds in absolute value.

Let us now turn to the signs in (2.12), for which we need explicit orientations of the Λ-stabilizers. We actually start by orienting the Γ-stabilizers.

Pick the vertex $\tilde{u}_0 := \langle b \rangle \in V(\mathcal{T})$. Then $\mathrm{Stab}_\Gamma(\tilde{u}_0) = \langle b \rangle$ and the positive edge $\tilde{d}_0 := \langle b^n \rangle \in E^+(\mathcal{T})$ has source $\tilde{u}_0$ and target $t\tilde{u}_0$. Since the Γ-stabilizers are canonically pairwise identified by conjugation (2.10), these choices induce a canonical conjugation-invariant orientation x_* of all the vertex and edge Γ-stabilizers: $x_{g\tilde{u}_0} := gbg^{-1}$ for $\mathrm{Stab}_\Gamma(g\tilde{u}_0)$ and $x_{g\tilde{d}_0} := gb^ng^{-1}$ for $\mathrm{Stab}_\Gamma(g\tilde{d}_0)$.

The inclusions $\mathrm{Stab}_\Gamma(\tilde{e}) \hookrightarrow \mathrm{Stab}_\Gamma(\mathtt{s}(\tilde{e}))$ and $\mathrm{Stab}_\Gamma(\tilde{e}) \hookrightarrow \mathrm{Stab}_\Gamma(\mathtt{t}(\tilde{e}))$ become multiplications by non-zero integers that we denote by $\mu_\Gamma^-(\tilde{e})$ and $\mu_\Gamma^+(\tilde{e})$. We have $\mu_\Gamma^-(\tilde{e}) = n$ since $x_{\tilde{e}} = x_{\mathtt{s}(\tilde{e})}^n$, and $\mu_\Gamma^+(\tilde{e}) = m$ since

$$x_{\tilde{e}} = gb^ng^{-1} = g(tbt^{-1})^m g^{-1} = x_{\mathtt{t}(\tilde{e})}^m.$$

The Λ-stabilizers have finite index in the corresponding Γ-stabilizers. We orient them coherently with the ambient Γ-stabilizers by using positive powers. The Λ-conjugations between Λ-stabilizers remain orientation-preserving, therefore by (2.11) the inclusion map $\mathrm{Stab}_\Lambda(\tilde{e}) \hookrightarrow \mathrm{Stab}_\Lambda(\mathtt{t}(\tilde{e}))$ becomes the multiplication by $\lambda_\Lambda^+(e)$. Similarly, the inclusion $\mathrm{Stab}_\Lambda(\tilde{e}) \hookrightarrow \mathrm{Stab}_\Lambda(\mathtt{s}(\tilde{e}))$ becomes multiplication by $\lambda_\Lambda^-(e)$. Since the orientations are coherent, we conclude that $\lambda_\Lambda^-(e)$ has the same sign as $\mu_\Gamma^-(e) = n$ and $\lambda_\Lambda^+(e)$ has the same sign as $\mu_\Gamma^+(e) = m$. ■

Corollary 2.15. *Let m and n be non-zero integers. Let $\Lambda \le \mathrm{BS}(m,n)$ be a subgroup such that $\Lambda \cap \langle b\rangle \neq \{\mathrm{id}\}$. The isomorphism type of Λ is completely determined by the oriented graph $\Lambda\backslash\mathcal{T}$.*

Proposition 2.16. *Let m and n be non-zero integers and let $\Lambda \le \mathrm{BS}(m,n)$ be a subgroup.*

(1) *If $\Lambda \cap \langle b\rangle \neq \{\mathrm{id}\}$, then either $\Lambda \simeq \mathbb{Z}$ is virtually a subgroup of $\langle b\rangle$ or Λ is not a free group.*

(2) *If $|m| = 1$ or $|n| = 1$, then the fundamental group of the underlying graph $\Lambda\backslash\mathcal{T}$ is a free group of rank ≤ 1.*

If $\Lambda \cap \langle b\rangle = \{\mathrm{id}\}$, then Λ is the fundamental group of the underlying graph $\Lambda\backslash\mathcal{T}$ (see Remark 2.8).

The first item of the proposition follows from standard techniques in ℓ^2-cohomology: if $\Lambda \cap \langle b\rangle \neq \{\mathrm{id}\}$, then Λ is the fundamental group of a graph of groups whose vertex and edge groups are isomorphic to $\mathbb{Z}$; all the ℓ^2-Betti numbers of such a group vanish. For the comfort of the reader, we propose a proof by hand.

Proof. We start with the first item. Recall that in a free group F, whenever non-trivial elements $g, h \in F$ satisfy $gh^kg^{-1} = h^l$ with $k \neq 0 \neq l$, then there is $a \in F$ such that g, h are both powers of a. In particular, such elements g, h always commute.

Now, assume that Λ is free and $\Lambda \cap \langle b\rangle \neq \{\mathrm{id}\}$, say $\Lambda \cap \langle b\rangle = \langle b^s\rangle$, where $s > 0$. Pick any $\lambda \in \Lambda$ and set $H_\lambda := \langle b^s\rangle \cap \lambda\langle b^s\rangle\lambda^{-1}$, which is the intersection of Λ with the stabilizer of the geodesic $[\langle b\rangle, \lambda\langle b\rangle]$ in $\mathcal{T}$. Observe that H_λ is a finite index subgroup of both $\langle b^s\rangle$ and $\lambda\langle b^s\rangle\lambda^{-1}$. Therefore, there are $k \neq 0 \neq l$ such that $\lambda b^{sk}\lambda^{-1} = b^{sl}$. As Λ is free, λ and b^s commute.

Consequently, the center of Λ contains $\langle b^s\rangle$. Thus, the rank of Λ is 1; in other words, Λ is infinite cyclic. It is now clear that $\langle b^s\rangle$ has finite index in both Λ and $\langle b\rangle$, so Λ is virtually a subgroup of $\langle b\rangle$.

Let us turn to the second item. The fundamental group of a graph of groups surjects onto the fundamental group of the underlying graph. The condition in item (2) implies the amenability of $\mathrm{BS}(m,n)$. Its subgroups thus cannot surject onto a non-amenable free group. ■

3. Bass–Serre graphs

3.1. Pre-actions

Let $m, n \in \mathbb{Z} \setminus \{0\}$ and $\mathrm{BS}(m,n) = \langle b, t \mid tb^m = b^n t\rangle$.

Recall that a partial bijection of a set X is a bijection between two subsets of X. Our actions are on the right; thus in a product of (partial) bijections $\sigma\tau$, the transformation σ is applied first.

Definition 3.1. Given a bijection β of a set X and a partial bijection τ of X, we say that τ is (β^n, β^m)-*equivariant* if $\tau\beta^m = \beta^n\tau$ as partial bijections, that is,

- $\mathrm{dom}(\tau)$ is β^n-invariant,
- $\mathrm{rng}(\tau)$ is β^m-invariant,
- $x\tau\beta^m = x\beta^n\tau$ for all $x \in \mathrm{dom}(\tau)$.

A *pre-action* of $\mathrm{BS}(m,n)$ on a set X is a couple (β, τ) where β is a bijection of X and τ is a (β^n, β^m)-equivariant partial bijection of X. The set X is called the *domain* of the pre-action. Such a pre-action is *saturated* if $\mathrm{dom}(\tau) = X = \mathrm{rng}(\tau)$.

Remark 3.2. Saturated pre-actions (β, τ) correspond to actions α of $\mathrm{BS}(m,n)$ on the same set X under the association $\beta \leftrightarrow \alpha(b)$ and $\tau \leftrightarrow \alpha(t)$.

Definition 3.3. Given a pre-action (β, τ) of $\mathrm{BS}(m,n)$, its *Schreier graph* is the oriented labeled graph $\mathbf{Sch}(\beta, \tau) = \mathcal{G}$ defined by

$$V(\mathcal{G}) := X \quad \text{and} \quad \begin{cases} E^+(\mathcal{G}) := X \times \{b\} \sqcup \mathrm{dom}(\tau) \times \{t\}, \\ E^-(\mathcal{G}) := X \times \{b^{-1}\} \sqcup \mathrm{rng}(\tau) \times \{t^{-1}\}, \end{cases}$$

where the label of any edge is its second component, and

- for all $x \in X$, we set
$$\mathtt{s}(x,b) := x, \quad \mathtt{t}(x,b) := x\beta \quad \text{and} \quad \overline{(x,b)} := (x\beta, b^{-1}),$$
- for all $x \in \mathrm{dom}(\tau)$, we set
$$\mathtt{s}(x,t) := x, \quad \mathtt{t}(x,t) := x\tau, \quad \text{and} \quad \overline{(x,t)} := (x\tau, t^{-1}).$$

Notice that the orientation of any edge (x,l) is determined by its label l and that the source of (x,l) is x, regardless of its orientation.

Noting that a $\mathrm{BS}(m,n)$-action is transitive if and only if the associated Schreier graph is connected, we make the following definition.

Definition 3.4. A pre-action of $\mathrm{BS}(m,n)$ is *transitive* if its Schreier graph is connected.

3.2. Bass–Serre graphs

We now introduce an important tool for our study. It is the labeled graph obtained from the Schreier graph defined in Section 3.1 by "shrinking each β-orbit to one point". We identify together the t-edges whose initial vertices belong to the same β^n-orbit. Note that their terminal vertices automatically belong to the same β^m-orbit.

We label the vertices by the cardinality of the corresponding β-orbit and the edges by the cardinality of the corresponding β^n-orbit. This is illustrated by Figure 2. The formal definition is as follows.

Definition 3.5. The *Bass–Serre graph* associated to a pre-action $\alpha = (\beta, \tau)$ of $\mathrm{BS}(m,n)$ on a set X is the oriented labeled graph $\mathbf{BS}(\alpha)$ defined by

$$V(\mathbf{BS}(\alpha)) := X/\langle\beta\rangle \quad \text{and} \quad \begin{cases} E^+(\mathbf{BS}(\alpha)) := \mathrm{dom}(\tau)/\langle\beta^n\rangle, \\ E^-(\mathbf{BS}(\alpha)) := \mathrm{rng}(\tau)/\langle\beta^m\rangle. \end{cases}$$

For every $x \in \operatorname{dom}\tau$, we set

$$\mathtt{s}(x\langle\beta^n\rangle) := x\langle\beta\rangle, \quad \mathtt{t}(x\langle\beta^n\rangle) := x\tau\langle\beta\rangle \quad \text{and} \quad \overline{x\langle\beta^n\rangle} := x\tau\langle\beta^m\rangle = x\langle\beta^n\rangle\tau.$$

We define the label map $L\colon V(\mathbf{BS}(\alpha)) \sqcup E(\mathbf{BS}(\alpha)) \to \mathbb{Z}_{\geq 1} \cup \{\infty\}$ by

$$L(x\langle\beta\rangle) := |x\langle\beta\rangle|, \quad L(x\langle\beta^n\rangle) := |x\langle\beta^n\rangle|, \quad L(y\langle\beta^m\rangle) := |y\langle\beta^m\rangle|.$$

Remark 3.6. For any $x \in \operatorname{dom}(\tau)$, the (β^n, β^m)-equivariant partial bijection τ induces a bijection from $x\langle\beta^n\rangle$ to $x\tau\langle\beta^m\rangle$. Thus, both the target and the opposite maps of $\mathbf{BS}(\alpha)$ are well defined and the label of each edge is equal to the label of its opposite.

Remark 3.7. We view the sets $E^+(\mathbf{BS}(\alpha))$ and $E^-(\mathbf{BS}(\alpha))$ as disjoint sets, even though we might have that $\operatorname{dom}(\tau)/\langle\beta^n\rangle \cap \operatorname{rng}(\tau)/\langle\beta^m\rangle \neq \varnothing$. Note that the source of an edge $x\langle\beta^k\rangle \in E^\pm(\mathbf{BS}(\alpha))$ is $x\langle\beta\rangle$ regardless of its orientation.

Remark 3.8. The groups $\mathrm{BS}(m,n)$ and $\mathrm{BS}(n,m)$ are isomorphic via $b \mapsto b$ and $t \mapsto t^{-1}$. For every pre-action (β, τ) of $\mathrm{BS}(m,n)$, the couple (β, τ^{-1}) is a pre-action of $\mathrm{BS}(n,m)$. At the level of Bass–Serre graphs, $\mathbf{BS}(\beta,\tau)$ and $\mathbf{BS}(\beta,\tau^{-1})$ coincide, except that the orientation is reversed.

Remark 3.9. In the case of a transitive $\mathrm{BS}(m,n)$-action, the graph underlying our Bass–Serre graph is the quotient of the Bass–Serre tree $\mathcal{T}$ by the stabilizer of any point of X, as will be explained in Section 3.6.

We now clarify what we meant by "shrinking each β-orbit to a point", by noting that we have the following simplicial map from the Schreier graph to the Bass–Serre graph of any pre-action.

Definition 3.10. The *projection* associated to a pre-action $\alpha = (\beta, \tau)$ is the application π_α given by

$$\begin{aligned} V(\mathbf{Sch}(\alpha)) &\to V(\mathbf{BS}(\alpha)), & x &\mapsto x\langle\beta\rangle, \\ E_t^+(\mathbf{Sch}(\alpha)) &\to E^+(\mathbf{BS}(\alpha)), & (x,t) &\mapsto x\langle\beta^n\rangle, \\ E_t^-(\mathbf{Sch}(\alpha)) &\to E^-(\mathbf{BS}(\alpha)), & (x,t^{-1}) &\mapsto x\langle\beta^m\rangle, \\ E_b(\mathbf{Sch}(\alpha)) &\to V(\mathbf{BS}(\alpha)), & (x,b^{\pm 1}) &\mapsto x\langle\beta\rangle, \end{aligned}$$

where $E_t^\pm(\mathbf{Sch}(\alpha))$ is the subset of edges in $\mathbf{Sch}(\alpha)$ whose label is t or t^{-1} respectively, and E_b is the subset of edges whose label is b or b^{-1}.

This projection is illustrated in Figure 2. Given any subgraph $\mathcal{G} \subseteq \mathbf{Sch}(\alpha)$ or path $\mathfrak{p}$ in $\mathbf{Sch}(\alpha)$, we obtain a subgraph $\pi_\alpha(\mathcal{G}) \subseteq \mathbf{BS}(\alpha)$ or a path $\pi_\alpha(\mathfrak{p})$ in $\mathbf{BS}(\alpha)$.

Note that for every vertex $v = x\langle\beta\rangle$,

$$|x\langle\beta^k\rangle| = \frac{|x\langle\beta\rangle|}{\gcd(|x\langle\beta\rangle|, k)},$$

thus the following facts hold:

- all the v-outgoing edges e have the same label, which is

$$L(e) = \frac{L(v)}{\gcd(L(v), n)},$$

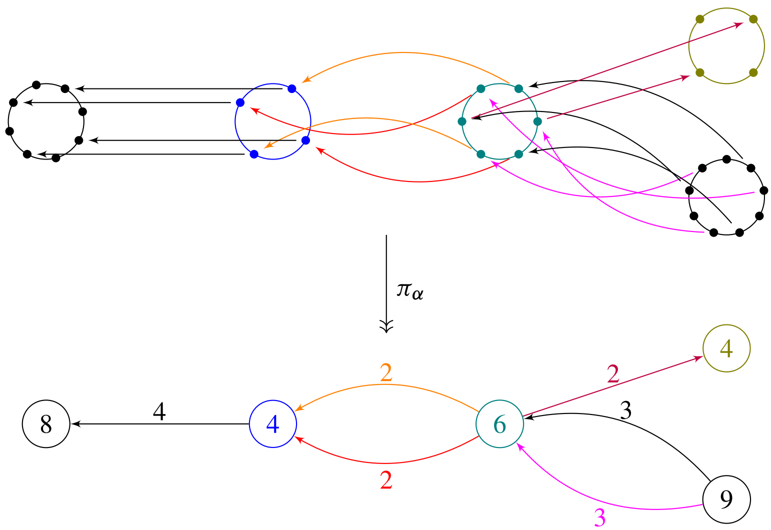

Figure 2. The projection from the Schreier graph onto the Bass–Serre graph of some non-saturated transitive $\mathrm{BS}(2,3)$-pre-action. The dotted circles represent the β-orbits in the Schreier graph.

- all the v-incoming edges e' have the same label, which is

$$L(e') = \frac{L(v)}{\gcd(L(v), m)}.$$

We also have the following relations between labels and degrees:

- The outgoing degree $\deg_{\text{out}}(v)$ is equal to the number of β^n-orbits contained in $x\langle\beta\rangle \cap \operatorname{dom}(\tau)$. Recall that $\operatorname{dom}(\tau)$ is β^n-invariant. Since $x\langle\beta\rangle$ contains exactly $\gcd(L(v), n)$ orbits under β^n, we get

$$\deg_{\text{out}}(v) \le \gcd(L(v), n),$$

  with equality if and only if $x\langle\beta\rangle \subseteq \operatorname{dom}(\tau)$.
- Similarly, the incoming degree $\deg_{\text{in}}(v)$ is equal to the number of β^m-orbits contained in $x\langle\beta\rangle \cap \operatorname{rng}(\tau)$, so

$$\deg_{\text{in}}(v) \le \gcd(L(v), m),$$

  with equality if and only if $x\langle\beta\rangle \subseteq \operatorname{rng}(\tau)$.

