UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH

Extensions of free groups: algebraic, geometric, and algorithmic aspects

Jordi Delgado Rodríguez

ADVERTIMENT La consulta d'aquesta tesi queda condicionada a l'acceptació de les següents condicions d'ús: La difusió d'aquesta tesi per mitjà del repositori institucional UPCommons (http://upcommons.upc.edu/tesis) i el repositori cooperatiu TDX (http://www.tdx.cat/) ha estat autoritzada pels titulars dels drets de propietat intel·lectual **únicament per a usos privats** emmarcats en activitats d'investigació i docència. No s'autoritza la seva reproducció amb finalitats de lucre ni la seva difusió i posada a disposició des d'un lloc aliè al servei UPCommons o TDX. No s'autoritza la presentació del seu contingut en una finestra o marc aliè a UPCommons (*framing*). Aquesta reserva de drets afecta tant al resum de presentació de la tesi com als seus continguts. En la utilització o cita de parts de la tesi és obligat indicar el nom de la persona autora.

ADVERTENCIA La consulta de esta tesis queda condicionada a la aceptación de las siguientes condiciones de uso: La difusión de esta tesis por medio del repositorio institucional UPCommons (http://upcommons.upc.edu/tesis) y el repositorio cooperativo TDR (http://www.tdx.cat/?locale-attribute=es) ha sido autorizada por los titulares de los derechos de propiedad intelectual **únicamente para usos privados enmarcados** en actividades de investigación y docencia. No se autoriza su reproducción con finalidades de lucro ni su difusión y puesta a disposición desde un sitio ajeno al servicio UPCommons No se autoriza la presentación de su contenido en una ventana o marco ajeno a UPCommons (*framing*). Esta reserva de derechos afecta tanto al resumen de presentación de la tesis como a sus contenidos. En la utilización o cita de partes de la tesis es obligado indicar el nombre de la persona autora.

WARNING On having consulted this thesis you're accepting the following use conditions: Spreading this thesis by the institutional repository UPCommons (http://upcommons.upc.edu/tesis) and the cooperative repository TDX (http://www.tdx.cat/?locale-attribute=en) has been authorized by the titular of the intellectual property rights **only for private uses** placed in investigation and teaching activities. Reproduction with lucrative aims is not authorized neither its spreading nor availability from a site foreign to the UPCommons service. Introducing its content in a window or frame foreign to the UPCommons service is not authorized (framing). These rights affect to the presentation summary of the thesis as well as to its contents. In the using or citation of parts of the thesis it's obliged to indicate the name of the author.

Extensions of free groups: algebraic, geometric, and algorithmic aspects

Jordi Delgado Rodríguez

Barcelona, June 2017
Version: v0.1

Universitat Politècnica de Catalunya

Departament de Matemàtiques
Facultat de Matemàtiques i Estadística
Programa de Doctorat en Matemàtiques

Extensions of free groups: algebraic, geometric, and algorithmic aspects

Jordi Delgado Rodríguez

PhD dissertation

Advisor: Enric Ventura Capell

Barcelona, June 2017

Jordi Delgado Rodríguez
Extensions of free groups: algebraic, geometric, and algorithmic aspects
PhD dissertation, Barcelona, June 2017
©2017 Jordi Delgado Rodríguez
All rights reserved

Mathematics Subject Classification (2010): 20E05, 20E06, 20E22, 20F05, 20F10, 20F36

Universitat Politècnica de Catalunya
Programa de Doctorat en Matemàtiques
Facultat de Matemàtiques i Estadística
Departament de Matemàtiques
Campus Nord, Carrer Jordi Girona, 1-3
08034 Barcelona

A mis padres

Abstract

In this work we use geometric techniques in order to study certain natural extensions of free groups, and solve several algorithmic problems on them.

To this end, we consider the family of free-abelian times free groups ($\mathbb{Z}^m \times \mathbb{F}_n$) as a seed towards further generalization in two main directions: semidirect products, and partially commutative groups (PC-groups).

The four principal projects of this thesis are the following:

Direct products of free-abelian and free groups [DV13; Del14b; DV17b]. We begin by studying the structure of the groups $\mathbb{Z}^m \times \mathbb{F}_n$, with special emphasis on their lattice of subgroups, and their endomorphisms (for which an explicit description is given, and both injectivity and surjectiveness are characterized); to then solve on them algorithmic problems involving both subgroups (membership problem, finite index problem, and subgroup and coset intersection problems) and endomorphisms (fixed points problem, Whitehead problems, and twisted-conjugacy problem).

Algorithmic recognition of infinite-cyclic extensions [Cav+17]. In the first part, we prove the algorithmic undecidability of several properties (finite generability, finite presentability, abelianity, finiteness, independence, triviality) of the base group of finitely presented cyclic extensions. In particular, we see that it is not possible to decide algorithmically if a finitely presented $\mathbb{Z}$-extension admits a finitely generated base group. This last result allows us to demonstrate the undecidability of the Bieri-Neumann-Strebel (BNS) invariant.

In the second part, we prove the equivalence between the isomorphism problem within the subclass of unique $\mathbb{Z}$-extensions, and the semi-conjugacy problem for certain type of outer automorphisms, which we characterize algorithmically.

Stallings automata for free-abelian by free groups [DV17a]. After recreating in a purely algorithmic language the classic theory of Stallings associating an automaton to each subgroup of the free group, we extend this theory to semidirect products of the form $\mathbb{Z}^m \rtimes \mathbb{F}_n$. Specifically, we associate to each subgroup of $\mathbb{Z}^m \rtimes \mathbb{F}_n$, an automaton ('enriched' with vectors in $\mathbb{Z}^m$), and we see that in the finitely generated case this construction is algorithmic and allows to solve the membership problem within this family of groups.

The geometric description obtained also shows (even in the case of direct products) not only that the intersection of finitely generated subgroups can be infinitely generated, but that even when it is finitely generated, the rank of the intersection can not be bound in terms of the ranks of the intersected subgroups. This fact is relevant because it denies any possible extension of the celebrated — and recently proven — Hanna Neumann conjecture in this direction.

Intersection problems for Droms groups [Del14a; DVZ17]. After characterizing those partially commutative groups satisfying the Howson property, we combine the algorithmic version of the Kurosh subgroup theorem given by S.V. Ivanov, with a generalization of some ideas from our work on $\mathbb{Z}^m \times \mathbb{F}_n$, to prove the solvability of the subgroup and coset intersection problems within the subfamily of Droms groups (that is, those PC- groups whose subgroups are always again partially commutative).

References

[Cav+17] B. Cavallo, J. Delgado, D. Kahrobaei, and E. Ventura. "Algorithmic Recognition of Infinite Cyclic Extensions". *Journal of Pure and Applied Algebra* 221.9 (Sept. 2017), pp. 2157–2179 (cit. on p. vii).

[Del14a] J. Delgado. "Some Characterizations of Howson PC-Groups". *Reports@SCM* 1.1 (Oct. 1, 2014), pp. 33–38 (cit. on p. viii).

[Del14b] J. Delgado. "{Whitehead Problems for Words in Fn x Zm". In: *Extended Abstracts Fall 2012*. Ed. by J. González-Meneses, M. Lustig, and E. Ventura. Trends in Mathematics. Springer International Publishing, Jan. 1, 2014, pp. 35–38 (cit. on p. vii).

[DV13] J. Delgado and E. Ventura. "Algorithmic Problems for Free-Abelian Times Free Groups". *Journal of Algebra* 391 (Oct. 1, 2013), pp. 256–283 (cit. on pp. vii, 25, 128, 208).

[DV17a] J. Delgado and E. Ventura. "Stallings Graphs for Free-Abelian by Free Groups". *(preprint)* (2017) (cit. on p. vii).

[DV17b] J. Delgado and E. Ventura. "Twisted-Conjugacy Problem for Free-Abelian Times Free Groups". *(preprint)* (2017) (cit. on p. vii).

[DVZ17] J. Delgado, E. Ventura, and A. Zakharov. "Intersection Problems for Droms Groups". *(preprint)* (2017) (cit. on p. viii).

Contents

Part I

Free times free-abelian groups

1 Definition and generalities

Our first approach to extensions of free and free-abelian groups is through direct products. Despite its naive appearance, the family of finitely generated free-abelian times free groups ($\mathbb{Z}^m \times \mathbb{F}_n$) constitutes not only a natural and interesting starting point by itself, but also a fruitful source of ideas for further generalization.

As we have seen in the introduction, finitely generated free-abelian groups, namely $\mathbb{Z}^m$, are classical and very well known. Free groups, on the other hand, although much wilder and complicated, have also been extensively studied in the literature since more than a hundred years ago.

The goal of this chapter is to investigate the structure of direct products of the form $\mathbb{Z}^m \times \mathbb{F}_n$, namely free-abelian times free groups. At first glance, it may seem that many questions and problems concerning $\mathbb{Z}^m \times \mathbb{F}_n$ will easily reduce to the corresponding questions or problems for $\mathbb{Z}^m$ and $\mathbb{F}_n$; and, in fact, this is the case when the problem considered is easy or rigid enough.

However, many questions that admit simple solutions when dealing with $\mathbb{Z}^m$ or $\mathbb{F}_n$ individually, require far more complex solutions over $\mathbb{Z}^m \times \mathbb{F}_n$. This is the case, for example, when one considers automorphisms; $\mathrm{Aut}(\mathbb{Z}^m \times \mathbb{F}_n)$ naturally contains $\mathrm{GL}_m(\mathbb{Z}) \times \mathrm{Aut}(\mathbb{F}_n)$, but there are many more automorphisms other than those preserving the factors $\mathbb{Z}^m$ and $\mathbb{F}_n$. This causes potential complications when studying problems involving automorphisms: apart from understanding the problem in both the free-abelian and the free parts, one has to be able to control how is it affected by the interaction between the two parts.

Another example of this phenomena is the study of intersections of subgroups. It is well known that every subgroup of $\mathbb{Z}^m$ is finitely generated. This is not true for free groups $\mathbb{F}_n$ with $n \geqslant 2$, but it is also a classical result that all these groups satisfy the Howson property: the intersection of two finitely generated subgroups is again finitely generated. This elementary property fails dramatically in $\mathbb{Z}^m \times \mathbb{F}_n$, when $m \geqslant 1$ and $n \geqslant 2$ (a very easy example reproduced below, already appears in [BK98] attributed to Moldavanski). Consequently, the problem of computing intersections of finitely generated subgroups of $\mathbb{Z}^m \times \mathbb{F}_n$ (including the preliminary decision on whether such intersection is finitely generated or not) becomes considerably more involved than the corresponding problems in $\mathbb{Z}^m$ and $\mathbb{F}_n$ (see Section 2.3).

Throughout the chapter we shall use the following notation and conventions. For $n \geqslant 1$, $[n]$ denotes the set of integers $\{1, \ldots, n\}$. Vectors from $\mathbb{Z}^m$ will always be understood as row vectors, and matrices $\mathbf{M}$ will always be thought as linear maps acting on the right, $\mathbf{v} \mapsto \mathbf{vM}$; accordingly, morphisms will always act on the right of the arguments, $x \mapsto x\alpha$. For notational coherence, we shall use uppercase boldface letters for matrices, and lowercase boldface letters for vectors (moreover, if $w \in \mathbb{F}_n$ then $\mathbf{w} \in \mathbb{Z}^n$ will typically denote its abelianization). We shall use lowercase Greek letters for endomorphisms of free groups, $\phi \colon \mathbb{F}_n \to \mathbb{F}_n$, and uppercase Greek letters for endomorphisms of free-abelian times free groups, $\Phi \colon \mathbb{Z}^m \times \mathbb{F}_n \to \mathbb{Z}^m \times \mathbb{F}_n$.

The chapter is organized as follows. In Section 1.1, we introduce the family of groups we are interested in, and we define several basic notions and properties shared by both families of free-abelian, and free groups, such as the concepts of rank and basis. In Section 1.2 we study the lattice of subgroups within this family, an we provide some computational results which will be useful in the next chapter. Finally, in Section 1.3, we give an explicit description of all automorphisms, monomorphisms and endomorphisms of free-abelian times free groups which will be used in Section 2.4 to study fixed point subgroups, in Section 2.5 to solve the Whitehead problems, and in Section 2.6 to solve the twisted-conjugacy problem within this family.

1.1 Free-abelian times free groups

Let $T = \{t_i : i \in I\}$ and $X = \{x_j : j \in J\}$ be disjoint (possibly empty) sets of symbols, and consider the group G given by the presentation

$$G = \langle T, X \mid [T, T \sqcup X] \rangle ,$$

where $[A, B]$ denotes the set of commutators of all elements from A with all elements from B. Calling $\mathbb{A}$ and $\mathbb{F}$ the subgroups of G generated, respectively, by T and X, it is easy to see that $\mathbb{A}$ is a free-abelian group with basis T, and $\mathbb{F}$ is a free group with basis X. We shall refer to the subgroups $\mathbb{A} = \langle T \rangle$ and $\mathbb{F} = \langle X \rangle$ as the *free-abelian part*, and the *free part* of the group G, respectively. Now, it is straightforward to see that G is the direct product of its free-abelian and free parts, namely

$$G = \langle T, X \mid [T, T \sqcup X] \rangle \simeq \mathbb{A} \times \mathbb{F} . \tag{1.1}$$

Definition 1.1.1. We say that a group is *free-abelian times free* (FATF) if it is isomorphic to one of the form (1.1).

Since each letter in T commutes with all generators, any word on $T \sqcup X$ can be rewritten moving the T-letters, say, to the left. So, every element from G

decomposes as a product of an element from $\mathbb{A}$ and an element from $\mathbb{F}$, in a unique way. After choosing a well ordering on the set T (whose existence for a general T is equivalent — assuming **ZF** — to the Axiom of Choice) we have a natural *normal form* for the elements in G, which we shall write as $t^{\mathbf{a}} w$, where $\mathbf{a} = (a_i)_i \in \bigoplus_{i \in I} \mathbb{Z} = \mathbb{Z}^{\oplus I}$, and $t^{\mathbf{a}}$ stands for the (finite) product $\prod_{i \in I} t_i^{a_i}$ (in the given order for T), and w is a reduced free word on X.

We shall mostly be interested in the finitely generated case, i.e., when T and X are both finite, say $\#I = m$, and $\#J = n$ respectively, with $m, n \geqslant 0$. In this case, $\mathbb{A}$ is the free-abelian group of rank m ($\mathbb{A} = \mathbb{Z}^m$), $\mathbb{F}$ is the free group of rank n ($\mathbb{F} = \mathbb{F}_n$), and our group G becomes

$$G = \mathbb{Z}^m \times \mathbb{F}_n = \langle t_1, \ldots, t_m, x_1, \ldots, x_n \mid t_i t_j = t_j t_i, \, t_i x_k = x_k t_i \rangle, \qquad (1.2)$$

where $i, j \in [m]$ and $k \in [n]$. The *normal form* for an element $g \in G$ is now

$$t^{\mathbf{a}} w := t_1^{a_1} \cdots t_m^{a_m} w(x_1, \ldots, x_n),$$

where $\mathbf{a} = (a_1, \ldots, a_m) \in \mathbb{Z}^m$ is a row vector, and $w = w(x_1, \ldots, x_n)$ is a reduced free word on the alphabet X. Note that the symbol t by itself has no real meaning, but it allows us to convert the notation for the abelian group $\mathbb{Z}^m$ from additive into multiplicative, by moving up the vectors (i.e., the entries of the vectors) to the level of exponents; this will be especially convenient when working in G, a noncommutative group in general.

Remark 1.1.2. Observe that the center of the group G is $\mathbb{A}$ unless $\mathbb{F}$ is infinite cyclic, in which case G is abelian, and so its center is the whole G. This exception will create some technical problems. For example, it causes the ranks of the free-abelian and free parts of G, say m and n, not to be invariants of the group G, since $\mathbb{Z}^m \times \mathbb{F}_1 \simeq \mathbb{Z}^{m+1} \times \mathbb{F}_0$. However, as one may expect, this is the only possible redundancy.