Remark 3.11. As a consequence of the last two items, the pre-action is an action if and only if, for every vertex v,

$$\deg_{\text{out}}(v) = \gcd(L(v), n) \quad \text{and} \quad \deg_{\text{in}}(v) = \gcd(L(v), m).$$

3.3. (m, n)-graphs

We now introduce an axiomatization of the Bass–Serre graphs we obtain from pre-actions. Recall that by convention $\gcd(\infty, k) = |k|$ for all $k \neq 0$.

Definition 3.12. An (m,n)-*graph* is an oriented labeled graph $\mathcal{G} = (V, E)$ with label map $L\colon V \sqcup E \to \mathbb{Z}_{\geq 1} \cup \{\infty\}$ such that:

- for every positive edge $e \in E^+$,

$$\frac{L(\mathtt{s}(e))}{\gcd(L(\mathtt{s}(e)), n)} = L(e) = \frac{L(\mathtt{t}(e))}{\gcd(L(\mathtt{t}(e)), m)}, \tag{3.13}$$

- for every negative edge $e \in E^-$, $L(e) = L(\bar{e})$,
- for every vertex $v \in V$, we have

$$\deg_{\text{out}}(v) \leq \gcd(L(v), n) \quad \text{and} \quad \deg_{\text{in}}(v) \leq \gcd(L(v), m). \tag{3.14}$$

Example 3.15. The Bass–Serre graph of any pre-action of $\mathrm{BS}(m,n)$ is an (m,n)-graph. The converse will be shown in Proposition 3.22.

Remark 3.16. Observe that an edge label is uniquely determined by the label of any of its vertices. The edge labels are thus redundant and are just calculation tools (see also Remark 4.7).

Example 3.17. Let us see how labels interact for $m = 2$ and $n = 3$. If e is an edge in a $(2,3)$-graph, then once we fix the label of one of the extremities, the other one can be chosen according to the following rules, using formula (3.13) for $L(e)$:

- If $\gcd(L(\mathtt{s}(e)), 2) = 1$, then $L(\mathtt{t}(e)) \in \{L(e),\ 2L(e)\}$.
- If $\gcd(L(\mathtt{s}(e)), 2) = 2$, then $L(\mathtt{t}(e)) = 2L(e)$.
- If $\gcd(L(\mathtt{t}(e)), 3) = 1$, then $L(\mathtt{s}(e)) \in \{L(e),\ 3L(e)\}$.
- If $\gcd(L(\mathtt{t}(e)), 3) = 3$, then $L(\mathtt{s}(e)) = 3L(e)$.

The reader is invited to consult the web page [12] to see the kinds of local constraints which occur in general. In Figure 3, we give an illustrative example.

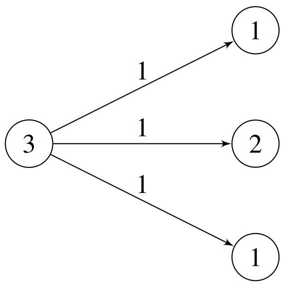

(a) Various choices for the label $L(t(e))$.

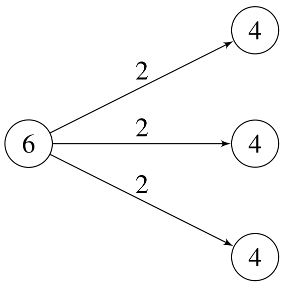

(b) No choice for the label $L(t(e))$.

Figure 3. Two examples of $(2,3)$-graphs.

Remark 3.18. As in Remark 3.8, every (m,n)-graph can be turned into an (n,m)-graph by flipping the orientations of its edges. Note that this operation does not affect the labels.

Remark 3.19. In a connected (m,n)-graph, the labels are, either all finite, or all ∞ by equation (3.13). This will be made more precise in Proposition 4.6. Observe that any oriented graph $\mathcal{G}$ satisfying $\deg_{\text{in}}(v) \le m$ and $\deg_{\text{out}}(v) \le n$, for every $v \in V(\mathcal{G})$, becomes an (m,n)-graph if we set all the labels to be infinite. However, one cannot always put finite labels, see Lemma 3.33.

Definition 3.20. Let $\mathcal{G}$ be an (m,n)-graph. A vertex v in $\mathcal{G}$ is *saturated* if the inequalities (3.14) are indeed equalities, i.e.,

$$\deg_{\text{out}}(v) = \gcd(L(v), n) \quad \text{and} \quad \deg_{\text{in}}(v) = \gcd(L(v), m).$$

The (m,n)-graph $\mathcal{G}$ is *saturated* if all its vertices are saturated.

Example 3.21. The Bass–Serre graph of a pre-action of $\mathrm{BS}(m,n)$ is saturated if and only if the pre-action is an action.

3.4. Realizing (m,n)-graphs as Bass–Serre graphs

Proposition 3.22. *Every (m,n)-graph $\mathcal{G}$ is the Bass–Serre graph of at least one pre-action of* $\mathrm{BS}(m,n)$. *Any such pre-action is an action if and only if $\mathcal{G}$ is saturated.*

The above proposition is a consequence of the following stronger statement where by definition, a *sub-(m,n)-graph* of an (m,n)-graph $\mathcal{G}$ is a subgraph $\mathcal{G}'$ labeled by the restriction of the label map of $\mathcal{G}$.

Proposition 3.23 (Extension of pre-actions from (m,n)-graphs). *Let $\mathcal{G}_1$ be the Bass–Serre graph of a pre-action α_1, and let $\mathcal{G}_2$ be an (m,n)-graph that contains $\mathcal{G}_1$ as a sub-(m,n)-graph. Then $\mathcal{G}_2$ is the Bass–Serre graph of a pre-action α_2 that extends α_1.*

Proof. We start with a pre-action (β_1, τ_1) on X_1, which yields the Bass–Serre graph $\mathcal{G}_1$. Let $W := V(\mathcal{G}_2) \setminus V(\mathcal{G}_1)$ and $X_2 := X_1 \sqcup \bigsqcup_{v \in W} X_v$, where each X_v is a set of cardinality $|X_v| = L(v)$. We extend β_1 to a permutation β_2 of X_2 by making it act as a cycle of length $L(v)$ on X_v.

By Zorn's lemma, it suffices to extend τ_1 when $\mathcal{G}_1$ only lacks one positive $\mathcal{G}_2$-edge. So, suppose $E^+(\mathcal{G}_1) \sqcup \{e\} = E^+(\mathcal{G}_2)$. Then, by inequality (3.14) in Definition 3.12,

$$\deg_{\text{out}}^{\mathcal{G}_1}(\mathtt{s}(e)) < \deg_{\text{out}}^{\mathcal{G}_2}(\mathtt{s}(e)) \le \gcd(L(\mathtt{s}(e)), n)$$

and, similarly,

$$\deg_{\text{in}}^{\mathcal{G}_1}(\mathtt{t}(e)) < \deg_{\text{in}}^{\mathcal{G}_2}(\mathtt{t}(e)) \le \gcd(L(\mathtt{t}(e)), m).$$

We can thus find a β_2^n-orbit $y\langle\beta_2^n\rangle$ contained in the β_2-orbit $\mathtt{s}(e)$ but disjoint from $\mathrm{dom}(\tau_1)$ and a β_2^m-orbit $z\langle\beta_2^m\rangle$ contained in the β_2-orbit $\mathtt{t}(e)$ but disjoint from $\mathrm{rng}(\tau_1)$.

Since these two orbits $y\langle\beta_2^n\rangle$ and $z\langle\beta_2^m\rangle$ share the same cardinal $L(e)$, we can define τ_2 as an extension of τ_1 which is also (β_2^n, β_2^m)-equivariant when restricted to $y\langle\beta_2^n\rangle$ by letting

$$y\,\beta_2^{kn}\,\tau_2 = z\,\beta_2^{km} \quad \text{for all } k \in \mathbb{Z}.$$

By construction, τ_2 is the desired extension. ■

The pre-action α_2 arising in Proposition 3.23 is definitively not unique in general. In a forthcoming work, we will characterize which (m,n)-graphs arise as Bass–Serre graphs of continuum many non-isomorphic actions. In particular, we will show that the (m,n)-graphs whose underlying graph has non-finitely generated fundamental group are of this kind. Such (m,n)-graphs always exist as long as $|m| \geq 2$ and $|n| \geq 2$. Here we give a simple example of a graph associated to continuum many non-isomorphic actions for $n = m = 2$.

Example 3.24. Let $\mathcal{G}$ be the $(2,2)$-graph whose underlying graph is such that $V(\mathcal{G}) = \mathbb{Z}$ and for every $z \in V(\mathcal{G})$, there are exactly two z-outgoing edges, one to z and the other to $z+1$, see Figure 4. That is, $\mathcal{G}$ is a line where every vertex has an extra loop. We set the labels of $\mathcal{G}$ to be all infinite.

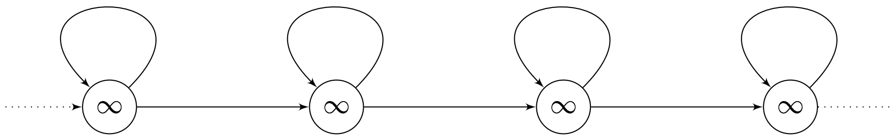

Figure 4. The (2,2)-graph $\mathcal{G}$.

Set $X := V(\mathcal{G}) \times \mathbb{Z} \cong \mathbb{Z} \times \mathbb{Z}$. For every function $f\colon \mathbb{Z} \to \mathbb{Z}$ such that for all $w < 0$, $f(w) = 0$ and $f(0) \neq 0$, we define an action α_f as follows, for all $(k,l) \in X$:

$$(k,l)\alpha_f(b) := (k, l+1), \qquad (k,l)\alpha_f(t) := \begin{cases} (k+1, l) & \text{if } l \text{ is odd,} \\ (k, l+f(k)) & \text{if } l \text{ is even.} \end{cases}$$

It is easy to check that all α_f are actions of $\mathrm{BS}(2,2)$ whose Bass–Serre graph is $\mathcal{G}$, that α_f and α_g are non-conjugate for $f \neq g$, and that there are continuum many such actions.

3.5. Additional properties of (m,n)-graphs

In this section, we collect some basic consequences of the definition of (m,n)-graphs. Observe that equation (3.13) is equivalent to the fact that

$$\max(|L(s(e))|_p - |n|_p, 0) = |L(e)|_p = \max(|L(\mathtt{t}(e))|_p - |m|_p, 0), \tag{3.25}$$

from which we obtain the following.

Remark 3.26. Consider an oriented labeled graph $\mathcal{G} = (V, E)$ with label map $L\colon V \sqcup E \to \mathbb{Z}_{\geq 1}$ satisfying $L(\bar{e}) = L(e)$ for every edge e. The labeled graph $\mathcal{G}$ is an (m,n)-graph if and only if the following two conditions hold:

- for every positive edge e and every prime p such that $|L(e)|_p \geq 1$,
$$|L(\mathtt{s}(e))|_p = |L(e)|_p + |n|_p \quad \text{and} \quad |L(\mathtt{t}(e))|_p = |L(e)|_p + |m|_p, \tag{3.27}$$
- for every positive edge e and every prime p such that $|L(e)|_p = 0$,
$$0 \leq |L(\mathtt{s}(e))|_p \leq |n|_p \quad \text{and} \quad 0 \leq |L(\mathtt{t}(e))|_p \leq |m|_p. \tag{3.28}$$

In particular, in an (m,n)-graph, $|L(\mathtt{s}(e))|_p > |n|_p$ if and only if $|L(\mathtt{t}(e))|_p > |m|_p$, and if one of these two equivalent conditions is met, then

$$|L(\mathtt{t}(e))|_p = |L(\mathtt{s}(e))|_p + |m|_p - |n|_p. \tag{3.29}$$

Lemma 3.30. *Let p be a prime number such that $|n|_p < |m|_p$ and let $\mathcal{G}$ be an (m,n)-graph. If $(e_k)_{k\geq 1}$ is any infinite path consisting of positive edges with $L(\mathtt{s}(e_1)) \neq \infty$ and $|L(\mathtt{s}(e_1))|_p > |n|_p$, then*

$$\lim_{k\to+\infty} |L(\mathtt{s}(e_k))|_p = +\infty.$$

If $(e_k)_{k\geq 1}$ is any infinite path consisting of negative edges with $L(\mathtt{s}(e_1)) \neq \infty$, then

$$\limsup_{k\to+\infty} |L(\mathtt{s}(e_k))|_p < |m|_p.$$

Proof. If $(e_k)_{k\geq 1}$ is an infinite path of positive edges such that $|L(\mathtt{s}(e_1))|_p > |n|_p$, then by a straightforward induction using Equation (3.29) we have that

$$|L(\mathtt{s}(e_k))|_p = |L(\mathtt{s}(e_1))|_p + k(|m|_p - |n|_p) \tag{3.31}$$

for all $k \geq 1$. The first result follows.

For the second one, let $(e_k)_{k\geq 1}$ be an infinite path consisting of negative edges. By exchanging the roles in equation (3.29), we have the claim:

if e is a negative edge, then $|L(\mathtt{s}(e))|_p > |m|_p$ if and only if $|L(\mathtt{t}(e))|_p > |n|_p$; and when this occurs, $|L(\mathtt{t}(e))|_p = |L(\mathtt{s}(e))|_p - |m|_p + |n|_p$.

Thus, $|L(\mathtt{s}(e_{k+1}))|_p = |L(\mathtt{t}(e_k))|_p < |L(\mathtt{s}(e_k))|_p$ as long as $|L(\mathtt{s}(e_k))|_p > |m|_p$. So there must be $k_0 \in \mathbb{N}$ such that $|L(\mathtt{s}(e_{k_0}))|_p \leq |m|_p$ (this could have already happened for $k_0 = 1$). From this point, we have $|L(\mathtt{s}(e_{k_0+1}))|_p = |L(\mathtt{t}(e_{k_0}))|_p \leq |n|_p < |m|_p$ and an inductive use of the claim gives $|L(\mathtt{s}(e_k))|_p \leq |n|_p < |m|_p$ for all $k > k_0$. This finishes the proof. ∎

Remark 3.32. It follows from equation (3.31) that any infinite path $(e_k)_{k\geq 1}$ consisting of positive edges with $L(\mathtt{s}(e_1)) \neq \infty$ and $|L(\mathtt{s}(e_1))|_p > |n|_p$ has to be a simple path.

Lemma 3.33. *If $|m| > |n|$ and $\mathcal{G}$ is an (m,n)-graph with a vertex of finite label, then there is a vertex $v \in V(\mathcal{G})$ such that $\deg_{\mathrm{in}}(v) < |m|$.*

Proof. Assume by contradiction that $\deg_{\mathrm{in}}(v) = |m|$ for all $v \in V(\mathcal{G})$. Then we can build inductively an infinite path $(e_k)_{k\in\mathbb{N}}$ consisting of negative edges with $L(\mathtt{s}(e_0))$ finite. By the previous lemma, this path goes through some vertex v_0 that $|L(v_0)|_p < |m|_p$. Then $\deg_{\mathrm{in}}(v_0) = \gcd(L(v_0), m) < |m|$, a contradiction. ∎

3.6. Bass–Serre theory for BS(m, n)

Take $m, n \in \mathbb{Z} \setminus \{0\}$. Set $\Gamma := \mathrm{BS}(m,n) = \langle b, t \mid tb^m t^{-1} = b^n \rangle$ and put $S := \{b, t\}$. Denote by $\mathcal{T}$ the associated Bass–Serre tree and remark that it is the underlying oriented graph of the Bass–Serre graph of the transitive and free action: $\mathcal{T} = \mathbf{BS}(\Gamma \curvearrowleft \Gamma)$.

Besides the Schreier graph, we can associate to each subgroup $\Lambda \leq \Gamma$ two decorated graphs:

- the Bass–Serre graph of the action $\Lambda\backslash\Gamma \curvearrowleft \Gamma$,
- the quotient graph of groups $\Lambda\backslash\mathcal{T}$ of the action $\Lambda \curvearrowright \mathcal{T}$.

Let us observe that the underlying oriented graphs of the two above decorated graphs are the same. Indeed, they are obtained as quotients of commuting actions as one can see in the following diagram, where by $\curvearrowleft^V \langle b\rangle$ we mean that $\langle b\rangle$ acts only on the set of vertices, where the $\swarrow$ arrows are graph morphisms obtained by quotienting by left Λ-actions, and where the dashed $\searrow$ arrows are projections as in Definition 3.10:

$$\begin{array}{ccccc}
 & & \Lambda \curvearrowright \mathbf{Cay}(\Gamma, S) \curvearrowleft^V \langle b\rangle & & \\
 & \swarrow & & \searrow & \\
\Lambda\backslash\mathbf{Cay}(\Gamma, S) \curvearrowleft^V \langle b\rangle & & & & \Lambda \curvearrowright \mathbf{BS}(\Gamma \curvearrowleft \Gamma) \\
\| & & & & \| \\
\mathbf{Sch}(\Lambda, S) \curvearrowleft^V \langle b\rangle & & & & \Lambda \curvearrowright \mathcal{T} \\
 & \searrow & & \swarrow & \\
 & & \mathbf{BS}(\Lambda\backslash\Gamma \curvearrowleft \Gamma) \simeq \Lambda\backslash\mathcal{T}. & &
\end{array}$$

Next, observe that, $\mathbf{BS}(\Lambda\backslash\Gamma \curvearrowleft \Gamma)$ being saturated, one has $\deg_{\mathrm{in}}(v) = \gcd(L(v), m)$ and $\deg_{\mathrm{out}}(v) = \gcd(L(v), n)$ for every vertex v in this graph. Hence, for every edge e, one has

$$\frac{L(\mathtt{s}(e))}{L(e)} = \gcd(L(\mathtt{s}(e)), n) = \deg_{\mathrm{out}}(\mathtt{s}(e)) \quad \text{and} \quad \frac{L(\mathtt{t}(e))}{L(e)} = \deg_{\mathrm{in}}(\mathtt{t}(e)).$$

Thus, Remark 2.8 and Proposition 2.9 can be immediately reformulated in terms of the labels of the Bass–Serre graph $\mathbf{BS}(\Lambda\backslash\Gamma \curvearrowleft \Gamma)$ as follows.

Proposition 3.34. *Let m and n be non-zero integers. Let $\mathcal{G}$ be a saturated connected (m, n)-graph and let Λ be a subgroup of $\Gamma = \mathrm{BS}(m, n)$ such that $\mathbf{BS}(\Lambda\backslash\Gamma \curvearrowleft \Gamma) \simeq \mathcal{G}$.*

(1) *If all labels of $\mathcal{G}$ are infinite, then Λ is a free group, namely, isomorphic to the fundamental group of the graph $\mathcal{G}$.*

(2) *If all labels of $\mathcal{G}$ are finite, then the quotient graph of groups arising from the action $\Lambda \curvearrowright \mathcal{T}$ is isomorphic to the graph of groups obtained by attaching a copy of $\mathbb{Z}$ to every vertex and every edge of $\mathcal{G}$, with structural maps of positive edges*

$$\mathbb{Z}_e \hookrightarrow \mathbb{Z}_{\mathtt{s}(e)}, \quad k \mapsto \frac{n \cdot L(e)}{L(\mathtt{s}(e))} \cdot k \quad \textit{and} \quad \mathbb{Z}_e \hookrightarrow \mathbb{Z}_{\mathtt{t}(e)}, \quad k \mapsto \frac{m \cdot L(e)}{L(\mathtt{t}(e))} \cdot k.$$

Then, combining Proposition 3.34 and Lemma 3.33, we get the following rephrasing of Corollary 2.15.

Corollary 3.35. *Let m and n be non-zero integers such that $|m| \neq |n|$. Then the isomorphism type of $\Lambda \leq \mathrm{BS}(m, n)$ depends only on the graph structure of $\mathbf{BS}(\Lambda)$.*

Proof. Recall that if an (m, n)-graph is saturated and has only infinite labels, then all vertices have incoming degree $|m|$ and outgoing degree $|n|$. Lemma 3.33 thus allows us to detect whether the Bass Serre graph of Λ contains infinite labels by purely looking at its graph structure: it has infinite labels if and only if all vertices have degree $|n| + |m|$. The result now follows from Proposition 3.34. ■

Remark 3.36. When $|m| = |n|$, the statement analogue to that of Corollary 3.35 fails since the central subgroup $\Lambda = \langle b^{2n} \rangle$ has the same Bass–Serre graph as the trivial subgroup $\{\mathrm{id}\}$.

4. Phenotype

In this section, we introduce a central invariant to understand transitive $\mathrm{BS}(m,n)$-(pre)-actions: the *phenotype* (see Definition 4.9). We first define the (m,n)-phenotype of a natural number. We then prove that given a transitive pre-action (τ, β), all cardinalities of β-orbits have the same phenotype.

4.1. Phenotypes of natural numbers

Recall that $\mathcal{P}$ denotes the set of prime numbers and that given $p \in \mathcal{P}$ and $k \in \mathbb{Z}$, we denote by $|k|_p$ the p-adic valuation of k.

Definition 4.1 (Phenotype of a natural number). Let $k \in \mathbb{Z}_{\geq 1}$. We set

$$\mathcal{P}_{m,n} := \{p \in \mathcal{P} : |m|_p = |n|_p\},$$
$$\mathcal{P}_{m,n}(k) := \{p \in \mathcal{P} : |m|_p = |n|_p \text{ and } |k|_p > |n|_p\}.$$

The (m,n)-*phenotype* of k, denoted by $\mathrm{Ph}_{m,n}(k)$, is the following positive integer:

$$\mathrm{Ph}_{m,n}(k) := \prod_{p \in \mathcal{P}_{m,n}(k)} p^{|k|_p}.$$

If $k = \infty$, we set $\mathrm{Ph}_{m,n}(k) := \infty$.

Example 4.2. If m and n are coprime, then for every $k \in \mathbb{Z}$,

$$\mathcal{P}_{m,n} = \{p \in \mathcal{P} : p \text{ does not divide } mn\},$$
$$\mathcal{P}_{m,n}(k) = \{p \in \mathcal{P} : p \text{ divides } k \text{ and } p \text{ does not divide } mn\}.$$

In this case, $\mathrm{Ph}_{m,n}(k)$ is the greatest divisor of k that is coprime to mn.

Example 4.3. If $m = 2^2 \cdot 3^2 \cdot 5$ and $n = 2^2 \cdot 3$, then $\mathcal{P}_{m,n} = \mathcal{P} \setminus \{3,5\}$ and

$$\mathcal{P}_{m,n}(k) = \begin{cases} \{p \in \mathcal{P} \colon p \text{ divides } k\} \setminus \{2,3,5\} & \text{if } 2^3 \text{ does not divide } k, \\ \{p \in \mathcal{P} \colon p \text{ divides } k\} \setminus \{3,5\} & \text{if } 2^3 \text{ divides } k. \end{cases}$$

For example, $\mathrm{Ph}_{m,n}(2 \cdot 3 \cdot 7) = 7$ and $\mathrm{Ph}_{m,n}(2^5 \cdot 3 \cdot 7) = 2^5 \cdot 7$.

Remark 4.4. If k and l both have phenotype q, then so do their lcm and gcd.

The following lemma will be useful in Section 5.

Lemma 4.5. *Let $q = \mathrm{Ph}_{m,n}(k)$ be a finite (m,n)-phenotype. Then $\mathrm{Ph}_{m,n}^{-1}(\{q\})$ is finite if and only if $|m| = |n|$.*

Proof. Assume first that $|m| \neq |n|$. In this case, there is a prime number p such that $|m|_p \neq |n|_p$. We get $\mathrm{Ph}_{m,n}(p^i k) = q$ for all i, hence $\mathrm{Ph}_{m,n}^{-1}(\{q\})$ is infinite.

If $|m| = |n|$, then $\mathcal{P}_{m,n} = \mathcal{P}$. If k and k' are two integers with the same phenotype, the only primes p for which the valuations of k and k' may differ are those for which $|k|_p \leq |m|_p$ and in this case $|k'|_p$ must also be bounded by $|m|_p$. There are only finitely many such k'. ∎

4.2. Phenotypes of (m,n)-graphs

If v is a vertex of an (m,n)-graph, we use the shorter expression "phenotype of the vertex v" to mean "phenotype of the label of the vertex v". The key feature of the notion of phenotype is the following statement.

Proposition 4.6. *Given a connected (m,n)-graph, all its vertices have the same (m,n)-phenotype.*

Proof. It is enough to check that for any positive edge e from v_- to v_+, the phenotypes of v_- and v_+ are the same. If the phenotype of one of them is infinite, then this is a direct consequence of equation (3.13) from Definition 3.12. Otherwise, remark that for every positive integer k and every $p \in \mathcal{P}_{m,n}$,

$$\left|\frac{k}{\gcd(k,n)}\right|_p > 0 \iff p \in \mathcal{P}_{m,n}(k).$$

Equation (3.13) implies

$$\left|\frac{L(v_-)}{\gcd(L(v_-),n)}\right|_p = |L(e)|_p = \left|\frac{L(v_+)}{\gcd(L(v_+),m)}\right|_p,$$

and hence $\mathcal{P}_{m,n}(L(v_-)) = \mathcal{P}_{m,n}(L(v_+))$. If $p \in \mathcal{P}_{m,n}(L(v_-))$, then $L(v_-)$ has higher p-valuation than m and n, so

$$|L(v_-)|_p - |n|_p = \left|\frac{L(v_-)}{\gcd(L(v_-),n)}\right|_p = \left|\frac{L(v_+)}{\gcd(L(v_+),m)}\right|_p = |L(v_+)|_p - |m|_p.$$

Since $|n|_p = |m|_p$, we conclude that for all $p \in \mathcal{P}_{m,n}(L(v_-)) = \mathcal{P}_{m,n}(L(v_+))$, we have $|L(v_-)|_p = |L(v_+)|_p$. Therefore, $L(v_-)$ and $L(v_+)$ share the same phenotype. ∎

Remark 4.7. One can prove that the edges of a connected (m,n)-graph also all have the same (m,n)-phenotype. However, it is a coarser invariant: there are connected graphs with different vertex phenotypes, but with the same edge phenotype. For example, fix

$$m = 2^2 \cdot 3^2 \cdot 5, \quad n = 2^2 \cdot 3$$

and consider the graph consisting of a single oriented edge e and its two endpoints. If the label of its origin is $L(\mathrm{s}(e)) = 2^3 \cdot 7$, then

$$L(e) = \frac{L(\mathrm{s}(e))}{\gcd(L(\mathrm{s}(e)),n)} = 2 \cdot 7 \quad \text{and} \quad \mathrm{Ph}_{m,n}(L(e)) = 7,$$

while $\mathrm{Ph}_{m,n}(L(\mathtt{s}(e)) = 2^3 \cdot 7$. If instead we set the label of its origin to be $L(\mathtt{s}(e)) = 2^4 \cdot 7$, then we get

$$L(e) = 2^2 \cdot 7 \quad \text{and} \quad \mathrm{Ph}(L(e)) = 7,$$

while $\mathrm{Ph}_{m,n}(L(\mathtt{s}(e)) = 2^4 \cdot 7 \neq 2^3 \cdot 7$. We will thus not use the phenotype of edges.

Proposition 4.6 allows us to define the phenotypes of connected (m, n)-graphs and transitive $\mathrm{BS}(m, n)$-pre-actions.

Definition 4.8. The *phenotype of a connected* (m, n)*-graph* $\mathcal{G}$ is the common phenotype of the labels of its vertices. We denote it $\mathbf{Ph}(\mathcal{G})$.

4.3. Phenotypes of BS(m, n)-actions

Recall that a pre-action is transitive if its Schreier graph is connected, which is equivalent to its Bass–Serre graph being connected.

Definition 4.9. The *phenotype of a transitive (pre)-action* α *of* $\mathrm{BS}(m, n)$ is the common phenotype of the cardinalities $\mathrm{Ph}_{m,n}(|x\langle b\rangle|)$ of its $\langle b\rangle$-orbits. We denote it $\mathbf{Ph}(\alpha)$.

By definition, the phenotype of any transitive (pre)-action coincides with the phenotype of its Bass–Serre graph.

Remark 4.10. Any $\mathrm{BS}(m, n)$-action with finite Bass–Serre graph and finite phenotype is necessarily an action on a finite set whose cardinality is the sum of the labels of the vertices.

For infinite phenotype, we have the following.

Lemma 4.11. *There exists an infinite phenotype transitive* $\mathrm{BS}(m, n)$*-action with finite Bass–Serre graph if and only if* $|m| = |n|$.

Proof. Consider an infinite phenotype $\mathrm{BS}(m, n)$-action with finite Bass–Serre graph $\mathcal{G}$. Since $\mathcal{G}$ is saturated, all its vertices have outgoing degree $|n|$ and incoming degree $|m|$. But there must be globally as many outgoing edges as incoming edges, so since $\mathcal{G}$ is finite, we must have $|n| = |m|$.

Conversely, if $|n| = |m|$, consider the bouquet of $|n|$ circles with edges and vertices labeled by ∞, and observe that this is a connected saturated (m, n)-graph. Proposition 3.22 provides a transitive action having this labeled bouquet of circles as its finite Bass–Serre graph of infinite phenotype. ∎

4.4. Merging pre-actions

In order to establish some of the main results of this article, we will need "cut and paste" operations on pre-actions, for instance:

- putting two prescribed pre-actions inside a single transitive action (useful for topological transitivity properties),
- modifying an action so as to add or remove a circuit in its Schreier graph (useful to get a new action that is close but distinct from the original one).

We now present these “cut and paste” operations. The main one is the following and the rest of this section will be devoted to its proof. Other useful results will appear in the course of the proof.

Theorem 4.12 (The merging machine). *Assume $|m| \geq 2$ and $|n| \geq 2$. Let α_1 and α_2 be two transitive non-saturated pre-actions of* $\mathrm{BS}(m,n)$ *with the same phenotype. There exists a transitive action α which contains copies of α_1 and α_2 with disjoint domains.*

Given a pre-action $\alpha = (\beta, \tau)$ and two sub-pre-actions α_1, α_2, let us recall that the domain of α is the set $\mathrm{dom}(\beta) = \mathrm{rng}(\beta)$. Notice that α_1 and α_2 have disjoint domains if and only if their Bass–Serre graphs $\mathbf{BS}(\alpha_1)$ and $\mathbf{BS}(\alpha_2)$ are disjoint (that is, have no common vertex) in $\mathbf{BS}(\alpha)$.

First, taking advantage of Proposition 3.23, we reduce to the case of (m,n)-graphs, for which the analogous result is the following.

Theorem 4.13 (The merging machine for (m,n)-graphs). *Assume $|m| \geq 2$ and $|n| \geq 2$. Let $\mathcal{G}_1$ and $\mathcal{G}_2$ be two connected and non-saturated (m,n)-graphs with the same phenotype. There exists a connected and saturated (m,n)-graph $\mathcal{G}$ which contains disjoint copies of $\mathcal{G}_1$ and $\mathcal{G}_2$.*

Remark 4.14. The hypothesis that both $|m|, |n| \geq 2$ is necessary. If $m = 1$ but $|n| \neq 1$, we can consider the $(1,n)$-graph consisting of a single vertex with infinite label and only one loop. This graph is not saturated but it cannot be connected to another copy of itself. Indeed, the reader can check that the only saturated graph containing it admits a unique circuit, namely, the loop itself.

Proof of Theorem 4.12 based on Theorem 4.13. The two Bass–Serre graphs $\mathbf{BS}(\alpha_1)$ and $\mathbf{BS}(\alpha_2)$ are connected non-saturated (m,n)-graphs with the same phenotype. Therefore, we can apply Theorem 4.13 to obtain a connected and saturated (m,n)-graph $\mathcal{G}$ which contains disjoint copies of $\mathbf{BS}(\alpha_1)$ and $\mathbf{BS}(\alpha_2)$.

Then we apply Proposition 3.23 to the pre-action $\alpha_1 \sqcup \alpha_2$, whose Bass–Serre graph $\mathbf{BS}(\alpha_1) \sqcup \mathbf{BS}(\alpha_2)$ is contained in $\mathcal{G}$, to ensure the existence of a $\mathrm{BS}(m,n)$-pre-action α which extends $\alpha_1 \sqcup \alpha_2$. Thus α extends both α_1 and α_2 with disjoint domains. Since $\mathcal{G}$ is connected and saturated, α is a transitive and saturated pre-action, i.e., it is a genuine transitive action of $\mathrm{BS}(m,n)$ that satisfies the requirements of Theorem 4.12. ∎

We now present some general results we will use in order to prove Theorem 4.13. We begin with two easy properties of phenotypes which will be useful in the proof.

Lemma 4.15. *For any $k \in \mathbb{Z}_{\geq 1}$, if $q = \mathrm{Ph}_{m,n}(k)$, then $\mathrm{Ph}_{m,n}(q) = q$ and $\gcd(q,n) = \gcd(q,m)$.*

Proof. We get directly from Definition 4.1 that $|q|_p = |k|_p$ if $p \in \mathcal{P}_{m,n}(k)$, and $|q|_p = 0$ for the other primes p. Consequently, we get $\mathcal{P}_{m,n}(q) = \mathcal{P}_{m,n}(k)$ and then $\mathrm{Ph}_{m,n}(q) = \mathrm{Ph}_{m,n}(k) = q$. Finally, since every prime p dividing q satisfies $|m|_p = |n|_p$ and $|n|_p < |q|_p$, we obtain

$$\gcd(q,n) = \prod_{p \in \mathcal{P} : p \mid q} p^{|n|_p} = \prod_{p \in \mathcal{P} : p \mid q} p^{|m|_p} = \gcd(q,m). \qquad ∎$$

In the following lemma, by welding two vertices we mean taking the quotient graph obtained by identifying these vertices. Its proof is a direct consequence of the definition of an (m,n)-graph, so we omit it.

Lemma 4.16 (Welding lemma). *Let $m,n \in \mathbb{Z} \setminus \{0\}$, let $\mathcal{G}$ be an (m,n)-graph and let v and w be two distinct vertices such that*

- $L := L(v) = L(w)$,
- $\deg_{\text{out}}(v) + \deg_{\text{out}}(w) \leq \gcd(n, L)$,
- $\deg_{\text{in}}(v) + \deg_{\text{in}}(w) \leq \gcd(m, L)$.

Welding together v and w delivers an (m,n)-graph.

Note that in this lemma $\mathcal{G}$ can be finite or infinite, connected or not. Together with the welding lemma, the following result will allow us to connect not saturated (m,n)-graphs via the well-known technique of arc welding.