Lemma 1.1.3. *Let $\mathbb{A}, \mathbb{A}'$ be arbitrary free-abelian groups, and let $\mathbb{F}, \mathbb{F}'$ be arbitrary free groups. If $\mathbb{F}$ and $\mathbb{F}'$ are not infinite cyclic, then*

$$\mathbb{A} \times \mathbb{F} \simeq \mathbb{A}' \times \mathbb{F}' \iff \operatorname{rk}(\mathbb{A}) = \operatorname{rk}(\mathbb{A}') \text{ and } \operatorname{rk}(\mathbb{F}) = \operatorname{rk}(\mathbb{F}').$$

Proof. Recall that the center of $\mathbb{A} \times \mathbb{F}$ is $\mathbb{A}$ (here is where $\mathbb{F} \not\simeq \mathbb{Z}$ is needed). Taking quotient by the center we obtain $(\mathbb{A} \times \mathbb{F})/\mathbb{A} \simeq \mathbb{F}$. The claimed result follows immediately. □

Last lemma leads to the following natural definition.

Definition 1.1.4. Let $G = \mathbb{A} \times \mathbb{F}$ be a free-abelian times free group and assume, without loss of generality, that $\mathbb{F} \not\simeq \mathbb{Z}$. Then, the pair of cardinals $(\operatorname{rk} \mathbb{A}, \operatorname{rk} \mathbb{F})$ is an invariant of G, which we shall refer to as the *rank* of G, denoted by $\operatorname{rk}(G)$.

Remark 1.1.5. We allow this abuse of terminology because the rank of G in the usual sense, namely the minimal cardinal of a set of generators, is precisely $\operatorname{rk} \mathbb{A} + \operatorname{rk} \mathbb{F}$: Indeed, let $\operatorname{rk} \mathbb{A} = \kappa$, and $\operatorname{rk} \mathbb{F} = \nu$. It is clear that G is generated by a set of $\kappa + \nu$ elements, and abelianizing, we get $G^{ab} = (\mathbb{A} \times \mathbb{F})^{ab} = \mathbb{A} \oplus \mathbb{F}^{ab}$, a free-abelian group of rank $\kappa + \nu$, so G cannot be generated by less than $\kappa + \nu$ elements.

Definition 1.1.6. Let $G = \mathbb{A} \times \mathbb{F}$ be a free-abelian times free group. A pair (B, A) of subsets of G is called a *basis* of G if the following three conditions are satisfied:

(i) B is an abelian basis of the center of G,

(ii) A is empty, or a free basis of a non-abelian free subgroup of G (note that this excludes the possibility $|A| = 1$),

(iii) $\langle B \cup A \rangle = G$.

In this case we shall also say that B and A are, respectively, the *free-abelian* and *free* components of (B, A). From (i), (ii) and (iii) it follows immediately

(iv) $\langle B \rangle \cap \langle A \rangle = \{1\}$,

(v) $B \cap A = \varnothing$,

since $\langle B \rangle \cap \langle A \rangle$ is contained in the center of G, but no nontrivial element of $\langle A \rangle$ belongs to it.

Usually, we shall abuse notation and just say that $B \cup A$ *is a basis* of G. Note that no information is lost because we can retrieve B as the elements in $B \cup A$ which belong to the center of G, and A as the remaining elements.

Observe that, by (i), (iii) and (iv) in the previous definition, if (B, A) is a basis of a free-abelian times free group G, then $G = \langle B \rangle \times \langle A \rangle$; and by (i) and (ii), $\langle B \rangle$ is a free-abelian group of rank $|B|$, and $\langle A \rangle$ is a free group of rank $|A| \neq 1$; hence, G admits the "canonical" presentation $G = \langle B, A \mid [B, B \sqcup A] \rangle$ and, by Lemma 1.1.3, $\operatorname{rk}(G) = (|B|, |A|)$. In particular, this implies that $(|B|, |A|)$ does not depend on the particular basis (B, A) chosen.

On the other hand, the first obvious example is $T \cup X$ being a basis of the group $G = \langle T, X \mid [T, T \sqcup X] \rangle$ (note that if $|X| \neq 1$ then $B = T$ and $A = X$, but if $|X| = 1$ then $B = T \cup X$ and $A = \varnothing$ due to the technical requirement in Lemma 1.1.3). We have proved the following.

Corollary 1.1.7. *Every free-abelian times free group* G *has a basis. Moreover, every basis* (B, A) *satisfies* $\operatorname{rk}(G) = (|B|, |A|)$. □

1.2 Subgroups of free-abelian times free groups

Let us focus now our attention to subgroups. It is very well known that every subgroup of a free-abelian group is free-abelian; and every subgroup of a free group is again free. These two facts lead, with a straightforward argument, to the same property for free-abelian times free groups (this fact will be crucial for the rest of the chapter).

Proposition 1.2.1. *Any subgroup* $H \leqslant \mathbb{Z}^m \times \mathbb{F}_n$ *admits the decomposition*

$$H \simeq (H \cap \mathbb{Z}^m) \times H\pi_{\mathbb{F}} . \tag{1.3}$$

In particular, subgroups of free-abelian times free group are again free-abelian times free.

Proof. Let T and X be arbitrary disjoint sets, let G be the free-abelian times free group given by presentation (1.1), and let $H \leqslant G$.

If $|X| = 0, 1$ then G is free-abelian, and so H is again free-abelian (with rank less than or equal to that of G); the result follows.

Assume $|X| \geqslant 2$. Let $\mathbb{A} = \langle T \rangle$ and $\mathbb{F} = \langle X \rangle$ be the free-abelian and free parts of G respectively, and let us consider the natural short exact sequence associated to the direct product structure of G:

$$1 \longrightarrow \mathbb{A} \overset{\iota}{\rightarrowtail} \mathbb{A} \times \mathbb{F} = G \overset{\pi_{\mathbb{F}}}{\twoheadrightarrow} \mathbb{F} \longrightarrow 1,$$

where ι is the inclusion, $\pi_{\mathbb{F}}$ is the projection $t^a w \mapsto w$, and therefore $\ker(\pi_{\mathbb{F}}) = \mathbb{A} = \operatorname{im}(\iota)$. Restricting this short exact sequence to $H \leqslant G$, we get

$$1 \longrightarrow \ker(\pi_{\mathbb{F}|H}) \overset{\iota}{\longrightarrow} H \overset{\pi_{\mathbb{F}}}{\longrightarrow} H\pi_{\mathbb{F}} \longrightarrow 1,$$

where $1 \leqslant \ker(\pi_{\mathbb{F}|H}) = H \cap \ker(\pi_{\mathbb{F}}) = H \cap \mathbb{A} \leqslant \mathbb{A}$, and $1 \leqslant H\pi_{\mathbb{F}} \leqslant \mathbb{F}$. Therefore, $\ker(\pi_{\mathbb{F}|H})$ is a free-abelian group, and $H\pi_{\mathbb{F}}$ is a free group. Since $H\pi_{\mathbb{F}}$ is free, $\pi_{\mathbb{F}|H}$ has a splitting

$$H \overset{\sigma}{\longleftarrow} H\pi_{\mathbb{F}}, \tag{1.4}$$

sending back each element of a chosen free basis for $H\pi_{\mathbb{F}}$ to an arbitrary preimage. Hence, σ is injective, $H\pi_{\mathbb{F}}\sigma \leqslant H$ is isomorphic to $H\pi_{\mathbb{F}}$, and straightforward calculations show that the following map is an isomorphism:

$$\begin{array}{rcl} \Theta_\sigma \colon H & \longrightarrow & \ker(\pi_{\mathbb{F}|H}) \times H\pi_{\mathbb{F}}\sigma \\ h & \longmapsto & \left(h(h\pi_{\mathbb{F}}\sigma)^{-1}, h\pi_{\mathbb{F}}\sigma\right). \end{array} \tag{1.5}$$

Thus $H \simeq \ker(\pi_{\mathbb{F}|H}) \times H\pi_{\mathbb{F}}\sigma$ is free-abelian times free and the result is proven. □

This proof shows a particular way of decomposing H into a direct product of a free-abelian subgroup and a free subgroup, which depends on the chosen splitting σ, namely

$$H = (H \cap \mathbb{A}) \times H\pi_{\mathbb{F}}\sigma. \tag{1.6}$$

We call the subgroups $H \cap \mathbb{A}$ and $H\pi_{\mathbb{F}}\sigma$, respectively, the *free-abelian* and *free parts of* H, *with respect to the splitting* σ. Note that the free-abelian and free parts of the subgroup $H = G$ with respect to the natural inclusion $\mathbb{F}_n \ni w \mapsto t^0 w \in G$ coincide with what we called the free-abelian and free parts of G.