Theorem 4.17 (Connecting lemma). *Assume $|m| \geq 2$ and $|n| \geq 2$. Let $k, \ell \in \mathbb{Z}_{\geq 1}$ be such that $\mathrm{Ph}_{m,n}(k) = \mathrm{Ph}_{m,n}(\ell)$, and let $\varepsilon_k, \varepsilon_\ell \in \{+,-\}$. There exists an (m,n)-graph $\mathcal{G}$, which is a simple edge path $(e_1, \ldots, e_h)$ of length $h \geq 1$, such that*

- *$L(\mathtt{s}(e_1)) = k$ and $L(\mathtt{t}(e_h)) = \ell$,*
- *the orientations of e_1 and e_h are given by $e_1 \in E(\mathcal{G})^{\varepsilon_k}$ and $e_h \in E(\mathcal{G})^{\varepsilon_\ell}$.*

Proof. Observe that every (m,n)-graph can be turned into an (n,m)-graph by flipping the orientations of its edges. Note that this operation does not affect the labels nor its phenotype. We thus can restrict ourselves to the case where the orientation ε_k of the first edge in the path is asked to be positive and no assumption is made on ε_ℓ. Let us set $q := \mathrm{Ph}_{m,n}(k) = \mathrm{Ph}_{m,n}(\ell)$.

We first treat the case $k = q = \ell$. Recall from Lemma 4.15 that $\mathrm{Ph}_{m,n}(q) = q$ and that we have $\gcd(m,q) = \gcd(n,q)$. Hence, there exists an (m,n)-graph with two vertices and a unique positive edge f_1 such that

$$L(\mathtt{s}(f_1)) = q = L(\mathtt{t}(f_1)) \quad \text{and} \quad L(f_1) = \frac{q}{\gcd(m,q)} = \frac{q}{\gcd(n,q)}.$$

If ε_ℓ is positive, then we are done. If not, create a vertex v with label $L(v) = \frac{q}{\gcd(n,q)} m$. We get $\gcd(m, L(v)) = |m|$, hence $\gcd(m, L(v)) \geq 2$. Therefore, we can equip v with two distinct incoming positive edges f_1 and f_2. Such edges have to be labeled by

$$\frac{L(v)}{\gcd(m, L(v))} = \frac{q}{\gcd(n,q)},$$

so that we can label $\mathtt{s}(f_1)$ and $\mathtt{s}(f_2)$ by q, and $(f_1, \bar{f_2})$ is the path we are looking for. The theorem is thus proved for $k = \ell = q$.

Let us now treat the case $k \neq q$ and $\ell = q$. Recall that $\mathcal{P}_{m,n}(k) = \{p \in \mathcal{P} : |m|_p = |n|_p$ and $|n|_p < |k|_p\}$ and that $\mathrm{Ph}_{m,n}(k) = \prod_{p \in \mathcal{P}_{m,n}(k)} p^{|k|_p}$. Thus, any number $L \in \mathbb{Z}_{\geq 1}$ with phenotype q admits a unique decomposition as follows:

$$L = q \cdot \prod_{\substack{p \in \mathcal{P} \setminus \mathcal{P}_{m,n}(k) \\ |m|_p \leq |n|_p}} p^{|L|_p} \prod_{\substack{p \in \mathcal{P} \\ |m|_p > |n|_p}} p^{|L|_p}. \tag{4.18}$$

In a first step, we construct (algorithmically) a simple path, consisting of positive edges with vertices $v_0, v_1, \dots, v_r$, such that v_0 has label k, and such that the decomposition of $L(v_r)$ reduces to

$$L(v_r) = q \cdot \prod_{p \in \mathcal{P} : |m|_p > |n|_p} p^{|L(v_r)|_p}, \tag{4.19}$$

that is, such that $|L(v_r)|_p = 0$ whenever $|m|_p \leq |n|_p$ and $p \notin \mathcal{P}_{m,n}(k)$.

To do so, starting with $i = 0$ and $L(v_0) = k$, while $L(v_i)$ has prime divisors p such that $|m|_p \leq |n|_p$ and $p \notin \mathcal{P}_{m,n}(k)$, we connect v_i to a new vertex v_{i+1} by a positive edge f_i. According to Remark 3.26, we label f_i by $|L(f_i)|_p := \max(|L(v_i)|_p - |n|_p, 0)$ and set

$$|L(v_{i+1})|_p := \begin{cases} |L(f_i)|_p + |m|_p & \text{if } |L(f_i)|_p \geq 1, \\ 0 & \text{if } |L(f_i)|_p = 0, \end{cases}$$

for every prime p. Then we replace i by $i + 1$, which terminates the “while” loop. Notice that we exit from the loop after finitely many steps. Indeed, given a prime p such that $|m|_p \leq |n|_p$ and $p \notin \mathcal{P}_{m,n}(k)$, we have:

- either $|L(f_1)|_p = 0$ in the case $|m|_p = |n|_p$ and $|k|_p \leq |n|_p$, which implies that $|L(v_i)|_p = 0$ for all $i \geq 1$,
- or $|L(v_{i+1})|_p = |L(v_i)|_p - |n|_p + |m|_p < |L(v_i)|_p$ whenever $|L(v_i)|_p \geq 1$ in the case $|m|_p < |n|_p$.

When we exit the “while” loop, Remark 3.26 guarantees that we have constructed an (m, n)-graph, and the loop condition ensures that the last vertex v_r satisfies $|L(v_r)|_p = 0$ whenever $|m|_p \leq |n|_p$ and $p \notin \mathcal{P}_{m,n}(k)$.

If we are lucky, we have $L(v_r) = q$. If not, in a second step, we notice that the same algorithm, exchanging the roles of m and n, produces a simple path consisting of negative edges from a vertex w_0 such that $L(w_0) = L(v_r)$ to a vertex w_s labeled by q. The decomposition (4.19) of $L(v_r) \neq q$ also shows that $\gcd(m, L(v_r)) \geq 2$, so vertices labeled $L(v_r)$ can have two distinct positive incoming edges. Using Lemma 4.16, we weld v_r and w_0 together and get a simple path from v_0 to w_s.

In any subcase, we now have a path $(e_1, \dots, e_{h'})$ such that e_1 is positive, $L(\mathtt{s}(e_1)) = k$, and $L(\mathtt{t}(e_{h'})) = q$. If $e_{h'}$ has the orientation prescribed by ε_ℓ, we are done; if not, using the case $k = q = \ell$, with the first edge having the same orientation as $e_{h'}$, and the last one having the orientation prescribed by ε_ℓ, we extend our path to a simple path $(e_1, \dots, e_h)$ with $L(\mathtt{s}(e_1)) = k$ and $L(\mathtt{t}(e_h)) = q$ such that e_1, e_h have the correct orientations. This concludes the case $\ell = q$ and $k \neq q$.

The case $k = q$ and $\ell \neq q$ is obtained by exchanging the roles of k and l in the above argument. Therefore, let us finally treat the case $k \neq q$ and $\ell \neq q$. The former cases furnish paths $(f_1, \dots, f_r)$ and $(f'_1, \dots, f'_s)$, that we may assume disjoint, such that

$$L(\mathtt{s}(f_1)) = k, \quad L(\mathtt{t}(f_r)) = q = L(\mathtt{s}(f'_1)), \quad L(\mathtt{t}(f'_s)) = \ell,$$

the orientations of f_1 and f'_s are given by ε_k and ε_ℓ respectively, and the orientations f_r and f'_1 coincide. Then we just weld the vertices $\mathtt{t}(f_r)$ and $\mathtt{s}(f'_1)$ together, and the path $(f_1, \dots, f_r, f'_1, \dots, f'_s)$ is as desired. ■

Remark 4.20. In Theorem 4.17, the assumption $|m| \geq 2$ and $|n| \geq 2$ is necessary. Indeed, Theorem 4.17 would be false for $n = 1$. If v is a vertex in a $(m,1)$-graph with $L(v) = 1$ and e is an edge such that $\mathtt{t}(e) = v$, then

$$1 = L(\mathtt{t}(e)) = \frac{L(\mathtt{t}(e))}{\gcd(L(\mathtt{t}(e)), m)} = \frac{L(\mathtt{s}(e))}{\gcd(L(\mathtt{s}(e)), 1)} = L(\mathtt{s}(e)).$$

Clearly any vertex with label 1 has at most one outgoing and one incoming edge. This implies that the labels of the vertices in any directed path which ends in v must be all 1. In other words, if we have any simple edge path as in Theorem 4.17 such that $\ell = 1$ and $\varepsilon_\ell = -$, then we must have that $k = 1$ (and $\varepsilon_k = +$).

Definition 4.21. Let $\mathcal{G}$ be a connected (m,n)-graph. A saturated extension $\mathcal{G}'$ of $\mathcal{G}$ is called a *forest-saturation* of $\mathcal{G}$ if it satisfies:

- the subgraph induced in $\mathcal{G}'$ by $V(\mathcal{G})$ is exactly $\mathcal{G}$,
- the subgraph induced in $\mathcal{G}'$ by $V(\mathcal{G}') \setminus V(\mathcal{G})$ is a forest $\mathcal{F}$,
- each connected component of $\mathcal{F}$ is connected to $\mathcal{G}$ by a single edge of $\mathcal{G}'$.

Lemma 4.22 (Forest-saturation lemma). *Let $\mathcal{G}$ be a connected (m,n)-graph. There is a forest-saturation $\mathcal{G}'$ of $\mathcal{G}$ such that all vertices of the forest $\mathcal{F}$ induced in $\mathcal{G}'$ by $V(\mathcal{G}') \setminus V(\mathcal{G})$ have degree $\geq 1 + \min(|m|, |n|)$ in $\mathcal{G}'$.*

The reader can observe in the following construction proving Lemma 4.22 that, while the labels of the new edges are prescribed, the axioms of (m,n)-graphs allows some choices concerning the labels of the new vertices. The systematic choice of the maximal label will be made for the new vertices among all those satisfying the transfer equation (3.13), that is, $L(\mathtt{s}(e))/\gcd(L(\mathtt{s}(e)), n) = L(e) = L(\mathtt{t}(e))/\gcd(L(\mathtt{t}(e)), m)$. Hence, the forest-saturation constructed in this proof is called the *maximal forest-saturation* of $\mathcal{G}$. Notice that other choices would have led to forest-saturations with different underlying graphs, by virtue of the relationship between labels and degrees (see Definition 3.20). These forest-saturations are further studied in the recent preprint [16].

Proof of Lemma 4.22. We can assume that the connected graph $\mathcal{G}$ is not yet saturated: it admits non-saturated vertices i.e., vertices v for which one of the inequalities in (3.14), $\deg_{\text{out}}(v) \leq \gcd(L(v), n)$ or $\deg_{\text{in}}(v) \leq \gcd(L(v), m)$, is strict. For every non-saturated vertex v of $\mathcal{G}$, we add:

- $(\gcd(L(v), n) - \deg_{\text{out}}(v))$-many new v-outgoing edges labeled $L_{\text{out}} := \frac{L(v)}{\gcd(n, L(v))}$ with extra target vertices labeled mL_{out}, and
- $(\gcd(L(v), m) - \deg_{\text{in}}(v))$-many new v-incoming edges labeled $L_{\text{in}} := \frac{L(v)}{\gcd(m, L(v))}$ with extra source vertices labeled nL_{in}.

We then iterate this construction. All the non-saturated vertices of the j-th step become saturated at the $(j+1)$-th one. The increasing union $\mathcal{G}'$ of these (m,n)-graphs is a saturated (m,n)-graph. The complement of $\mathcal{G}$ in it is a forest since at each step, each new edge has a new vertex as one of its vertices. The label of each new vertex v is an integer multiple of either m or n. Thus, the degree $\deg_{\text{out}}(v) + \deg_{\text{in}}(v) = \gcd(L(v), n) + \gcd(L(v), m)$ of v is larger than $1 + \min(|m|, |n|)$ as expected. ∎

Proof of Theorem 4.13. By hypothesis, for $i = 1, 2$, there is a non-saturated vertex v_i in $\mathcal{G}_i$, that is, a vertex for which one of the inequalities (3.14) is strict. If $\deg_{\mathrm{in}}(v_i) < \gcd(L(v_i), m)$, then let $\epsilon_i := +$; otherwise, let $\epsilon_i := -$. The labels of v_1 and v_2 having identical phenotypes, the connecting lemma (Theorem 4.17) furnishes an (m,n)-graph $\mathcal{G}_0$ which is a simple edge path $(e_1, \ldots, e_h)$ such that $L(\mathtt{s}(e_1)) = L(v_1)$ and $L(\mathtt{t}(e_h)) = L(v_2)$, and the orientations of e_1 and e_h are given by $-\epsilon_1$ and ϵ_2, respectively.

We then consider the disjoint union $\mathcal{G}_1 \sqcup \mathcal{G}_0 \sqcup \mathcal{G}_2$. We claim that we can merge the vertices v_1 and $\mathtt{s}(e_1)$ thanks to the welding Lemma 4.16. Indeed, the choice of orientation for e_1 and the form of $\mathcal{G}_0$ (a path of edges) are made for the assumptions of Lemma 4.16 to hold. Then we can merge v_2 and $\mathtt{t}(e_h)$, applying Lemma 4.16 again (this time, using the fact that the orientation of e_h is well chosen). This produces a connected (m,n)-graph $\mathcal{G}_3$ which contains disjoint copies of $\mathcal{G}_1$ and $\mathcal{G}_2$.

It now suffices to apply the saturation Lemma 4.22 to $\mathcal{G}_3$ so as to obtain a connected saturated (m,n)-graph $\mathcal{G}$ that satisfies the requirements of Theorem 4.13. ∎

5. Perfect kernel and dense orbits

5.1. Perfect kernels of Baumslag–Solitar groups

In the case $|m| = 1$ or $|n| = 1$, it follows from the proof of Corollary 8.4 in [3] that $\mathrm{Sub}(\mathrm{BS}(m,n))$ is countable, hence the perfect kernel $\mathcal{K}(\mathrm{BS}(m,n))$ is empty. Our main theorem describes the perfect kernels in the remaining cases.

Theorem 5.1. *Let $m, n \in \mathbb{Z}$ with $|m| \geq 2$ and $|n| \geq 2$. We have*

$$\mathcal{K}(\mathrm{BS}(m,n)) = \{\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n)) : \Lambda\backslash \mathrm{BS}(m,n)/\langle b\rangle \text{ is infinite}\}.$$

Let us temporarily give a name to the set appearing in Theorem 5.1:

$$\mathcal{J} = \mathcal{J}(m,n) := \{\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n)) : \Lambda\backslash \mathrm{BS}(m,n)/\langle b\rangle \text{ is infinite}\},$$

and recall that $\mathrm{Sub}_{[\infty]}(\Gamma)$ denotes the space of infinite index subgroups of Γ.

Given an action α of Γ on a space X and a point $v \in X$, we have already introduced the notation $[\alpha, v]$ for the action α *pointed* at v.

Remark 5.2. In terms of pointed transitive actions, $\mathcal{J}(m,n)$ is the set of pointed transitive actions with infinitely many b-orbits, whence $\mathcal{J} = \{[\alpha, v] \colon \mathbf{BS}(\alpha) \text{ is infinite}\}$. Moreover,

- if $|m| \neq |n|$, we have $\mathcal{J}(m,n) = \mathrm{Sub}_{[\infty]}(\mathrm{BS}(m,n))$, since every infinite action has an infinite Bass–Serre graph by Lemma 4.11;
- if $|m| = |n|$, we have $\mathcal{J}(m,n) = \pi^{-1}(\mathrm{Sub}_{[\infty]}(\mathrm{BS}(m,n)/\langle b^m\rangle))$, where π is the homomorphism from $\mathrm{BS}(m,n)$ to its quotient by the normal subgroup $\langle b^m\rangle = \langle b^n\rangle$. Indeed, since $\langle b^m\rangle$ has finite index in $\langle b\rangle$, we get that $\Lambda\backslash \mathrm{BS}(m,n)/\langle b\rangle$ is finite if and only if $\Lambda\backslash \mathrm{BS}(m,n)/\langle b^m\rangle$ is finite.

Therefore, Theorem 5.1 can be rephrased in two ways, as follows.

Theorem 5.3. *Let $m, n \in \mathbb{Z}$ with $|m| \geq 2$ and $|n| \geq 2$.*

(1) *In terms of pointed transitive actions, the perfect kernel corresponds exactly to actions whose Bass–Serre graph is infinite*:

$$\mathcal{K}(\mathrm{BS}(m,n)) = \{[\alpha, v] : \mathbf{BS}(\alpha) \textit{ is infinite}\}.$$

(2) *In terms of subgroups*:

- *if $|m| \neq |n|$, the perfect kernel is equal to the space of infinite index subgroups*

$$\mathcal{K}(\mathrm{BS}(m,n)) = \mathrm{Sub}_{[\infty]}(\mathrm{BS}(m,n)),$$

- *if $|m| = |n|$, we have*

$$\mathcal{K}(\mathrm{BS}(m,n)) = \pi^{-1}(\mathrm{Sub}_{[\infty]}(\mathrm{BS}(m,n)/\langle b^m \rangle)),$$

where π is the homomorphism from $\mathrm{BS}(m,n)$ *to its quotient by the normal subgroup $\langle b^m \rangle = \langle b^n \rangle$.*

Proof of Theorem 5.1. Our aim is to prove that $\mathcal{K}(\mathrm{BS}(m,n)) = \mathcal{J}(m,n)$. It will be convenient to write one inclusion in terms of pointed transitive actions and the other in terms of subgroups.