Furthermore, Proposition 1.2.1 and the decomposition (1.6) give a characterization of the bases, rank, and all possible isomorphism classes of such an arbitrary subgroup H.

Corollary 1.2.2. *With the above notation, a subset* $E \subseteq H \leqslant G = \mathbb{A} \times \mathbb{F}$ *is a basis of* H *if and only if*

$$E = E_{\mathbb{A}} \sqcup E_{\mathbb{F}},$$

where $E_{\mathbb{A}}$ *is an abelian basis of* $H \cap \mathbb{A}$, *and* $E_{\mathbb{F}}$ *is a free basis of* $H\pi_{\mathbb{F}}\sigma$, *for a certain splitting* σ *as in* (1.4).

Proof. The implication to the left is straightforward, with $E = B \sqcup A$, and $(B, A) = (E_{\mathbb{A}}, E_{\mathbb{F}})$ except for the case $\operatorname{rk}(\mathbb{F}) = 1$, where we have $(B, A) = (E_{\mathbb{A}} \sqcup E_{\mathbb{F}}, \varnothing)$.

Suppose now that $E = B \sqcup A$ is a basis of H in the sense of Definition 1.1.6, and let us look at the decomposition (1.6), for suitable σ. If $\operatorname{rk}(H\pi_{\mathbb{F}}) = 1$, then H is abelian, B is an abelian basis for H, $A = \varnothing$, and all but exactly one of the elements in B belong to $H \cap \mathbb{A}$ (i.e., have normal forms using only letters from T); in this case the result is clear, taking $E_{\mathbb{F}}$ to be just that special element. Otherwise, $\mathbb{A}(H) = H \cap \mathbb{A}$ having B as an abelian basis; take $E_{\mathbb{A}} = B$ and $E_{\mathbb{F}} = A$. It is clear that the projection $\pi_{\mathbb{F}} \colon H \twoheadrightarrow H\pi_{\mathbb{F}}$, $t^a u \mapsto u$, restricts to an isomorphism $\pi_{\mathbb{F}|\langle A \rangle} \colon \langle A \rangle \to H\pi_{\mathbb{F}}$ since no nontrivial element in $\langle A \rangle$ belongs to $\ker \pi_{\mathbb{F}} = H \cap \mathbb{A}$. Therefore, taking $\sigma = \pi_{\mathbb{F}|\langle A \rangle}^{-1}$, $E_{\mathbb{F}}$ is a free basis of $H\pi_{\mathbb{F}}\sigma$. □

Corollary 1.2.3. *Let* G *be the free-abelian times free group given by presentation* (1.1), *and let* $\operatorname{rk}(G) = (\kappa, \sigma)$. *Every subgroup* $H \leqslant G$ *is again free-abelian times free with* $\operatorname{rk}(H) = (\kappa', \sigma')$ *where,*

(i) *in case of* $\sigma = 0$: $0 \leqslant \kappa' \leqslant \kappa$ *and* $\sigma' = 0$;

(ii) *in case of* $\sigma \geqslant 2$: *either* $0 \leqslant \kappa' \leqslant \kappa + 1$ *and* $\sigma' = 0$, *or* $0 \leqslant \kappa' \leqslant \kappa$ *and* $0 \leqslant \sigma' \leqslant \max\{\sigma, \aleph_0\}$ *and* $\sigma' \neq 1$.

Furthermore, for every such (κ', σ'), *there is a subgroup* $H \leqslant G$ *such that* $\operatorname{rk}(H) = (\kappa', \sigma')$. □

Throughout the rest of the chapter, we shall concentrate on the finitely generated case. From Proposition 1.2.1 we can easily deduce the following corollary, which will be useful later.

Corollary 1.2.4. *A subgroup* $H \leqslant \mathbb{Z}^m \times \mathbb{F}_n$ *is finitely generated if and only if its projection to the free part* $H\pi_{\mathbb{F}}$ *is finitely generated.* □

Remark 1.2.5. Note that the proof of Proposition 1.2.1, at least in the finitely generated case, is completely algorithmic; i.e., if H is given by a finite set of generators, one can effectively choose a splitting σ, and compute a basis of the free-abelian and free parts of H (w.r.t. σ). This fact will be crucial for the rest of the chapter, and we make it more precise in the following proposition.

Proposition 1.2.6. *Let* $G = \mathbb{Z}^m \times \mathbb{F}_n$ *be a finitely generated free-abelian times free group. There is an algorithm which, given a subgroup* $H \leqslant G$ *by a finite family of generators, computes a basis for* H *and writes both, the new elements in terms of the old generators, and the old generators in terms of the new basis.*

Proof. If $n = |X| = 0, 1$ then G is free-abelian and the problem is a straightforward exercise in linear algebra. So, let us assume $n \geqslant 2$.

We are given a finite set of generators for H, say $t^{\mathbf{c_1}}w_1, \ldots, t^{\mathbf{c_p}}w_p$, where $p \geqslant 1$, $\mathbf{c_1}, \ldots, \mathbf{c_p} \in \mathbb{Z}^m$ are row vectors, and $w_1, \ldots, w_p \in \mathbb{F}_n$ are reduced words on $X = \{x_1, \ldots, x_n\}$. Applying suitable Nielsen transformations, see [LS01], we can algorithmically transform the p-tuple $(w_1, \ldots, w_p)$ of elements from $\mathbb{F}_n$, into another of the form $(u_1, \ldots, u_{n'}, 1, \ldots, 1)$, where $\{u_1, \ldots, u_{n'}\}$ is a free basis of $\langle w_1, \ldots, w_p \rangle = H\pi_{\mathbb{F}}$, and $0 \leqslant n' \leqslant p$. Furthermore, reading along the Nielsen process performed, we can effectively compute expressions of the new elements as words on the old generators, say $u_j = \eta_j(w_1, \ldots, w_p)$, $j \in [n']$, as well as expressions of the old generators in terms of the new free basis, say $w_i = \nu_i(u_1, \ldots, u_{n'})$, for $i \in [p]$.

Now, the map $\sigma \colon H\pi_{\mathbb{F}} \to H$, $u_j \mapsto \eta_j(t^{\mathbf{c_1}}w_1, \ldots, t^{\mathbf{c_p}}w_p)$ can serve as a splitting in the proof of Proposition 1.2.1, since $\eta_j(t^{\mathbf{c_1}}w_1, \ldots, t^{\mathbf{c_p}}w_p) = t^{\mathbf{a_j}}\eta_j(w_1, \ldots, w_p) = t^{\mathbf{a_j}}u_j \in H$, where $\mathbf{a_j}$, $j \in [n']$, are integral linear combinations of $\mathbf{c_1}, \ldots, \mathbf{c_p}$.

It only remains to compute an abelian basis for $\ker(\pi_{\mathbb{F}|H}) = H \cap \mathbb{Z}^m$. For each one of the given generators $h = t^{\mathbf{c_i}}w_i$, calculate $h(h\pi_{\mathbb{F}}\sigma)^{-1} = t^{\mathbf{d_i}}$ (here, we shall need the words ν_i computed before). Using the isomorphism Θ_σ from the proof of Proposition 1.2.1, we deduce that $\{t^{\mathbf{d_1}}, \ldots, t^{\mathbf{d_p}}\}$ generate $H \cap \mathbb{Z}^m$; it only remains to use a standard linear algebra procedure, to extract from here an abelian basis $\{t^{\mathbf{b_1}}, \ldots, t^{\mathbf{b_{m'}}}\}$ for $H \cap \mathbb{Z}^m$. Clearly, $0 \leqslant m' \leqslant m$.

We immediately get a basis (B, A) for H (with just a small technical caution): if $n' \neq 1$, take $B = \{t^{\mathbf{b_1}}, \ldots, t^{\mathbf{b_{m'}}}\}$ and $A = \{t^{\mathbf{a_1}}u_1, \ldots, t^{\mathbf{a_{n'}}}u_{n'}\}$; and if $n' = 1$ take $B = \{t^{\mathbf{b_1}}, \ldots, t^{\mathbf{b_{m'}}}, t^{\mathbf{a_1}}u_1\}$ and $A = \varnothing$.