Let us first prove the inclusion $\mathcal{K}(\mathrm{BS}(m,n)) \supseteq \mathcal{J}$. It suffices to show that no element of $\mathcal{J}$ is isolated in $\mathcal{J}$. Recall the definition of the topology in terms of pointed actions, see Section 2.2 and, in particular, equation (2.7). Let us fix a pointed transitive action (α_0, v) representing an element of $\mathcal{J}$ and a radius $R \geq 0$. We will show that the basic neighborhood $\mathcal{N}([\alpha_0, v], R)$ contains at least two distinct elements of $\mathcal{J}$.

Let (β, τ) be the pre-action obtained by restricting α_0 to the union of the b-orbits of the vertices of the ball of radius $R+1$ centered at v in the Schreier graph of α_0. The Bass–Serre graph of (β, τ) is the projection in $\mathbf{BS}(\alpha_0)$ (see Definition 3.10) of this ball, hence it is finite. Since $\mathbf{BS}(\alpha_0)$ is infinite, the pre-action (β, τ) is not saturated.

We now build two (m,n)-graphs $\mathcal{G}_1, \mathcal{G}_2$ that extend the finite non-saturated Bass–Serre graph $\mathcal{G}$ of (β, τ) in two different ways. First, let $\mathcal{G}_1$ be a forest-saturation of $\mathcal{G}$ given by Lemma 4.22. In particular, the subgraph induced in $\mathcal{G}_1$ by $V(\mathcal{G}_1) \setminus V(\mathcal{G})$ is a forest whose vertices have degree at least $3 \leq 1 + \min(|m|, |n|)$ in $\mathcal{G}_1$.

We then construct $\mathcal{G}_2$ by modifying $\mathcal{G}_1$. Let us pick a vertex $v \in V(\mathcal{G}_1) \setminus V(\mathcal{G})$. The subgraph induced in $\mathcal{G}_1$ by $V(\mathcal{G}_1) \setminus \{v\}$ has at least 3 connected components. Choose two connected components disjoint from $\mathcal{G}$ and remove them. In the resulting (m,n)-graph $\mathcal{G}_1'$, the vertex v is the only one that is not saturated: two edges are missing.

Theorem 4.17 gives us an (m,n)-graph, which is a simple edge path $\mathcal{P}$ whose extremities have the same label as v and for which the orientations of the end edges are compatible with that of the missing edges of v. We then apply twice the welding lemma, Lemma 4.16, so as to weld the two extremities of $\mathcal{P}$ to v. We eventually define $\mathcal{G}_2$ to be a forest-saturation of the graph that we obtained. Observe that $\mathcal{G}_1$ is not isomorphic to $\mathcal{G}_2$, since the fundamental groups of their underlying graphs are free groups of distinct ranks.

Finally, we extend (β, τ) to pre-actions α_1 and α_2, whose Bass–Serre graphs are $\mathcal{G}_1$ and $\mathcal{G}_2$, respectively, thanks to Proposition 3.23. Since $\mathcal{G}_1, \mathcal{G}_2$ are saturated, α_1, α_2 are actually actions by Example 3.21. We already remarked that $\mathcal{G}_1$ is not isomorphic to $\mathcal{G}_2$, so the pointed transitive actions (α_1, v) and (α_2, v) are not isomorphic: $[\alpha_1, v] \neq [\alpha_2, v]$. Moreover, the balls of radius R centered at the basepoints in the Schreier graphs of $\alpha_0, \alpha_1, \alpha_2$ all coincide by construction with that of (β, τ), so $[\alpha_1, v]$ and $[\alpha_2, v]$ are both in $\mathcal{N}([\alpha_0, v], R)$.

Let us now turn to the inclusion $\mathcal{K}(\mathrm{BS}(m,n)) \subseteq \mathcal{J}$. Let us pick a subgroup $\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n)) \setminus \mathcal{J}(m,n)$, and let us prove that it is not in the perfect kernel.

If $|m| \neq |n|$, then Λ has finite index in $\mathrm{BS}(m,n)$ by Remark 5.2, hence it is isolated in $\mathrm{Sub}(\mathrm{BS}(m,n))$.

If $|m| = |n|$, then $\pi(\Lambda)$ has finite index in $\mathrm{BS}(m,n)/\langle b^m \rangle$ by Remark 5.2, hence it is finitely generated. Therefore, the set

$$\mathcal{V} := \{\Lambda' \in \mathrm{Sub}(\mathrm{BS}(m,n)) : \pi(\Lambda') \geq \pi(\Lambda)\}$$

is a neighborhood of Λ, since it contains the basic neighborhood

$$\mathcal{V}(S, \varnothing) = \{\Lambda' \in \mathrm{Sub}(\mathrm{BS}(m,n)) \colon S \subseteq \Lambda'\},$$

where $S \subseteq \Lambda$ is a finite set such that $\pi(S)$ generates $\pi(\Lambda)$.

Now, for any $\Lambda' \in \mathcal{V}$, the subgroup $\pi(\Lambda')$ has finite index in $\mathrm{BS}(m,m)/\langle b^m \rangle$. Hence, $\pi(\Lambda')$ is finitely generated, so Λ' itself is finitely generated since it is written as an extension with cyclic kernel:

$$1 \to \langle b^m \rangle \cap \Lambda' \to \Lambda' \to \pi(\Lambda') \to 1.$$

Therefore, all subgroups of $\mathcal{V}$ are finitely generated, which implies that $\mathcal{V}$ is countable and hence Λ is not in $\mathcal{K}(\mathrm{BS}(m,n))$. ■

Corollary 5.4. *If $|m| \geq 2$, $|n| \geq 2$ and $|m| \neq |n|$, then*

$$\mathbf{Ph}^{-1}(\infty) \subseteq \mathcal{K}(\mathrm{BS}(m,n));$$

in other words, every infinite phenotype subgroup is in the perfect kernel.

Proof. Any subgroup with infinite phenotype has infinite index and hence it belongs to $\mathcal{K}(\mathrm{BS}(m,n))$ according to Theorem 5.3. ■

5.2. Phenotypical decomposition of the perfect kernel

Let us now turn to a description of the internal structure of $\mathcal{K}(\mathrm{BS}(m,n))$.

Notation 5.5. Let $m, n \in \mathbb{Z} \setminus \{-1, 0, 1\}$. We denote by $\mathcal{Q}_{m,n}$ the set of all possible (m,n)-phenotypes, that is, $\mathcal{Q}_{m,n} := \mathrm{Ph}_{m,n}(\mathbb{Z}_{\geq 1} \cup \{\infty\})$.

Definition 5.6. The phenotype of a subgroup $\Lambda \leq \mathrm{BS}(m,n)$ is the (m,n)-phenotype of the index of $\Lambda \cap \langle b \rangle$ in $\langle b \rangle$:

$$\mathbf{Ph}(\Lambda) = \mathbf{Ph}(\Lambda \cap \langle b \rangle) := \mathrm{Ph}_{m,n}([\langle b \rangle : \Lambda \cap \langle b \rangle]).$$

This yields a function $\mathbf{Ph} \colon \mathrm{Sub}(\mathrm{BS}(m,n)) \to \mathcal{Q}_{m,n} \subseteq \mathbb{Z}_{\geq 1} \cup \{\infty\}$.

In particular, $\mathbf{Ph}(\langle b^k \rangle) = \mathrm{Ph}_{m,n}(k)$ for $k \in \mathbb{Z}_{\geq 1}$, and the phenotype of the trivial subgroup is infinite.

Remark 5.7. The index $[\langle b\rangle : \Lambda \cap \langle b\rangle]$ is the cardinal of the $\langle b\rangle$-orbit of the point Λ in the action $\Lambda\backslash \mathrm{BS}(m,n) \curvearrowleft \mathrm{BS}(m,n)$. Hence, $\mathbf{Ph}(\Lambda)$ is the phenotype of this action (as given in Definition 4.9). Since the latter does not depend on the basepoint, the function **Ph** is invariant under conjugation.

It follows from the definitions that if $\mathbf{Ph}(\Lambda) = \mathbf{Ph}(\Lambda')$, then $\mathbf{Ph}(\Lambda) = \mathbf{Ph}(\Lambda \cap \Lambda')$, see Remark 4.4.

Proposition 5.8. *In the partition of the space of subgroups of* $\mathrm{BS}(m,n)$ *according to their phenotype*

$$\mathrm{Sub}(\mathrm{BS}(m,n)) = \bigsqcup_{q\in\mathcal{Q}_{m,n}} \mathbf{Ph}^{-1}(q),$$

the pieces are non-empty and satisfy:

(1) *For every finite* $q \in \mathcal{Q}_{m,n}$, *the piece* $\mathbf{Ph}^{-1}(q)$ *is open*; *it is also closed if and only if* $|m| = |n|$.

(2) *For* $q = \infty$, *the piece* $\mathbf{Ph}^{-1}(\infty)$ *is closed and not open.*

In particular, the function $\mathbf{Ph}\colon \mathrm{Sub}(\mathrm{BS}(m,n)) \to \mathbb{Z}_{\geq 1} \cup \{+\infty\}$ is Borel. It is continuous if and only if $|m| = |n|$.

Proof. Given $k \in \mathbb{Z}_{\geq 1}$, we set

$$A_k := \{\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n)) : \Lambda \cap \langle b\rangle = \langle b^k\rangle\}.$$

Writing A_k as

$$A_k = \{\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n) : b^k \in \Lambda,\ b^i \notin \Lambda \text{ for every } 1 \leq i < k\}$$

makes it clear that A_k is clopen for every $k \in \mathbb{Z}_{\geq 1}$. Moreover, $\langle b^k\rangle \in A_k$, so in particular A_k is not empty. Now, by definition, for every $q \in \mathbb{Z}_{\geq 1}$, we have

$$\mathbf{Ph}^{-1}(q) = \bigsqcup_{k\in \mathrm{Ph}_{m,n}^{-1}(q)} A_k. \tag{5.9}$$

Hence, $\mathbf{Ph}^{-1}(q)$ is open for every finite q and, by taking the complement, $\mathbf{Ph}^{-1}(\infty)$ is closed.

Take a sequence of positive integers $(k_i)_{i\in\mathbb{N}}$ tending to ∞. Observe that the subgroups $\{\langle b^{k_i}\rangle\}_i$ have finite phenotype and converge to the trivial subgroup which has infinite phenotype. Therefore, $\mathbf{Ph}^{-1}(\infty)$ is not open. Moreover, if $\mathrm{Ph}_{m,n}^{-1}(q)$ is not finite, we can choose all the k_i's with phenotype q; the same argument shows that $\mathbf{Ph}^{-1}(q)$ is not closed. Finally, the clopen decomposition (5.9) shows that $\mathbf{Ph}^{-1}(q)$ is closed as long as $\mathrm{Ph}_{m,n}^{-1}(q)$ is finite. By Lemma 4.5, $\mathrm{Ph}_{m,n}^{-1}(q)$ is finite exactly when $|m| = |n|$. ∎

We now restrict the above partition to the perfect kernel

$$\mathcal{K}(\mathrm{BS}(m,n)) = \bigsqcup_{q\in\mathcal{Q}_{m,n}} \mathcal{K}_q(\mathrm{BS}(m,n)), \tag{5.10}$$

where

$$\mathcal{K}_q(\mathrm{BS}(m,n)) := \mathcal{K}(\mathrm{BS}(m,n)) \cap \mathbf{Ph}_{m,n}^{-1}(q). \tag{5.11}$$

Remark 5.12. Observe that each $\mathcal{K}_q(\mathrm{BS}(m,n))$ is not empty. Indeed, it contains $\langle b^q\rangle$ which belongs to the perfect kernel by Theorem 5.1. Moreover, in the proof of Theorem 5.1, the (m,n)-graphs we construct have the same phenotype, so every element of $\mathcal{K}_q(\mathrm{BS}(m,n))$ is actually a non-trivial limit of elements of $\mathcal{K}_q(\mathrm{BS}(m,n))$. We conclude that $\mathcal{K}_q(\mathrm{BS}(m,n))$ is equal to the perfect kernel of $\mathbf{Ph}_{m,n}^{-1}(q)$.

Let us show that the action of $\mathrm{BS}(m,n)$ by conjugation on each term is *topologically transitive* in the following sense.

Definition 5.13. An action by homeomorphisms of a group Γ on a topological space X is called *topologically transitive* if for every nonempty open sets U and V, there is a point whose Γ-orbit intersects both U and V.

Theorem 5.14. *Let m,n be integers such that $|m|,|n|\geq 2$. Then for every phenotype $q\in\mathcal{Q}_{m,n}$, the action by conjugation of* $\mathrm{BS}(m,n)$ *on the invariant subspace* $\mathcal{K}_q(\mathrm{BS}(m,n))$ *is topologically transitive.*

Proof. We again use the definition of the topology in terms of pointed actions, see Section 2.2 and, in particular, equation (2.7). So let us fix two pointed actions (α_1,v_1) and (α_2,v_2) in $\mathcal{K}_q(\mathrm{BS}(m,n))$, take $R>0$, and consider the basic open sets $\mathcal{N}([\alpha_1,v_1],R)$ and $\mathcal{N}([\alpha_2,v_2],R)$. We need to construct a pointed action whose orbit meets both open sets.

As in the proof of Theorem 5.1, for $i=1,2$, we let (β_i,τ_i) be the pre-action obtained by restricting α_i to the union of the b-orbits of the vertices of the balls $B(v_i,R+1)$ of radius $R+1$ centered at v_i in the Schreier graph of α_i. The Bass–Serre graph of (β_i,τ_i) is finite. Since $\mathbf{BS}(\alpha_i)$ is infinite, the pre-action (β_i,τ_i) is not saturated.

Moreover, (β_1,τ_1) and (β_2,τ_2) have the same phenotype, so we can apply the merging machine (Theorem 4.12) to obtain an action α whose Schreier graph contains (copies of) the balls $B(v_i,R+1)$.

Pointing α at the copy of v_1 that we denote by v, we have $(\mathbf{Sch}(\alpha),v)\simeq_R(\mathbf{Sch}(\alpha_1),v_1)$. By transitivity of α, there is $\gamma\in\mathrm{BS}(m,n)$ such that $v\alpha(\gamma)$ is the copy of v_2, and thus $(\mathbf{Sch}(\alpha),v\alpha(\gamma))\simeq_R(\mathbf{Sch}(\alpha_2),v_2)$. In particular, the orbit of $[\alpha,v]$ meets both open sets $\mathcal{N}([\alpha_1,v_1],R)$ and $\mathcal{N}([\alpha_2,v_2],R)$. ∎

Corollary 5.15. *Let m,n be integers such that $|m|,|n|\geq 2$. Then for every $q\in\mathcal{Q}_{m,n}$, there is a dense G_δ subset of* $\mathcal{K}_q(\mathrm{BS}(m,n))$ *consisting of subgroups with dense conjugacy class in* $\mathcal{K}_q(\mathrm{BS}(m,n))$.

Proof of Corollary 5.15. By Proposition 5.8, each $\mathcal{K}_q(\mathrm{BS}(m,n))$ is Polish as an open or a closed subset of the Polish space $\mathcal{K}(\mathrm{BS}(m,n))$.

The corollary now follows from a well-known characterization of topological transitivity in Polish spaces: If (U_i) is a countable base of non-empty open subsets, then the set $\bigcap_{i\in\mathbb{N}}U_i\Gamma$ of points with dense orbit is a dense G_δ by the Baire theorem. ∎

5.3. Closed invariant subsets with a fixed finite phenotype

Given a finite phenotype q, we will show that there is a largest closed invariant subset inside the (open but not closed when $|m|\neq|n|$) set of subgroups of phenotype q. The following lemma is key.

Lemma 5.16. *Let $|m| \neq |n|$, and let $L \in \mathbb{Z}_{\geq 1}$ satisfying*:

$$\exists p \in \mathcal{P}, \quad |m|_p \neq |n|_p \quad \text{and} \quad |L|_p > \min(|m|_p, |n|_p).$$

Then for any saturated (m,n)-graph which contains L as a label, the range of the label map is unbounded.

Proof. By symmetry, we may as well assume that $|n|_p < |m|_p$ for a fixed prime p, and so $|L|_p > |n|_p$. Let $v_0 \in V(\mathcal{G})$ have label L. Since our Bass–Serre graph $\mathcal{G}$ is saturated, every vertex has at least one outgoing edge. We can thus build inductively an infinite path $(e_k)_{k\in\mathbb{N}}$ consisting of positive edges with $\mathtt{s}(e_0) = v_0$. The conclusion then follows directly from Lemma 3.30. ∎

Remark 5.17. When $|n| = |m|$, the lemma fails because labels are bounded. Indeed, if L_0 is a label, then all labels in the same connected component must satisfy the inequality $|L|_p \leq \max(|L_0|_p, |m|_p, |n|_p)$ for all primes p by equation (3.29) and the discussion that precedes it.