On the other hand, as a side product of the calculations done, we have the expressions $t^{\mathbf{a_j}}u_j = \eta_j(t^{\mathbf{c_1}}w_1, \ldots, t^{\mathbf{c_p}}w_p)$, $j \in [n']$. And we can also compute expressions of the $t^{\mathbf{b_i}}$'s in terms of the $t^{\mathbf{d_i}}$'s, and of the $t^{\mathbf{d_i}}$'s in terms of the $t^{\mathbf{c_i}}w_i$'s. Hence we can compute expressions for each one of the new elements in terms of the old generators.

For the other direction, we also have the expressions $w_i = \nu_i(u_1, \ldots, u_{n'})$, for $i \in [p]$. Hence, $\nu_i(t^{\mathbf{a_1}}u_1, \ldots, t^{\mathbf{a_{n'}}}u_{n'}) = t^{\mathbf{e_i}}w_i$ for some $\mathbf{e_i} \in \mathbb{Z}^m$. But

$$H \ni (t^{\mathbf{c_i}}w_i)(t^{\mathbf{e_i}}w_i)^{-1} = t^{\mathbf{c_i}-\mathbf{e_i}} \in \mathbb{Z}^m,$$

so we can compute integers $\lambda_1, \ldots, \lambda_{m'}$ such that $\mathbf{c_i} - \mathbf{e_i} = \lambda_1\mathbf{b_1} + \cdots + \lambda_{m'}\mathbf{b_{m'}}$. Thus,

$$t^{\mathbf{c_i}}w_i = t^{\mathbf{c_i}-\mathbf{e_i}}t^{\mathbf{e_i}}w_i = t^{\lambda_1\mathbf{b_1}+\cdots+\lambda_{m'}\mathbf{b_{m'}}}t^{\mathbf{e_i}}w_i = (t^{\mathbf{b_1}})^{\lambda_1}\cdots(t^{\mathbf{b_{m'}}})^{\lambda_{m'}}\nu_i(t^{\mathbf{a_1}}u_1, \ldots, t^{\mathbf{a_{n'}}}u_{n'}),$$

for $i \in [p]$. This concludes the proof. □

To conclude this section, let us introduce some notation that will be useful later. Let H be a finitely generated subgroup of $G = \mathbb{Z}^m \times \mathbb{F}_n$, and consider a basis for H,

$$\{t^{\mathbf{b_1}}, \ldots, t^{\mathbf{b_{m'}}},\ t^{\mathbf{a_1}}u_1, \ldots, t^{\mathbf{a_{n'}}}u_{n'}\}, \tag{1.7}$$

where $0 \leqslant m' \leqslant m$, $\{\mathbf{b_1}, \ldots, \mathbf{b_{m'}}\}$ is an abelian basis of $H \cap \mathbb{Z}^m \leqslant \mathbb{Z}^m$, $0 \leqslant n'$, $\mathbf{a_1}, \ldots, \mathbf{a_{n'}} \in \mathbb{Z}^m$, and $\{u_1, \ldots, u_{n'}\}$ is a free basis of $H\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$. Let us denote by L the abelian subgroup $\langle \mathbf{b_1}, \ldots, \mathbf{b_{m'}} \rangle \leqslant \mathbb{Z}^m$ (with additive notation, i.e., these are true vectors with m integral coordinates each), and let us denote by $\mathbf{A}$ the $n' \times m$ integral matrix whose rows are the $\mathbf{a_i}$'s,

$$\mathbf{A} = \begin{bmatrix} \mathbf{a_1} \\ \vdots \\ \mathbf{a_{n'}} \end{bmatrix} \in \mathcal{M}_{n' \times m}(\mathbb{Z}).$$

If ω is a word on n' letters (i.e., an element of the abstract free group $\mathbb{F}_{n'}$), we will denote by $\omega(u_1, \ldots, u_{n'})$ the element of $H\pi_{\mathbb{F}}$ obtained by replacing the i-th letter in ω by u_i, $i \in [n']$. And we shall use boldface, $\boldsymbol{\omega}$, to denote the abstract abelianization of ω, which is an integral vector with n' coordinates, $\boldsymbol{\omega} \in \mathbb{Z}^{n'}$ (not to be confused with the image of $\omega(u_1, \ldots, u_{n'}) \in \mathbb{F}_n$ under the abelianization map $\mathbb{F}_n \twoheadrightarrow \mathbb{Z}^n$). Straightforward calculations show the result below.

Lemma 1.2.7. *With the previous notations, we have*

$$H = \{\, t^{\mathbf{a}}\, \omega(u_1, \ldots, u_{n'}) \,:\, \omega \in \mathbb{F}_{n'}, \mathbf{a} \in \boldsymbol{\omega}\mathbf{A} + L \,\},$$

a convenient description of H. □

Definition 1.2.8. Given a subgroup $H \leqslant \mathbb{Z}^m \times \mathbb{F}_n$, and an element $w \in \mathbb{F}_n$, we define the *abelian completion of w in* H as

$$\mathcal{C}_H(w) = \{\mathbf{a} \in \mathbb{Z}^m \,:\, t^{\mathbf{a}} w \in H\} \subseteq \mathbb{Z}^m.$$

Corollary 1.2.9. *With the above notation, for every $w \in \mathbb{F}_n$ we have*

(a) *if $w \notin H\pi_{\mathbb{F}}$, then $\mathcal{C}_H(w) = \varnothing$,*

(b) *if $w \in H\pi_{\mathbb{F}}$, then $\mathcal{C}_H(w) = \boldsymbol{\omega}\mathbf{A} + L$, where $\boldsymbol{\omega}$ is the abelianization of the word $\omega \in \mathbb{F}_{n'}$ which expresses $w \in \mathbb{F}_n$ in terms of the free basis $\{u_1, \ldots, u_{n'}\}$, i.e., $w = \omega(u_1, \ldots, u_{n'})$ (note the difference between w and ω).*

So, the completion $\mathcal{C}_H(w) \subseteq \mathbb{Z}^m$ is either empty, or a coset of $L = H \cap \mathbb{Z}^m$. □

1.3 Endomorphisms

In this section we will study the endomorphisms of a finitely generated free-abelian times free group $G = \mathbb{Z}^m \times \mathbb{F}_n$ (with the notation from presentation (1.2)). Without loss of generality, we assume $n \neq 1$.

To clarify notation, we shall use lowercase Greek letters to denote endomorphisms of $\mathbb{F}_n$, and uppercase Greek letters to denote endomorphisms of $G = \mathbb{Z}^m \times \mathbb{F}_n$.

Proposition 1.3.1. *Let $G = \mathbb{Z}^m \times \mathbb{F}_n$ with $n \neq 1$. The following is a complete list of all endomorphisms of G:*

I. $\Psi_{\phi,\mathbf{Q},\mathbf{P}} \,:\, t^{\mathbf{a}} u \mapsto t^{\mathbf{aQ}+\mathbf{uP}} u\phi$, *where $\phi \in \operatorname{End}(\mathbb{F}_n)$, $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$, and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$.*

II. $\Psi_{z,\mathbf{l},\mathbf{h},\mathbf{Q},\mathbf{P}} \,:\, t^{\mathbf{a}} u \mapsto t^{\mathbf{aQ}+\mathbf{uP}} z^{\mathbf{al}^{\mathsf{T}}+\mathbf{uh}^{\mathsf{T}}}$, *where $1 \neq z \in \mathbb{F}_n$ is not a proper power, $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$, $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$, $\mathbf{0} \neq \mathbf{l} \in \mathbb{Z}^m$, and $\mathbf{h} \in \mathbb{Z}^n$.*

(In both cases, $\mathbf{u} \in \mathbb{Z}^n$ denotes the abelianization of the word $u \in \mathbb{F}_n$.) We will refer to them as endomorphisms of type I and II *respectively.*

Proof. It is routine to check that all maps of types I and II are, indeed, endomorphisms of G.

To see that this is a complete list, let $\Psi \colon G \to G$ be an arbitrary endomorphism of G. Looking at the normal form of the images of the x_i's and t_j's, we have

$$\Psi \colon \left\{ \begin{array}{rcl} x_i & \longmapsto & t^{\mathbf{p}_i} w_i \\ t_j & \longmapsto & t^{\mathbf{q}_j} z_j\,, \end{array} \right. \tag{1.8}$$

where $\mathbf{p}_i, \mathbf{q}_j \in \mathbb{Z}^m$ and $w_i, z_j \in \mathbb{F}_n$, $i \in [n]$, $j \in [m]$. Let us distinguish two cases.

Case 1: $z_j = 1$ for all $j \in [m]$. Denoting by ϕ the endomorphism of $\mathbb{F}_n$ given by $x_i \mapsto w_i$, and by $\mathbf{P}$ and $\mathbf{Q}$ the following integral matrices (of sizes $n \times m$ and $m \times m$, respectively)

$$\mathbf{P} = \begin{bmatrix} p_{11} & \cdots & p_{1m} \\ \vdots & \ddots & \vdots \\ p_{n1} & \cdots & p_{nm} \end{bmatrix} = \begin{bmatrix} \mathbf{p_1} \\ \vdots \\ \mathbf{p_n} \end{bmatrix} \quad \text{and} \quad \mathbf{Q} = \begin{bmatrix} q_{11} & \cdots & q_{1m} \\ \vdots & \ddots & \vdots \\ q_{m1} & \cdots & q_{mm} \end{bmatrix} = \begin{bmatrix} \mathbf{q_1} \\ \vdots \\ \mathbf{q_m} \end{bmatrix},$$

we can write

$$\Psi \colon \left\{ \begin{array}{lcl} u & \longmapsto & t^{\mathbf{uP}}\, u\phi \\ t^{\mathbf{a}} & \longmapsto & t^{\mathbf{aQ}}, \end{array} \right.$$

where $u \in \mathbb{F}_n$, and $\mathbf{a} \in \mathbb{Z}^m$. So, $(t^{\mathbf{a}}\, u)\Psi = t^{\mathbf{aQ}+\mathbf{uP}}\, u\phi$, and $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$ of type I.

Case 2: $z_k \neq 1$ for some $k \in [m]$. In order to Ψ to be well defined, $t^{\mathbf{p_i}}\, w_i$ and $t^{\mathbf{q_j}}\, z_j$ must all commute with $t^{\mathbf{q_k}}\, z_k$, and so w_i and z_j with $z_k \neq 1$, for all $i \in [n]$ and $j \in [m]$. This means that $w_i = z^{h_i}$, $z_j = z^{l_j}$ for some integers $h_i, l_j \in \mathbb{Z}$, $i \in [n]$, $j \in [m]$, with $l_k \neq 0$, and some $z \in \mathbb{F}_n$ not being a proper power. Hence,

$$(t^{\mathbf{a}}\, u)\Psi \;=\; (t^{\mathbf{a}}\Psi)\,(u\Psi) \;=\; (t^{\mathbf{aQ}}\, z^{\mathbf{al}^{\mathsf{T}}})\,(t^{\mathbf{uP}}\, z^{\mathbf{uh}^{\mathsf{T}}}) \;=\; t^{\mathbf{aQ}+\mathbf{uP}}\, z^{\mathbf{al}^{\mathsf{T}}+\mathbf{uh}^{\mathsf{T}}},$$

and $\Psi = \Psi_{z,\mathbf{l},\mathbf{h},\mathbf{Q},\mathbf{P}}$ of type II. This completes the proof. □

Remark 1.3.2. Note that if $n = 0$ then type I and type II endomorphisms do coincide. Otherwise, type II endomorphisms will be seen to be neither injective nor surjective.

The following proposition gives a quite natural characterization of which endomorphisms of type I are injective, and which are surjective. It is important to note that the matrix $\mathbf{P}$ plays absolutely no role in this matter.

Proposition 1.3.3. *Let Ψ be an endomorphism of $G = \mathbb{Z}^m \times \mathbb{F}_n$, with $n \geqslant 2$. Then,*

(i) *Ψ is a monomorphism if and only if it is of type I, $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$, with ϕ a monomorphism of $\mathbb{F}_n$, and $\det(\mathbf{Q}) \neq 0$,*

(ii) *Ψ is an epimorphism if and only if it is of type I, $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$, with ϕ an epimorphism of $\mathbb{F}_n$, and $\det(\mathbf{Q}) = \pm 1$.*

(iii) *Ψ is an automorphism if and only if it is of type I, $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$, with $\phi \in \mathrm{Aut}(\mathbb{F}_n)$ and $\mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})$; in this case, $(\Psi_{\phi,\mathbf{Q},\mathbf{P}})^{-1} = \Psi_{\phi^{-1},\mathbf{Q}^{-1},-\mathbf{M}^{-1}\mathbf{P}\mathbf{Q}^{-1}}$, where $\mathbf{M} \in \mathrm{GL}_n(\mathbb{Z})$ is the abelianization of ϕ.*

Proof. (i) Suppose that Ψ is injective. Then Ψ cannot be of type II since, if it were, the commutator of any two elements in $\mathbb{F}_n$ ($n \geqslant 2$) would be in the kernel of Ψ. Hence, $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$ for some $\phi \in \mathrm{End}(\mathbb{F}_n)$, $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$, and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$.

Since $t^{\mathbf{a}}\Psi = t^{\mathbf{aQ}}$, the injectivity of Ψ implies that of $\mathbf{a} \mapsto \mathbf{aQ}$; hence, $\det(\mathbf{Q}) \neq 0$. Finally, in order to prove the injectivity of ϕ, let $u \in \mathbb{F}_n$ with $u\phi = 1$. Note that the endomorphism of $\mathbb{Q}^m$ given by $\mathbf{Q}$ is invertible so, in particular, there exists $\mathbf{v} \in \mathbb{Q}^m$ such that $\mathbf{vQ} = \mathbf{uP}$; write $\mathbf{v} = \frac{1}{b}\mathbf{a}$ for some $\mathbf{a} \in \mathbb{Z}^m$ and $b \in \mathbb{Z}$, $b \neq 0$, and we have $\mathbf{aQ} = b\,\mathbf{vQ} = b\,\mathbf{uP}$; thus, $(t^{\mathbf{a}}u^{-b})\Psi = t^{\mathbf{aQ}}(t^{\mathbf{uP}}1)^{-b} = t^{\mathbf{aQ}-b\mathbf{uP}} = t^{\mathbf{0}} = 1$. Hence, $t^{\mathbf{a}}u^{-b} = 1$ and so, $u = 1$.

Conversely, let $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$ be of type I, with ϕ a monomorphism of $\mathbb{F}_n$ and $\det(\mathbf{Q}) \neq \mathbf{0}$, and let $t^{\mathbf{a}}\,u \in G$ be such that $1 = (t^{\mathbf{a}}\,u)\Psi = t^{\mathbf{aQ}+\mathbf{uP}}\,u\phi$. Then, $u\phi = 1$ and so, $u = 1$; and $\mathbf{0} = \mathbf{aQ} + \mathbf{uP} = \mathbf{aQ}$ and so, $\mathbf{a} = \mathbf{0}$. Hence, Ψ is injective.

(ii) Suppose that Ψ is surjective. Since the image of an endomorphism of type II followed by the projection $\pi_{\mathbb{F}}$ onto $\mathbb{F}_n$, $n \geqslant 2$, is contained in $\langle z \rangle$ (and so is cyclic), Ψ cannot be of type II. Hence, $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$ for some $\phi \in \mathrm{End}(\mathbb{F}_n)$, $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$, and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$.

Given $v \in \mathbb{F}_n \leqslant G$ there must be $t^{\mathbf{a}}u \in G$ such that $(t^{\mathbf{a}}u)\Psi = v$ and so $u\phi = v$. Thus $\phi\colon \mathbb{F}_n \to \mathbb{F}_n$ is surjective. On the other hand, for every $j \in [m]$, let δ_j be the canonical vector of $\mathbb{Z}^m$ with 1 at coordinate j, and let $t^{\mathbf{b}_j}u_j \in G$ be a preimage by Ψ of $t_j = t^{\delta_j}$. We have $(t^{\mathbf{b}_j}u_j)\Psi = t^{\delta_j}$, i.e., $u_j\phi = 1$, $\mathbf{u_j} = \mathbf{0}$ and $\mathbf{b_j Q} = \mathbf{b_j Q} + \mathbf{u_j P} = \delta_j$. This means that the matrix $\mathbf{B}$ with rows $\mathbf{b_j}$ satisfies $\mathbf{BQ} = \mathbf{I_m}$ and thus, $\det(\mathbf{Q}) = \pm 1$.