Let q be a finite (m,n)-phenotype. In order to describe which saturated (m,n)-graphs have unbounded labels, we now define

$$s(q,m,n) := q \cdot \prod_{\substack{p\in\mathcal{P} \\ |q|_p=0; \\ |m|_p=|n|_p>0}} p^{|m|_p} \cdot \prod_{\substack{p\in\mathcal{P} \\ |m|_p\neq|n|_p}} p^{\min\{|n|_p,|m|_p\}}. \tag{5.18}$$

Remark 5.19. The definition is motivated by the fact that $s(q,m,n)$ is the largest label of phenotype q which does not satisfy the hypothesis of Lemma 5.16. As we will see in the proof of Theorem 5.20, a saturated (m,n)-graph with phenotype q has unbounded labels if and only if one of its labels does not divide $s(q,m,n)$.

Proposition 5.8 implies that every subgroup (or pointed action) that lies in the closure of the set of subgroups of phenotype q has phenotype either q or ∞, and phenotype ∞ can occur only when $|m| \neq |n|$. We can now characterize the subgroups Λ with phenotype q whose orbit approaches subgroups with infinite phenotype.

Theorem 5.20. *Let m,n be integers such that $|m|,|n| \geq 2$ and denote by $q \in \mathcal{Q}_{m,n} \setminus \{\infty\}$ a finite (m,n)-phenotype. Let $s = s(q,m,n)$ be as in equation (5.18). Then the space*

$$\mathcal{MC}_q := \mathbf{Ph}^{-1}(q) \cap \{\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n)) : \Lambda \geq \langle\langle b^s \rangle\rangle\}$$

of subgroups of phenotype q containing the normal subgroup $\langle\langle b^s \rangle\rangle$ satisfies the following properties:

1. *$\mathcal{MC}_q$ is the largest closed $\mathrm{BS}(m,n)$-invariant subset of $\mathrm{Sub}(\mathrm{BS}(m,n))$ contained in $\mathbf{Ph}^{-1}(q)$; in particular, all normal subgroups of phenotype q and all finite index subgroups of phenotype q contain $\langle\langle b^s \rangle\rangle$.*
2. *If $|m| = |n|$, then $\mathcal{MC}_q = \mathbf{Ph}^{-1}(q)$.*
3. *For every $\Lambda \in \mathbf{Ph}^{-1}(q) \setminus \mathcal{MC}_q$, the orbit of Λ accumulates to $\mathbf{Ph}^{-1}(\infty)$.*
4. *If $|m| \neq |n|$, then $\mathcal{MC}_q \cap \mathcal{K}_q(\mathrm{BS}(m,n))$ has empty interior in $\mathcal{K}_q(\mathrm{BS}(m,n))$.*
5. *If $\gcd(m,n) = 1$, then $s = q$ and $\mathcal{MC}_q \cap \mathcal{K}(\mathrm{BS}(m,n)) = \{\langle\langle b^q \rangle\rangle\}$; in particular, $\langle\langle b^q \rangle\rangle$ is the unique normal subgroup of phenotype q of infinite index.*

Proof of Theorem 5.20. The proofs of (2) and (3) rely on the following claim.

Claim. *For any* $\Lambda \in \mathbf{Ph}^{-1}(q) \setminus \mathcal{MC}_q$*, there are a prime* p *such that* $|m|_p \neq |n|_p$ *and a vertex label* L *in the Bass–Serre graph of* Λ *such that* $|L|_p > |s|_p$*.*

Proof of the claim. Observe that a subgroup Λ contains $\langle\langle b^s \rangle\rangle$ if and only if all the b-orbits of the corresponding action $\Lambda\backslash \mathrm{BS}(m,n) \curvearrowleft \mathrm{BS}(m,n)$ have cardinality which divides s. So if $\Lambda \in \mathbf{Ph}^{-1}(q) \setminus \mathcal{MC}_q$, we can fix a prime p such that $|L|_p > |s|_p$, and we will prove that $|m|_p \neq |n|_p$.

Assume by contradiction that $|m|_p = |n|_p$. Then $|s|_p \geq |m|_p = |n|_p$. If $|m|_p = 0$, then the inequality clearly holds, otherwise, by equation (5.18),

- if p divides $q = \mathrm{Ph}_{m,n}(s)$, then $|s|_p = |q|_p > |m|_p = |n|_p$,
- if p does not divide $q = \mathrm{Ph}_{m,n}(s)$, then $|s|_p = |m|_p = |n|_p$.

Thus, we have $|L|_p > |m|_p = |n|_p$, in other words, $p \in \mathcal{P}_{m,n}(L)$ (see Definition 4.1). Hence, we have $|\mathrm{Ph}_{m,n}(L)|_p = |L|_p > |s|_p \geq |\mathrm{Ph}_{m,n}(s)|_p$. This is a contradiction since both phenotypes are equal to q. $\square_{\text{claim}}$

We can now easily prove (2) by the contrapositive. By the above claim, if $\mathcal{MC}_q \neq \mathbf{Ph}^{-1}(q)$, then there is a prime p such that $|m|_p \neq |n|_p$, in particular, $|m| \neq |n|$.

Let us prove (3). Let $\Lambda \in \mathbf{Ph}^{-1}(q) \setminus \mathcal{MC}_q$. The claim above provides a prime p such that $|m|_p \neq |n|_p$ and the Bass–Serre graph of Λ admits a vertex label L such that $|L|_p > |s|_p$. It follows from equation (5.18) that $|s|_p = \min(|m|_p, |n|_p)$, therefore we have $|L|_p > \min(|m|_p, |n|_p)$. Lemma 5.16 thus applies, and so there is a sequence of vertices in the Bass–Serre graph of Λ whose labels tend to $+\infty$. In other words, there is a sequence $(\gamma_i)_{i\geq 0}$ such that the index of $\gamma_i \Lambda \gamma_i^{-1} \cap \langle b \rangle$ in $\langle b \rangle$ tends to $+\infty$. By compactness, we may assume that this sequence converges, and its limit Δ cannot contain a non-zero power of b since $[\langle b \rangle : \gamma_i \Lambda \gamma_i^{-1} \cap \langle b \rangle] \to +\infty$. Hence, Δ has infinite phenotype, which proves (3).

We now prove (1). We first claim that $\mathcal{MC}_q$ is closed in $\mathrm{Sub}(\mathrm{BS}(m,n))$. Indeed, the set

$$\mathcal{B}_s := \{\Lambda \in \mathrm{Sub}(\mathrm{BS}(m,n)) : \Lambda \geq \langle\langle b^s \rangle\rangle\}$$

is a countable intersection of basic clopen sets and hence it is closed. Then notice that $\mathcal{B}_s$ intersects only finitely many sets $\mathbf{Ph}^{-1}(q')$, since q' must be finite and divide s. Proposition 5.8 claims that the $\mathbf{Ph}^{-1}(q')$ are open, hence

$$\mathcal{MC}_q = \mathcal{B}_s \setminus \bigcup_{\substack{q' \neq q \\ q' \text{ divides } s}} \mathbf{Ph}^{-1}(q')$$

is closed. Also note that $\mathcal{MC}_q$ is obviously $\mathrm{BS}(m,n)$-invariant. Finally, item (3) implies that every closed $\mathrm{BS}(m,n)$-invariant subset of $\mathbf{Ph}^{-1}(q)$ is contained in $\mathcal{MC}_q$. This proves that $\mathcal{MC}_q$ is the largest closed $\mathrm{BS}(m,n)$-invariant subset of $\mathrm{Sub}(\mathrm{BS}(m,n))$ contained in $\mathbf{Ph}^{-1}(q)$. Since all normal subgroups and all finite index subgroups have finite (hence closed) orbits, the remaining statement in item (1) is clear.

Let us prove item (4). Suppose $|n| \neq |m|$ and let p be a prime number such that $|m|_p \neq |n|_p$. By definition, $\mathrm{Ph}_{m,n}(sp) = \mathrm{Ph}_{m,n}(s) = q$, so $\langle b^{sp} \rangle \in \mathcal{K}_q(\mathrm{BS}(m,n)) \setminus \mathcal{MC}_q$. Consider a subgroup $\Lambda \in \mathcal{K}_q(\mathrm{BS}(m,n))$ whose orbit is dense in $\mathcal{K}_q(\mathrm{BS}(m,n))$, as provided by Corollary 5.15. Since the orbit of Λ accumulates to $\langle b^{sp} \rangle \notin \mathcal{MC}_q$ and $\mathcal{MC}_q$

is invariant and closed, the latter does not contain any point of that orbit. Hence, the complement $\mathcal{K}_q(\mathrm{BS}(m,n)) \setminus \mathcal{MC}_q$ contains the dense orbit of Λ. We conclude that $\mathcal{MC}_q \cap \mathcal{K}_q(\mathrm{BS}(m,n))$ has empty interior in $\mathcal{K}_q(\mathrm{BS}(m,n))$.

We finally prove item (5). The equality $s = q$ follows immediately from formula (5.18) for $s(q,m,n)$. We have the presentation

$$\mathrm{BS}(m,n)/\langle\langle b^q \rangle\rangle = \langle \bar{b}, \bar{t} : \bar{t}\bar{b}^m\bar{t}^{-1} = \bar{b}^n, \bar{b}^q = 1\rangle.$$

Since $\gcd(q,m) = \gcd(q,n) = 1$, the elements $\bar{b}^m$ and $\bar{b}^n$ both generate $\langle \bar{b}\rangle$ in the quotient group $\mathrm{BS}(m,n)/\langle\langle b^q \rangle\rangle$. We thus have a natural semi-direct product decomposition

$$\mathrm{BS}(m,n)/\langle\langle b^q \rangle\rangle \cong \mathbb{Z}/q\mathbb{Z} \rtimes \mathbb{Z} = \langle \bar{b}\rangle \rtimes \langle \bar{t}\rangle.$$

Consider $\Lambda \in \mathcal{MC}_q$ in the perfect kernel; by definition, it contains $\langle\langle b^q \rangle\rangle$. It suffices to prove that the image $\Lambda_q := \Lambda/\langle\langle b^q \rangle\rangle$ of Λ in $\langle \bar{b}\rangle \rtimes \langle \bar{t}\rangle$ is trivial. Since $\mathbf{Ph}(\Lambda) = q$, the index $[\langle b\rangle : \Lambda \cap \langle b\rangle]$ is a multiple of q, so we have $\Lambda_q \cap \langle \bar{b}\rangle = \{\mathrm{id}\}$. Thus, Λ_q is mapped injectively in the quotient $\langle \bar{b}\rangle \rtimes \langle \bar{t}\rangle/\langle \bar{b}\rangle \simeq \mathbb{Z}$. If this image were not $\{0\}$, then Λ would have finite index in $\mathrm{BS}(m,n)$, contradicting that Λ is in the perfect kernel. The group Λ_q is thus trivial as wanted. ∎

Remark 5.21. In terms of actions, $\mathcal{MC}_q$ is the set of classes $[\alpha, v]$ all of whose cardinalities of b-orbits divide s and have phenotype q.

Proposition 5.22. *Let $m,n \in \mathbb{Z} \setminus \{0\}$ and $k \in \mathbb{Z}_{\geq 1}$. Let*

$$G_{m,n,k} := \mathrm{BS}(m,n)/\langle\langle b^k \rangle\rangle = \langle \bar{t}, \bar{b} \mid \bar{t}\bar{b}^m\bar{t}^{-1} = \bar{b}^n, \bar{b}^k = 1\rangle$$

and let

$$r(k) := \max\{r' \in \mathbb{N} : r' \text{ divides } k \text{ and } \gcd(r',m) = \gcd(r',n)\}.$$

Then:

(1) *b has order $r(k)$ in the quotient $G_{m,n,k}$; in particular, $\langle\langle b^k \rangle\rangle = \langle\langle b^{r(k)} \rangle\rangle$.*

(2) *The group $G_{m,n,k} = G_{m,n,r(k)}$ is the HNN extension of $\mathbb{Z}/r(k)\mathbb{Z} = \langle \bar{b}\rangle$ with respect to the relation $\bar{t}\bar{b}^m\bar{t}^{-1} = \bar{b}^n$.*

(3) $\mathrm{Ph}_{m,n}(k) = \mathrm{Ph}_{m,n}(r(k)) = \mathbf{Ph}(\langle\langle b^k \rangle\rangle)$.

Remark 5.23. It follows from item (1) in the above proposition that

$$r(k) = [\langle b\rangle : \langle\langle b^k \rangle\rangle \cap \langle b\rangle].$$

It is a routine computation, working prime number by prime number, to check that

$$r(k) = \prod_{\substack{p\in\mathcal{P}\\ |m|_p = |n|_p}} p^{|k|_p} \cdot \prod_{\substack{p\in\mathcal{P}\\ |m|_p \neq |n|_p}} p^{\min(|k|_p, |m|_p, |n|_p)} \tag{5.24}$$

In particular, $r(k)$ is a multiple of all the r's which divide k and satisfy the equation $\gcd(r',m) = \gcd(r',n)$. Moreover, $r(r(k)) = r(k)$.

Remark 5.25. It also follows from items (1) and (3) of the above proposition that the set of integers k of phenotype q such that $r(k) = k$ parametrizes the normal subgroups of the form $\langle\langle b^{k'}\rangle\rangle$ of phenotype q. Comparing equations (5.24) and (5.18), one can check that this is exactly the set of integers k that are multiple of q and that divide $s(q, m, n)$, i.e., $k = q \cdot j$, where

(1) $|j|_p = 0$ for $p \in \mathcal{P}$ such that $|m|_p = |n|_p = 0$,

(2) $|j|_p \leq |m|_p$ for $p \in \mathcal{P}$ such that $|m|_p = |n|_p > 0$ and $|q|_p = 0$,

(3) $|j|_p \leq \min\{|n|_p, |m|_p\}$ for $p \in \mathcal{P}$ such that $|m|_p \neq |n|_p$.

Proof of Proposition 5.22. Set $r := r(k)$. Since $\bar{b}^m$ and $\bar{b}^n$ are conjugate in $G_{m,n,k}$, they have the same order:

$$\frac{\operatorname{ord}(\bar{b})}{\gcd(\operatorname{ord}(\bar{b}), m)} = \operatorname{ord}(\bar{b}^m) = \operatorname{ord}(\bar{b}^n) = \frac{\operatorname{ord}(\bar{b})}{\gcd(\operatorname{ord}(\bar{b}), n)}.$$

Thus, $\gcd(\operatorname{ord}(\bar{b}), m) = \gcd(\operatorname{ord}(\bar{b}), n)$. Moreover, $\operatorname{ord}(\bar{b})$ divides k. So, by the definition of r, the order $\operatorname{ord}(\bar{b})$ divides r and hence $b^r \in \langle\langle b^k\rangle\rangle$. On the other hand, $b^k \in \langle b^r\rangle$, so $\langle\langle b^r\rangle\rangle = \langle\langle b^k\rangle\rangle$ and $G_{m,n,k} = G_{m,n,r}$.

Since $\gcd(r, m) = \gcd(r, n)$, we have that the subgroups generated by $\tilde{b}^m$ and $\tilde{b}^n$ in the group $\mathbb{Z}/r\mathbb{Z} = \langle \tilde{b} : \tilde{b}^r = 1\rangle$ are isomorphic. We can thus consider the HNN-extension of $\mathbb{Z}/r\mathbb{Z} = \langle \tilde{b} : \tilde{b}^r = 1\rangle$ with the relation $\tilde{t}\tilde{b}^m\tilde{t}^{-1} = \tilde{b}^n$. This admits the presentation $\langle \tilde{t}, \tilde{b} \mid \tilde{t}\tilde{b}^m\tilde{t}^{-1} = \tilde{b}^n, \tilde{b}^r = 1\rangle$ and it is hence isomorphic to $G_{m,n,r}$.

By the normal form theorem for HNN-extensions, the vertex group injects, i.e., $\bar{b}$ has order exactly r. Finally, formula (5.24) implies that

$$\mathrm{Ph}_{m,n}(k) = \mathrm{Ph}_{m,n}(r) = \mathbf{Ph}(\langle\langle b^r\rangle\rangle). \qquad \blacksquare$$

Theorem 5.26. *Let $m, n \in \mathbb{Z} \setminus \{0\}$ and q be a finite phenotype.*

(1) *If $\gcd(m, n) = 1$, then the perfect kernel contains a unique normal subgroup of phenotype q, namely, $\langle\langle b^q\rangle\rangle$.*

(2) *If $\gcd(m, n) \neq 1$, then the perfect kernel contains continuum many normal subgroups of phenotype q.*

Proof. The case $\gcd(m, n) = 1$ follows from item (5) of Theorem 5.20. Therefore, let us assume that $\gcd(m, n) \neq 1$.

Consider a prime p which divides both m and n. Then either $|q|_p \neq 0$ and we set $k := q$, otherwise, set $k := qp$. In both cases, remark that $\mathrm{Ph}_{m,n}(k) = q$ and $\gcd(k, m) = \gcd(k, n)$, and hence $r(k) = k$. Then Proposition 5.22 yields that $\bar{b}$ has order k in $G_{m,n,k}$. Furthermore, since $k_0 := \gcd(k, m) = \gcd(k, n) > 1$, the elements $\bar{b}^n$ and $\bar{b}^m$ are not generators of the subgroup $\langle \bar{b}\rangle$; the group $G_{m,n,k}$ is not a semi-direct product. We claim that $G_{m,n,k}$ is not amenable. Indeed, we can write the group $G_{m,n,k}$ as the amalgamated free product

$$G_{m,n,k} = \langle \bar{t}, \bar{c} \mid \bar{t}(\bar{c})^{m/k_0}\bar{t}^{-1} = (\bar{c})^{n/k_0}, (\bar{c})^{k/k_0} = 1\rangle *_{\bar{c}=\bar{b}^{k_0}} \langle \bar{b} \mid \bar{b}^k\rangle,$$

and one can easily check that $G_{m,n,k}$ admits as a quotient the non-amenable free product $\langle \tilde{t}\rangle * \langle \tilde{b} \mid \tilde{b}^{k_0}\rangle$.