Conversely, let $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$ be of type I, with ϕ being an epimorphism of $\mathbb{F}_n$ and $\det(\mathbf{Q}) = \pm 1$. Since $\mathbb{F}_n$ is hopfian, $\phi \in \mathrm{Aut}(\mathbb{F}_n)$ and we can consider $\Upsilon = \Psi_{\phi^{-1},\mathbf{Q}^{-1},-\mathbf{M}^{-1}\mathbf{P}\mathbf{Q}^{-1}}$, where $\mathbf{M} \in \mathrm{GL}_n(\mathbb{Z})$ is the abelianization of ϕ. For every $t^{\mathbf{a}}\,u \in G$, we have

$$(t^{\mathbf{a}}\,u)\Upsilon\Psi = \big(t^{\mathbf{a}\mathbf{Q}^{-1}-\mathbf{u}\mathbf{M}^{-1}\mathbf{P}\mathbf{Q}^{-1}}\,(u\phi^{-1})\big)\Psi = t^{\mathbf{a}-\mathbf{u}\mathbf{M}^{-1}\mathbf{P}+\mathbf{u}\mathbf{M}^{-1}\mathbf{P}}\,u = t^{\mathbf{a}}\,u\,.$$

Hence, Ψ is surjective.

(iii) The equivalence is a direct consequence of (i) and (ii). To see the actual value of Ψ^{-1} it remains to compute the composition in the reverse order:

$$(t^{\mathbf{a}}\,u)\Psi\Upsilon = \big(t^{\mathbf{aQ}+\mathbf{uP}}(u\phi)\big)\Upsilon = t^{\mathbf{a}+\mathbf{uP}\mathbf{Q}^{-1}-\mathbf{u}\mathbf{M}\mathbf{M}^{-1}\mathbf{P}\mathbf{Q}^{-1}}\,u = t^{\mathbf{a}}\,u\,. \qquad \square$$

From these characterizations for an endomorphism to be mono, epi or auto, we immediately have the following corollary.

Corollary 1.3.4. *$\mathbb{Z}^m \times \mathbb{F}_n$ is hopfian and not cohopfian.* □

Remark 1.3.5. This result was already known: in [Gre90] and [Hum94] it is proved that finitely generated partially commutative groups (this includes groups of the form $G = \mathbb{Z}^m \times \mathbb{F}_n$) are residually finite and so, hopfian. However, our proof is more direct and explicit in the sense of giving complete characterizations of the injectivity and surjectivity of a given endomorphism of G. We remark that the hopfian property for $\mathbb{Z}^m \times \mathbb{F}_n$ does not follow directly from that of free-abelian and free groups (both very well known): in [Tyr71], the author constructs a direct product of two hopfian groups which is *not* hopfian.

Below, we detail the expressions obtained after composing endomorphisms of types I and II in all possible ways.

Lemma 1.3.6. *Composition in the monoid* $\mathrm{End}(\mathbb{Z}^m \times \mathbb{F}_n)$ *can be summarized in the following four cases:*

(a) Type I *followed by* type I*:*

$$\Psi_{\phi_1,\mathbf{Q_1},\mathbf{P_1}} \Psi_{\phi_2,\mathbf{Q_2},\mathbf{P_2}} = \Psi_{\phi_1\phi_2\,,\,\mathbf{Q_1Q_2}\,,\,\mathbf{P_1Q_2}+\mathbf{M_1P_2}}\,. \tag{1.9}$$

(b) Type I *followed by* type II*:*

$$\Psi_{\phi_1,\mathbf{Q_1},\mathbf{P_1}} \Psi_{w_2,\mathbf{l_2},\mathbf{h_2},\mathbf{Q_2},\mathbf{P_2}} = \Psi_{w_2\,,\,\mathbf{l_2Q_1^T}\,,\,\mathbf{l_2P_1^T}+\mathbf{h_2M_1^T}\,,\,\mathbf{Q_1Q_2}\,,\,\mathbf{P_1Q_2}+\mathbf{M_1P_2}}\,. \tag{1.10}$$

(c) Type II *followed by* type I*:*

$$\Psi_{w_1,\mathbf{l_1},\mathbf{h_1},\mathbf{Q_1},\mathbf{P_1}} \Psi_{\phi_2,\mathbf{Q_2},\mathbf{P_2}} = \Psi_{w_1\phi_2\,,\,\mathbf{l_1}\,,\,\mathbf{h_1}\,,\,\mathbf{Q_1Q_2}+\mathbf{l_1^T w_1 P_2}\,,\,\mathbf{P_1Q_2}+\mathbf{h_1^T w_1 P_2}}\,. \tag{1.11}$$

(d) Type II *followed by* type II*:*

$$\begin{aligned}\Psi_{w_1,\mathbf{l_1},\mathbf{h_1},\mathbf{Q_1},\mathbf{P_1}} \Psi_{w_2,\mathbf{l_2},\mathbf{h_2},\mathbf{Q_2},\mathbf{P_2}} &= \\ &= \Psi_{w_2\,,\,\mathbf{l_2Q_1^T}+\mathbf{h_2w_1^Tl_1}\,,\,\mathbf{l_2P_1^T}+\mathbf{h_2w_1^Th_1}\,,\,\mathbf{Q_1Q_2}+\mathbf{l_1^Tw_1P_2}\,,\,\mathbf{P_1Q_2}+\mathbf{h_1^Tw_1P_2}}\,.\end{aligned} \tag{1.12}$$

where, $\mathbf{M_1}$ *denotes the (matrix of) the abelianization of* $\phi_1 \in \mathrm{End}(\mathbb{F}_n)$*.* □

So, composition in $\mathrm{End}(\mathbb{Z}^m \times \mathbb{F}_n)$ is closed by both types of endomorphisms, whereas type II absorbs crossed products. In particular, since the identity map clearly belongs to type I, we conclude that type I endomorphisms constitute a submonoid of $\mathrm{End}(\mathbb{Z}^m \times \mathbb{F}_n)$ which, moreover, contains the cases most interesting to us (recall Proposition 1.3.3).

For later use, next lemma summarizes how to operate type I endomorphisms (compose, invert and take a power); it can be easily proved by following routine computations. can find similar expressions for the composition of two type II endomorphisms, or one of each.

Lemma 1.3.7. *Let* $\Psi_{\phi,\mathbf{Q},\mathbf{P}}$ *and* $\Psi_{\phi',\mathbf{Q'},\mathbf{P'}}$ *be two type I endomorphisms of* $G = \mathbb{Z}^m \times \mathbb{F}_n$*,* $n \neq 1$*, and denote by* $\mathbf{M} \in \mathcal{M}_n(\mathbb{Z})$ *the (matrix of the) abelianization of* $\phi \in \mathrm{End}(\mathbb{F}_n)$*. Then,*

(i) $\Psi_{\phi,\mathbf{Q},\mathbf{P}} \cdot \Psi_{\phi',\mathbf{Q}',\mathbf{P}'} = \Psi_{\phi\phi',\mathbf{Q}\mathbf{Q}',\mathbf{P}\mathbf{Q}'+\mathbf{M}\mathbf{P}'}$,

(ii) *for all* $k \geqslant 1$, $(\Psi_{\phi,\mathbf{Q},\mathbf{P}})^k = \Psi_{\phi^k,\mathbf{Q}^k,\mathbf{P}_k}$, *where* $\mathbf{P}_\mathbf{k} = \sum_{i=1}^{k} \mathbf{M}^{i-1}\mathbf{P}\,\mathbf{Q}^{k-i}$,

(iii) $\Psi_{\phi,\mathbf{Q},\mathbf{P}}$ *is invertible if and only if* $\phi \in \operatorname{Aut}(\mathbb{F}_n)$ *and* $\mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})$*; in this case,* $(\Psi_{\phi,\mathbf{Q},\mathbf{P}})^{-1} = \Psi_{\phi^{-1},\mathbf{Q}^{-1},-\mathbf{M}^{-1}\mathbf{P}\mathbf{Q}^{-1}}$.