Since $G_{m,n,k}$ is the fundamental group of a finite graph of finite groups, it admits a finite index normal subgroup F which is a finitely generated free group by Proposition 11, p. 120, in [25]. Since $G_{m,n,k}$ is non-amenable, this normal free subgroup is not amenable.

Every characteristic subgroup N of F is itself normal in $G_{m,n,k}$. Thus, the pull-back under the quotient map $\mathrm{BS}(m,n) \twoheadrightarrow G_{m,n,k}$ is a normal subgroup $\tilde{N} \lhd \mathrm{BS}(m,n)$. Since the intersection of F with the finite group $\langle \bar{b} \rangle$ is trivial, the same holds for its characteristic subgroups, that is, $N \cap \langle \bar{b} \rangle = \{\mathrm{id}\}$. Therefore, the order of the image of b in $G_{m,n,k}/N = \mathrm{BS}(m,n)/\tilde{N}$ is the same as in $G_{m,n,k}$, namely, k. In other words, we have $\tilde{N} \cap \langle b \rangle = \langle b^k \rangle$. By definition,

$$\mathbf{Ph}(\tilde{N}) = \mathrm{Ph}_{m,n}([\langle b \rangle : \tilde{N} \cap \langle b \rangle] = \mathrm{Ph}_{m,n}(k) = q.$$

There are continuum many characteristic subgroups N in the finitely generated free subgroup F by [9] (see also [8]). At most countably many of them lie outside the perfect kernel, so the theorem follows. ∎

6. Limits of finite phenotype subgroups

In this section, we characterize the subgroups of infinite phenotype of $\mathrm{BS}(m,n)$ which arise as limits of finite phenotype subgroups. We will use a version of the straightforward fact that finitely generated subgroups always form a dense set in the space of subgroups.

Lemma 6.1. *Let $m,n \in \mathbb{Z} \setminus \{0\}$. For every phenotype $q \in \mathcal{Q}_{m,n}$, the finitely generated subgroups of phenotype q are dense in $\mathbf{Ph}^{-1}(q)$.*

Proof. Let Λ be a non-finitely generated subgroup of phenotype q. Let $k \in \mathbb{Z}_{\geq 0}$ be such that $\Lambda \cap \langle b \rangle = \langle b^k \rangle$. The group Λ can be written as the increasing union of finitely generated subgroups all containing b^k. They have the same phenotype as Λ. ∎

6.1. Limits of subgroups with fixed finite phenotype

Recall from Proposition 5.8 that, for q finite, $\mathbf{Ph}^{-1}(q)$ is open while $\mathbf{Ph}^{-1}(\infty)$ is closed, and from Theorem 5.20 (3) that the orbit of any $\Lambda \in \mathbf{Ph}^{-1}(q) \setminus \mathcal{MC}_q$ accumulates to $\mathbf{Ph}^{-1}(\infty)$. We now determine the set of such accumulation points in $\mathbf{Ph}^{-1}(\infty)$, which is exactly the set of subgroups contained in the normal closure $\langle\!\langle b \rangle\!\rangle$ of $\langle b \rangle$ but having trivial intersection with $\langle b \rangle$ itself (since they belong to $\mathbf{Ph}^{-1}(\infty)$).

Theorem 6.2. *Suppose $|m| \neq |n|$ and let q be a finite phenotype. Then*

$$\overline{\mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty) = \{\Lambda \in \mathbf{Ph}^{-1}(\infty) \colon \Lambda \leq \langle\!\langle b \rangle\!\rangle\}.$$

We need two preparatory lemmas. We start with an easy consequence of the defining relation $tb^m = b^n t$ of $\mathrm{BS}(m,n)$.

Notation 6.3. Given $\gamma \in \mathrm{BS}(m,n)$, let us denote:

- by κ_γ the *t-length* of γ, namely, the number of occurrences of $t^{\pm 1}$ in the normal form of γ,

- by Σ_γ the number of occurrences of t minus the number of occurrences of t^{-1} in the normal form of γ, which is often called the *t-height* of γ.

Remark that Σ_γ is the image of γ in $\mathrm{BS}(m,n)/\langle\langle b\rangle\rangle \cong \mathbb{Z}$. In particular, $\Sigma_\gamma = 0$ if and only if $\gamma \in \langle\langle b\rangle\rangle$.

Lemma 6.4. *Fix $\gamma \in \mathrm{BS}(m,n)$. Let $A \in \mathbb{Z}$ be such that, for all primes $p \in \mathcal{P}$,*

- *if $|m|_p = |n|_p$, then $|A|_p \geq |m|_p$,*
- *otherwise, $|A|_p \geq \kappa_\gamma |m|_p$ and $|A|_p \geq \kappa_\gamma |n|_p$.*

Then there is $B \in \mathbb{Z}$, such that $\gamma b^A = b^B \gamma$, where $|B|$ is determined by

$$|B|_p = |A|_p + \Sigma_\gamma(|n|_p - |m|_p) \quad \textit{for all } p \in \mathcal{P}.$$

Proof. This follows from a straightforward induction on κ_γ using the relation $tb^m = b^n t$. We leave the details to the reader. ∎

The proof of the inclusion in Theorem 6.2 from left to right relies on the following lemma.

Lemma 6.5. *Fix $\gamma \notin \langle\langle b\rangle\rangle$ and let q be a finite phenotype. There is an integer $R = R(q,\gamma)$ such that every subgroup Λ of phenotype q containing γ must also contain b^R.*

Proof. Up to replacing γ by its inverse, let us assume $\Sigma_\gamma > 0$. We first define the integer

$$M := \max\{|m|_p, |n|_p : p \in \mathcal{P}\},$$

and then we let

$$R := q\Big(\prod_{\substack{p\in\mathcal{P}\\ |m|_p+|n|_p>0}} p\Big)^{\kappa_\gamma M}.$$

Fix Λ of phenotype q. Since q is finite, we have $\langle b\rangle \cap \Lambda = \langle b^N\rangle$ with $N > 0$. We have to show that N divides R. Notice that $\mathrm{Ph}_{m,n}(N) = q$, thus N decomposes as

$$N = q \cdot p_1^{l_1} \cdots p_k^{l_k} p_{k+1}^{l_{k+1}} \cdots p_r^{l_r},$$

where $r \geq 0$ and $l_1, \dots, l_r \geq 1$, while the p_i are distinct prime numbers coprime with q, see Definition 4.1. Moreover, we order them so that $p_1, \dots, p_k \in \mathcal{P}_{m,n} \setminus \mathcal{P}_{m,n}(N)$ and $p_{k+1}, \dots, p_r \in \mathcal{P} \setminus \mathcal{P}_{m,n}$.

Observe that $|m|_{p_i} = |n|_{p_i} \geq |N|_{p_i} = l_i \geq 1$ when $p_i \in \mathcal{P}_{m,n} \setminus \mathcal{P}_{m,n}(N)$ and $|m|_{p_i} \neq |n|_{p_i}$ when $p_i \in \mathcal{P} \setminus \mathcal{P}_{m,n}$. Hence, $|m|_{p_i} + |n|_{p_i} > 0$ for every $i \in \{1, \dots, r\}$. Consequently, to establish that N divides R, it suffices to prove

$$l_i \leq \kappa_\gamma M \quad \text{for all } i \in \{1, \dots, r\}. \tag{6.6}$$

Observe that $\kappa_\gamma \geq 1$, since $\gamma \notin \langle\langle b\rangle\rangle$. For $i \in \{1, \dots, k\}$, equation (6.6) holds, since $p_i \in \mathcal{P}_{m,n} \setminus \mathcal{P}_{m,n}(N)$, thus

$$l_i \leq |m|_{p_i} = |n|_{p_i} \leq M \leq \kappa_\gamma M.$$

Let us hence fix $i \in \{k+1, \dots, r\}$ and suppose, by contradiction, that $l_i > \kappa_\gamma M$. Consider

$$N' = N \times (p_1 \cdots p_k)^M (p_{k+1} \cdots \widehat{p_i} \cdots p_r)^{\kappa_\gamma M},$$

where by $\widehat{p_i}$ we mean that the factor p_i is removed from the product. Clearly, $b^{N'} \in \Lambda$ and $|N'|_{p_i} = l_i$. Put

$$\varepsilon := \operatorname{sign}(|m|_{p_i} - |n|_{p_i}).$$

Note that $p_i \notin \mathcal{P}_{m,n}$, hence $|m|_{p_i} \neq |n|_{p_i}$, so $\varepsilon \neq 0$. Since we assumed $|N|_{p_i} = l_i \geq \kappa_\gamma M$, we also have $|N'|_{p_i} \geq \kappa_\gamma M$. It is then clear that N' satisfies the assumption of Lemma 6.4, so $\gamma^\varepsilon b^{N'} \gamma^{-\varepsilon} = b^{N''}$, where

$$\begin{aligned} |N''|_{p_i} = l_i + \Sigma_{\gamma^\varepsilon}(|n|_{p_i} - |m|_{p_i}) &= l_i + \varepsilon \Sigma_\gamma (|n|_{p_i} - |m|_{p_i}) \\ &= l_i - \Sigma_\gamma \big||m|_{p_i} - |n|_{p_i}\big| < l_i. \end{aligned}$$

Clearly, $b^{N''} \in \Lambda$, hence $b^{N''} \in \langle b^N \rangle$. But $|N''|_{p_i} < |N|_{p_i}$, a contradiction. We thus have established equation (6.6), which finishes the proof. ∎

Proof of Theorem 6.2. Set

$$\mathcal{L} := \{\Lambda \in \mathbf{Ph}^{-1}(\infty); \Lambda \leq \langle\!\langle b \rangle\!\rangle\}.$$

We first show the inclusion

$$\overline{\mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty) \subseteq \mathcal{L}.$$

Take $\Delta \in \mathbf{Ph}^{-1}(\infty) \setminus \mathcal{L}$ and $\gamma \in \Delta \setminus \langle\!\langle b \rangle\!\rangle$. By Lemma 6.5, there is an R such that every subgroup Λ of phenotype q containing γ also contains b^R. Thus, the clopen neighborhood of Δ, given by

$$\mathcal{O} := \{\Lambda \in \operatorname{Sub}(\operatorname{BS}(m,n)) : \gamma \in \Lambda,\ b^R \notin \Lambda\},$$

does not intersect $\mathbf{Ph}^{-1}(q)$. Thus, Δ is not in the closure of $\mathbf{Ph}^{-1}(q)$.

We now show the reverse inclusion

$$\mathcal{L} \subseteq \overline{\mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty).$$

Remark that, as in Lemma 6.1, the finitely generated elements of $\mathcal{L}$ are dense in $\mathcal{L}$: every element of $\mathcal{L}$ is an increasing union of finitely generated subgroups which have to be in $\mathcal{L}$ as well. So, take $\Lambda = \langle S \rangle \in \mathcal{L}$, where S is finite; we will show that Λ is a limit of subgroups with phenotype q. Set $\kappa := \max_{\gamma \in S} \kappa_\gamma$, where κ_γ is the t-length of γ (see Notation 6.3). Set $M := \max\{|m|_p, |n|_p : p \in \mathcal{P}\}$. Note that $\mathcal{P} \setminus \mathcal{P}_{m,n}$ is finite, since it is composed of primes p such that $|m|_p + |n|_p > 0$, and that $|m|_p = 0$ for all but finitely many primes p. Hence, for $j \geq 1$, we can define the integer

$$N_j := q \cdot \prod_{p \in \mathcal{P}_{m,n} \setminus \mathcal{P}_{m,n}(q)} p^{|m|_p} \cdot \prod_{p \in \mathcal{P} \setminus \mathcal{P}_{m,n}} p^{j\kappa M}.$$

Observe that $\mathrm{Ph}_{m,n}(N_j) = q$.

Since $\Lambda \leq \langle\langle b \rangle\rangle$, the height Σ_γ is zero (see Notation 6.3) for every $\gamma \in S$, whence, for every $\gamma \in S$ and every j, Lemma 6.4 gives $\gamma b^{N_j} = b^{\pm N_j}\gamma$. Thus, $\Lambda = \langle S \rangle$ normalizes $\langle b^{N_j} \rangle$. Moreover, Λ has trivial intersection with $\langle b^{N_j} \rangle$ because it has infinite phenotype. In particular, for $j = 1$, we have a natural isomorphism

$$\Phi : \Lambda \ltimes \langle b^{N_1} \rangle \to \langle \Lambda, b^{N_1} \rangle.$$

Since N_1 divides N_j, we get

$$\Phi(\Lambda \ltimes \langle b^{N_j} \rangle) = \langle \Lambda, b^{N_j} \rangle.$$

Observe that Φ induces a homeomorphism

$$\mathrm{Sub}(\Lambda \ltimes \langle b^{N_1} \rangle) \to \mathrm{Sub}(\langle \Lambda, b^{N_1} \rangle) \subseteq \mathrm{Sub}(\mathrm{BS}(m,n)),$$

and that the sequence of subgroups $(\Lambda \ltimes \langle b^{N_j} \rangle)_{j\geq 1}$ converges to $\Lambda \ltimes \{\mathrm{id}\}$. Therefore, we have that $\langle \Lambda, b^{N_j} \rangle$ converges to Λ. Since

$$\mathbf{Ph}(\langle \Lambda, b^{N_j} \rangle) = \mathrm{Ph}_{m,n}(N_j) = q,$$

the group Λ is the limit of a sequence of elements of phenotype q as wanted. ∎

6.2. Limits of subgroups with varying finite phenotype

In Theorem 6.2, we showed that $\overline{\mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty)$ does not depend on the finite phenotype q. We will now consider the closure of all subgroups with finite phenotype, and we will first analyze what happens if $|m| = |n|$.

Proposition 6.7. *Let m and n be integers such that $|m| = |n| \geq 2$. Then*

$$\mathbf{Ph}^{-1}(\infty) \subseteq \overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)}.$$

In other words, every subgroup with infinite phenotype is a limit of subgroups with finite (variable) phenotypes.

Proof. Let us fix $\Lambda \in \mathbf{Ph}^{-1}(\infty)$. Note that $\langle b^n \rangle$ is normalized by Λ thanks to the relation $tb^nt^{-1} = b^{\pm n}$. We now proceed as in the second part of the proof of Theorem 6.2. The group $\langle \Lambda, b^{jn} \rangle$ has finite phenotype, it is isomorphic to $\Lambda \ltimes \langle b^{jn} \rangle$ and the sequence of subgroups $(\langle \Lambda, b^{jn} \rangle)_{j\geq 1}$ converges to Λ. ∎

The situation is completely different in the case $|m| \neq |n|$.

Proposition 6.8. *Let m and n be integers such that $|m| \neq |n|$ and $|m|, |n| \geq 2$. Then*

$$\mathbf{Ph}^{-1}(\infty) \not\subseteq \overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)}.$$

In other words, there are subgroups with infinite phenotype that are not limits of subgroups with finite (variable) phenotypes.

Let us recall from Corollary 5.4 that

$$\mathbf{Ph}^{-1}(\infty) = \mathcal{K}_\infty(\mathrm{BS}(m,n)) \quad \text{whenever } |m| \neq |n|.$$

Hence, the subgroups given by the proposition lie in fact in $\mathcal{K}_\infty(\mathrm{BS}(m,n))$.

In the proof of Proposition 6.8, we will need a lemma and a proposition.

Lemma 6.9. *Let m,n be integers such that $|m| \neq |n|$ and $|m|$, $|n| \geq 2$. Let $k := \gcd(m,n)$. Let $\Lambda \leq \mathrm{BS}(m,n)$ be a subgroup containing the following elements*:

$$t, btb^{-1}, \dots, b^{k-1}tb^{-(k-1)}.$$

If Λ has finite phenotype, then Λ has finite index in $\mathrm{BS}(m,n)$.

Proof. Let α be the action $\Lambda\backslash\mathrm{BS}(m,n) \curvearrowleft \mathrm{BS}(m,n)$. Since the phenotype is finite, it is sufficient to show that the Bass–Serre graph $\mathbf{BS}(\alpha)$ is finite (see Remark 4.10).

Since Λ contains t, there is a loop in $\mathbf{BS}(\alpha)$ at the vertex $v := \Lambda\langle b\rangle$. In particular, equation (3.13) gives

$$\frac{L(v)}{\gcd(L(v),m)} = \frac{L(v)}{\gcd(L(v),n)}.$$

As Λ has finite phenotype, $L(v)$ is finite, so $\gcd(L(v),m) = \gcd(L(v),n)$. Moreover, as $\mathbf{BS}(\alpha)$ is a saturated (m,n)-graph, we obtain

$$\deg_{\mathrm{in}}(v) = \gcd(L(v),m) = \gcd(L(v),n) = \deg_{\mathrm{out}}(v).$$

This number, that we will denote d, is the greatest common divisor of m, n and $L(v)$. Hence, d divides $k = \gcd(m,n)$.

The d outgoing edges at v are exactly $\Lambda\langle b^n\rangle, \Lambda b\langle b^n\rangle, \dots, \Lambda b^{d-1}\langle b^n\rangle$. As $d \leq k$, the subgroup Λ contains $t, btb^{-1}, \dots, b^{d-1}tb^{-(d-1)}$. Since

$$\Lambda b^j t = (\Lambda b^j t b^{-j}) b^j = \Lambda b^j,$$

the element t fixes all the points $\Lambda, \Lambda b, \dots, \Lambda b^{d-1} \in \Lambda\backslash\mathrm{BS}(m,n)$. The terminal vertex of the edge $\Lambda b^j\langle b^n\rangle$ is precisely the vertex $\Lambda b^j t\langle b\rangle = \Lambda b^j\langle b\rangle = v$ (see Definition 3.5), so all outgoing edges at v are loops.

Since the outgoing degree at v is equal to the incoming degree, all incoming edges at v are loops as well. Therefore, $\mathbf{BS}(\alpha)$ consists only of the vertex v and d loops. It is thus finite as wanted. ∎

Proposition 6.10. *Let m and n be integers with $|m|, |n| \geq 2$. Let Λ be a finitely generated subgroup of infinite phenotype and infinite Bass–Serre graph. Then there is a sequence of conjugates of Λ which converges to* $\{\mathrm{id}\}$. *In particular, Λ does not contain any non-trivial normal subgroup of* $\mathrm{BS}(m,n)$.

Proof. First recall that Λ is free. Indeed, having infinite phenotype, it acts freely on the Bass–Serre tree $\mathcal{T}$ of $\mathrm{BS}(m,n)$. Taking the class $\langle b\rangle$ as a base point in $\mathcal{T}$, the subgroup Λ is the fundamental group of the quotient graph $\Lambda\backslash\mathcal{T}$ based at $\Lambda\langle b\rangle$. This quotient graph is equal to the Bass–Serre graph of Λ, see Section 3.6, so it is infinite. Since moreover Λ is finitely generated, it consists of a finite graph to which are attached finitely many infinite trees. Moving the basepoint along one of those infinite trees toward infinity amounts to conjugating Λ by a certain sequence of elements γ_i of $\mathrm{BS}(m,n)$ for which we claim that $\gamma_i \Lambda \gamma_i^{-1} \to \{\mathrm{id}\}$. Indeed, each non-trivial element of $\gamma_i \Lambda \gamma_i^{-1}$ is represented by a long path in the tree, followed by a closed path in the finite graph and the long path back to the new basepoint. All such elements have a uniformly large t-length which tends to $+\infty$ with i: their t-length is bounded below by twice the t-length of γ_i minus the diameter of the finite graph. In particular, for any finite set $F \subset \Gamma \setminus \{\mathrm{id}\}$ and large enough i, all the elements of $\gamma_i \Lambda \gamma_i^{-1}$ have t-length larger than all those of F, so $\gamma_i \Lambda \gamma_i^{-1} \cap F = \varnothing$. This proves that $\gamma_i \Lambda \gamma_i^{-1} \to \{\mathrm{id}\}$, as wanted. ■

Proof of Proposition 6.8. Consider the group $\Lambda := \langle t, btb^{-1}, \dots, b^{k-1}tb^{-(k-1)} \rangle$. Observe that, by Britton's lemma (see, e.g., Chapter IV.2 of [23]), it is a free group freely generated by $t, btb^{-1}, \dots, b^{k-1}tb^{-(k-1)}$. Every non-trivial element of Λ contains at least one $t^{\pm 1}$ in its normal form, in particular, $\Lambda \cap \langle b \rangle = \{\mathrm{id}\}$: the phenotype of Λ is infinite. We claim that

$$\Lambda \notin \overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)}.$$

Suppose that $(\Lambda_i)_{i \geq 0}$ is a sequence of subgroups of finite (variable) phenotypes converging to Λ. For i large enough, we have $t, btb^{-1}, \dots, b^{k-1}tb^{-(k-1)} \in \Lambda_i$, and thus the subgroup Λ_i has finite index by Lemma 6.9. However, recall that since $|m| \neq |n|$, the group $\mathrm{BS}(m,n)$ is not residually finite [24]. Therefore, there is a non-trivial normal subgroup $N \trianglelefteq \mathrm{BS}(m,n)$ contained in every finite index subgroup, and we have $N \leq \Lambda$ since $\Lambda_i \to \Lambda$. This is impossible by Proposition 6.10. ■

Corollary 6.11. *Let m and n be integers such that $|m| \neq |n|$ and $|m|, |n| \geq 2$. Then*

$$\overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty)$$

has empty interior in $\mathbf{Ph}^{-1}(\infty)$.

Proof. Recall again that $\mathbf{Ph}^{-1}(\infty) = \mathcal{K}_\infty(\mathrm{BS}(m,n))$, see Corollary 5.4. In this space, the subset

$$\mathcal{K}_\infty(\mathrm{BS}(m,n)) \setminus \overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)}$$

is open and Proposition 6.8 implies that it is non-empty. By Corollary 5.15, this open subset contains a subgroup Λ whose orbit is dense in $\mathcal{K}_\infty(\mathrm{BS}(m,n))$. Therefore,

$$\overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)}$$

has empty interior in $\mathcal{K}_\infty(\mathrm{BS}(m,n))$. ■

Proposition 6.12. *Let m and n be integers such that $|m|, |n| \geq 2$. For any finite phenotype q_0, the following inclusion is strict:*

$$\overline{\mathbf{Ph}^{-1}(q_0)} \cap \mathbf{Ph}^{-1}(\infty) \subsetneq \overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty).$$

Observe that Proposition 6.12 is trivially true if $|m| = |n|$. Indeed, Proposition 6.7 implies that the right-hand side is equal to $\mathbf{Ph}^{-1}(\infty)$. Since Proposition 5.8 yields that $\mathbf{Ph}^{-1}(q_0)$ is closed, the left-hand side is empty.

Proof of Proposition 6.12. For a prime p which divides neither m nor n, let us define $\Lambda_p := \langle b^p, t\rangle$. Then Λ_p clearly has phenotype p (and index p in $\mathrm{BS}(m,n)$). Let Λ be an accumulation point of the sequence (Λ_p). Then, by construction, Λ has infinite phenotype, so it is in the set

$$\overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty).$$

However, it contains $t \notin \langle\langle b\rangle\rangle$, so it is not in $\overline{\mathbf{Ph}^{-1}(q_0)}$ by Theorem 6.2. ∎

Corollary 6.13. *Let m and n be integers such that $|m|, |n| \geq 2$. The following inclusion is strict*:

$$\bigcup_{q \text{ finite}} \overline{\mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty) \subsetneq \overline{\bigcup_{q \text{ finite}} \mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty).$$

Proof. If $|m| = |n|$, then as already remarked the left-hand side is empty.

If $|m| \neq |n|$, recall from Theorem 6.2 that $\overline{\mathbf{Ph}^{-1}(q)} \cap \mathbf{Ph}^{-1}(\infty)$ is independent of q. The corollary thus follows from Proposition 6.12. ∎

We can also give a statement analogous to Proposition 6.12 in the perfect kernel, which is less easy to obtain.

Theorem 6.14. *Let m and n be integers such that $|m|, |n| \geq 2$. For any finite phenotype q_0, the following inclusion is strict*:

$$\overline{\mathcal{K}_{q_0}(\mathrm{BS}(m,n))} \cap \mathcal{K}_\infty(\mathrm{BS}(m,n)) \subsetneq \overline{\bigcup_{q \text{ finite}} \mathcal{K}_q(\mathrm{BS}(m,n))} \cap \mathcal{K}_\infty(\mathrm{BS}(m,n)).$$

Proof. For a fixed prime p which divides neither m nor n, let us define a pre-action (β_p, τ_p) as follows. Consider three β_p-cycles say o_1, o_2 and o_3, of cardinals pn, p and pm, respectively. Then fix basepoints $y_i \in o_i$ for $i = 1, 2, 3$. Remark that o_1 splits into $|n| \geq 2$ β_p^n-orbits of size p and that o_3 splits into $|m| \geq 2$ β_p^m-orbits of size p. Therefore, we can define τ_p by setting

$$y_1 \beta_p^{jn} \tau_p := y_2 \beta_p^{jm}, \quad y_2 \beta_p^{jn} \tau_p := y_3 \beta_p^{jm} \quad \text{and} \quad y_1 \beta_p^{-1+jn} \tau_p := y_3 \beta_p^{1+jm}.$$

Clearly the phenotype of such a pre-action is p and the associated Bass–Serre graph

$$\mathcal{G}_{0,p} := \mathbf{BS}(\beta_p, \tau_p)$$

is a triangle. Set $x_p := y_1$ and note that for every p, we have

$$x_p \tau_p \tau_p \beta_p \tau_p^{-1} \beta_p = x_p.$$

By Lemma 4.22, we can then extend $\mathcal{G}_{0,p}$ to a saturated (m,n)-graph $\mathcal{G}_p$, see Figure 5, and by Proposition 3.23, we can extend the pre-action (β_p, τ_p) to an action α_p whose Bass–Serre graph is $\mathcal{G}_p$.

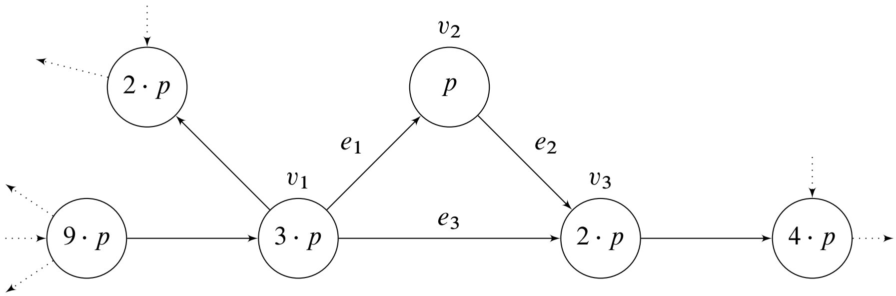

Figure 5. A (2,3)-graph $\mathscr{G}_p$, where $m = 2$ and $n = 3$.

Define Λ_p to be the stabilizer of the action α_p at x_p and remark that $t^2bt^{-1}b \in \Lambda_p$. Moreover, by construction, $\mathbf{Ph}(\Lambda_p) = p$.

By compactness, we find an accumulation point Λ of the sequence $(\Lambda_p)_p$. Since Λ_p has phenotype p, the subgroup Λ has infinite phenotype. Since $\underline{t^2bt^{-1}b} \in \Lambda_p$ for every p, we have that $t^2bt^{-1}b \in \Lambda$. Moreover, $t^2bt^{-1}b \notin \langle\langle b \rangle\rangle$, so $\Lambda \notin \mathbf{Ph}^{-1}(q_0)$ by Theorem 6.2. Therefore, the proof is complete. ∎

Acknowledgments. We are very grateful to both referees for their work and their detailed remarks which helped us to improve the paper.

Funding. A. Carderi acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 281869850 (RTG 2229). D. Gaboriau is supported by the CNRS and partially supported by the LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR). F. Le Maître acknowledges funding by the ANR projects ANR-17-CE40-0026 AGRUME and ANR-19-CE40-0008 AODynG.

References

[1] Azuelos, P. and Gaboriau, D.: Perfect kernel and dynamics: from Bass–Serre theory to hyperbolic groups. *Math. Ann.* **391** (2025), no. 3, 4733–4789. MR 4865255 Zbl 7986095

[2] Baumslag, G. and Solitar, D.: Some two-generator one-relator non-Hopfian groups. *Bull. Amer. Math. Soc.* **68** (1962), no. 3, 199–201. MR 142635 Zbl 0108.02702

[3] Becker, O., Lubotzky, A. and Thom, A.: Stability and invariant random subgroups. *Duke Math. J.* **168** (2019), no. 12, 2207–2234. MR 3999445 Zbl 1516.20007

[4] Bendixson, I.: Quelques théorèmes de la théorie des ensembles de points. *Acta Math.* **2** (1883), no. 1, 415–429. MR 1554609 Zbl 15.0455.01

[5] Bontemps, S.: Perfect kernel of generalized Baumslag–Solitar groups. Preprint 2024, arXiv:2411.03221v1.

[6] Bontemps, S., Gaboriau, D., Le Maître, F. and Stalder, Y.: On the space of subgroups of Baumslag–Solitar groups III: The Cantor–Bendixson rank. In preparation, 2025.

[7] Bowen, L., Grigorchuk, R. and Kravchenko, R.: Invariant random subgroups of lamplighter groups. *Israel J. Math.* **207** (2015), no. 2, 763–782. MR 3359717 Zbl 1334.43006

[8] Bowen, L., Grigorchuk, R. and Kravchenko, R.: Characteristic random subgroups of geometric groups and free abelian groups of infinite rank. *Trans. Amer. Math. Soc.* **369** (2017), no. 2, 755–781. MR 3572253 Zbl 1373.20093

[9] Bryant, R. M.: Characteristic subgroups of free groups. In *Proceedings of the Second International Conference on the theory of groups (Australian Nat. Univ., Canberra, 1973)*, pp. 141–149. Lecture Notes in Math. 372, Springer, Berlin-New York, 1974. MR 357609 Zbl 0288.20028

[10] Cantor, G.: Ueber unendliche, lineare Punktmannichfaltigkeiten, part 6. *Math. Ann.* **23** (1884), no. 4, 453–488. MR 1510266

[11] Carderi, A., Gaboriau, D. and Le Maître, F.: On dense totipotent free subgroups in full groups. *Geom. Topol.* **27** (2023), no. 6, 2297–2318. MR 4634748 Zbl 1533.37012

[12] Carderi, A., Gaboriau, D., Le Maître, F. and Stalder, Y.: How to build (m, n)-graphs. Zenodo, published October 25, 2022.

[13] Dudkin, F. A.: Subgroups of Baumslag–Solitar groups. *Algebra Logic* **48** (2009), no. 1, 1–19; translation from *Algebra Logika* **48** (2009), no. 1, 3–30. MR 2526434 Zbl 1245.20028

[14] Fima, P., Le Maître, F., Moon, S. and Stalder, Y.: A characterization of high transitivity for groups acting on trees. *Discrete Anal.* (2022), article no. 8, 63 pp. MR 4481405

[15] Forester, M.: Splittings of generalized Baumslag–Solitar groups. *Geom. Dedicata* **121** (2006), no. 1, 43–59. MR 2276234 Zbl 1117.20023

[16] Gaboriau, D., Le Maître, F. and Stalder, Y.: On the space of subgroups of Baumslag–Solitar groups II: High transitivity. Preprint 2024, arXiv:2410.23224v2.

[17] Gal, Ś. R. and Januszkiewicz, T.: New a-T-menable HNN-extensions. *J. Lie Theory* **13** (2003), no. 2, 383–385. MR 2003149 Zbl 1041.20028

[18] Gelman, E.: Subgroup growth of Baumslag–Solitar groups. *J. Group Theory* **8** (2005), no. 6, 801–806. MR 2179671 Zbl 1105.20018

[19] Glasner, Y., Kitroser, D. and Melleray, J.: From isolated subgroups to generic permutation representations. *J. Lond. Math. Soc. (2)* **94** (2016), no. 3, 688–708. MR 3614924 Zbl 1406.20034

[20] Kechris, A. S.: *Classical descriptive set theory*. Grad. Texts in Math. 156, Springer, New York, 1995. MR 1321597 Zbl 0819.04002

[21] Levitt, G.: On the automorphism group of generalized Baumslag–Solitar groups. *Geom. Topol.* **11** (2007), 473–515. MR 2302496 Zbl 1143.20014

[22] Levitt, G.: Quotients and subgroups of Baumslag–Solitar groups. *J. Group Theory* **18** (2015), no. 1, 1–43. MR 3297728 Zbl 1317.20030

[23] Lyndon, R. C. and Schupp, P. E.: *Combinatorial group theory*. Classics Math., Springer, Berlin, 2001. MR 1812024 Zbl 0997.20037

[24] Meskin, S.: Nonresidually finite one-relator groups. *Trans. Amer. Math. Soc.* **164** (1972), 105–114. MR 285589 Zbl 0245.20028

[25] Serre, J.-P.: *Trees*. Springer, Heidelberg, 1980. Zbl 0548.20018 MR 0607504

[26] Skipper, R. and Wesolek, P.: On the Cantor–Bendixson rank of the Grigorchuk group and the Gupta–Sidki 3 group. *J. Algebra* **555** (2020), 386–405. MR 4082048 Zbl 1457.20026

[27] Stalder, Y.: Moyennabilité intérieure et extensions HNN. *Ann. Inst. Fourier (Grenoble)* **56** (2006), no. 2, 309–323. MR 2226017 Zbl 1143.20013

Received April 2, 2024; revised August 11, 2024.

Alessandro Carderi
No longer affiliated with any mathematical institution.

Damien Gaboriau
Unité de Mathématiques Pures et Appliquées, École Normale Supérieure de Lyon
46 allée d'Italie, 69364 Lyon, France;
gaboriau@ens-lyon.fr

François Le Maître
Université Bourgogne Europe, CNRS, IMB UMR 5584
21000 Dijon, France;
flemaitre@math.cnrs.fr

Yves Stalder
Université Clermont Auvergne, CNRS, LMBP
63000 Clermont-Ferrand, France;
yves.stalder@uca.fr