(iv) *For every* $\mathbf{a} \in \mathbb{Z}^m$ *and* $u \in \mathbb{F}_n$*, the right conjugation by* $t^{\mathbf{a}}u$ *is* $\Gamma_{t^{\mathbf{a}}u} = \Psi_{\gamma_u,\mathbf{I}_m,\mathbf{0}}$*, where* γ_u *is the right conjugation by* u *in* $\mathbb{F}_n$*,* $v \mapsto u^{-1}vu$*,* $\mathbf{I}_\mathbf{m}$ *is the identity matrix of size* m*, and* $\mathbf{0}$ *is the zero matrix of size* $n \times m$*.* □

In the rest of the section, we shall use this information to derive the structure of $\operatorname{Aut}(G)$, where $G = \mathbb{Z}^m \times \mathbb{F}_n$, $m \geqslant 1$, $n \geqslant 2$.

Theorem 1.3.8. *For* $G = \mathbb{Z}^m \times \mathbb{F}_n$*, with* $m \geqslant 1$ *and* $n \geqslant 2$*, the group* $\operatorname{Aut}(G)$ *is isomorphic to the semidirect product* $\mathcal{M}_{n\times m}(\mathbb{Z}) \rtimes (\operatorname{Aut}(\mathbb{F}_n) \times \mathrm{GL}_m(\mathbb{Z}))$ *with respect to the natural action. In particular,* $\operatorname{Aut}(G)$ *is finitely presented.*

Proof. First of all note that, for every $\phi, \phi' \in \operatorname{Aut}(\mathbb{F}_n)$, every $\mathbf{Q}, \mathbf{Q}' \in \mathrm{GL}_m(\mathbb{Z})$, and every $\mathbf{P}, \mathbf{P}' \in \mathcal{M}_{n\times m}(\mathbb{Z})$, we have

$$\begin{aligned}
\Psi_{\phi,\mathbf{I}_m,\mathbf{0}} \cdot \Psi_{\phi',\mathbf{I}_m,\mathbf{0}} &= \Psi_{\phi\phi',\mathbf{I}_m,\mathbf{0}}\,,\\
\Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{0}} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q}',\mathbf{0}} &= \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q}\mathbf{Q}',\mathbf{0}}\,,\\
\Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}'} &= \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}+\mathbf{P}'}\,.
\end{aligned}$$

Hence, the three groups $\operatorname{Aut}(\mathbb{F}_n)$, $\mathrm{GL}_m(\mathbb{Z})$, and $\mathcal{M}_{n\times m}(\mathbb{Z})$ (this last one with the addition of matrices), are all subgroups of $\operatorname{Aut}(G)$ via the three natural inclusions: $\phi \mapsto \Psi_{\phi,\mathbf{I}_m,\mathbf{0}}$, $\mathbf{Q} \mapsto \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{0}}$, and $\mathbf{P} \mapsto \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}}$, respectively. Furthermore, for every $\phi \in \operatorname{Aut}(\mathbb{F}_n)$ and every $\mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})$, it is clear that $\Psi_{\phi,\mathbf{I}_m,\mathbf{0}} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{0}} = \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{0}} \cdot \Psi_{\phi,\mathbf{I}_m,\mathbf{0}}$; hence $\operatorname{Aut}(\mathbb{F}_n) \times \mathrm{GL}_m(\mathbb{Z})$ is a subgroup of $\operatorname{Aut}(G)$ in the natural way.

On the other hand, for every $\phi \in \operatorname{Aut}(\mathbb{F}_n)$, every $\mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})$, and every $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$, we have

$$(\Psi_{\phi,\mathbf{I}_m,\mathbf{0}})^{-1} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}} \cdot \Psi_{\phi,\mathbf{I}_m,\mathbf{0}} = \Psi_{\phi^{-1},\mathbf{I}_m,\mathbf{0}} \cdot \Psi_{\phi,\mathbf{I}_m,\mathbf{P}} = \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{M}^{-1}\mathbf{P}}\,, \tag{1.13}$$

where $\mathbf{M} \in \mathrm{GL}_n(\mathbb{Z})$ is the abelianization of ϕ, and

$$(\Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{0}})^{-1} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{0}} = \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q}^{-1},\mathbf{0}} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{P}\mathbf{Q}} = \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}\mathbf{Q}}\,. \tag{1.14}$$

In particular, $\mathcal{M}_{n\times m}(\mathbb{Z})$ is a normal subgroup of $\operatorname{Aut}(G)$. But $\operatorname{Aut}(\mathbb{F}_n)$, $\mathrm{GL}_m(\mathbb{Z})$ and $\mathcal{M}_{n\times m}(\mathbb{Z})$ altogether generated the whole $\operatorname{Aut}(G)$, as can be seen with the equality

$$\Psi_{\phi,\mathbf{Q},\mathbf{P}} = \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{I}_m,\mathbf{P}\mathbf{Q}^{-1}} \cdot \Psi_{\mathrm{id}_{\mathbb{F}_n},\mathbf{Q},\mathbf{0}} \cdot \Psi_{\phi,\mathbf{I}_m,\mathbf{0}}\,. \tag{1.15}$$

Thus, $\mathrm{Aut}(G)$ is isomorphic to the semidirect product $\mathcal{M}_{n\times m}(\mathbb{Z}) \rtimes (\mathrm{Aut}(\mathbb{F}_n) \times \mathrm{GL}_m(\mathbb{Z}))$, with the action of $\mathrm{Aut}(\mathbb{F}_n) \times \mathrm{GL}_m(\mathbb{Z})$ on $\mathcal{M}_{n\times m}(\mathbb{Z})$ given by equations (1.13) and (1.14).

Since $\mathcal{M}_{n\times m}(\mathbb{Z})$, $\mathrm{Aut}(\mathbb{F}_n)$, and $\mathrm{GL}_m(\mathbb{Z})$ are finitely presented, $\mathrm{Aut}(G)$ is finitely presented too. This completes the proof. □

Remark 1.3.9. Finite presentability of $\mathrm{Aut}(G)$ was previously known as a particular case of a more general result: in [Lau95], M. Laurence gave a finite family of generators for the group of automorphisms of any finitely generated partially commutative group, in terms of the underlying graph. It turns out that, when particularizing this to free-abelian times free groups, Laurence's generating set for $\mathrm{Aut}(G)$ is essentially the same as the one obtained here, after deleting some obvious redundancy.

Later, in [Day09], M. Day builds a kind of peak reduction technique for such groups, from which he deduces finite presentation for its group of automorphisms. However, our Theorem 1.3.8 provides the explicit structure of the automorphism group of a free-abelian times free group.

Algorithmic problems 2

We shall dedicate this chapter to solve several algorithmic problems in $G = \mathbb{Z}^m \times \mathbb{F}_n$. The first approach will be to reduce them to the analogous problems on each part, $\mathbb{Z}^m$ and $\mathbb{F}_n$, and then apply the vast existing literature for free-abelian and free groups. In some cases, the solutions for the free-abelian and free parts will naturally build up a solution for G, while in some others the interaction between both parts will be more intricate; it depends on how complicated the relation between the two parts becomes.

From the algorithmic point of view, the statement "let G be a group" is not sufficiently precise. The algorithmic behavior of G may depend on how it is given to us. For free-abelian times free groups, we will always assume that they are finitely generated and given to us with the standard presentation (1.2). We will also assume that elements, subgroups, homomorphisms and any other objects associated with the group are given to us in terms of this presentation.

The chapter is organized as follows: we start recalling the folklore solutions to the three classical Dehn problems for free-abelian times free groups, to then move to the subgroup membership problem in Section 2.1, which is an easy consequence of the computability of basis in this family. In the next two sections we study some other more interesting algorithmic problems: the finite index subgroup problem in Section 2.2, and the subgroup and coset intersection problems in Section 2.3. In the second half of the chapter we consider algorithmic problems involving endomorphisms. Namely we study the fixed subgroup of an endomorphism in Section 2.4, the Whitehead problems in Section 2.5, and finally in Section 2.6 the twisted-conjugacy problem and its connection with orbit decidability.

From the existence and computability of normal forms for its elements, one can easily deduce the solvability of the word and conjugacy problems for $\mathbb{Z}^m \times \mathbb{F}_n$. The third of Dehn's problems is also easy within our family of groups.

Proposition 2.0.1. *Let* $G = \mathbb{Z}^m \times \mathbb{F}_n$. *Then*

(i) *the word problem for* G *is solvable,*

(ii) *the conjugacy problem for* G *is solvable,*

(iii) *the isomorphism problem is solvable within the family of finitely generated free-abelian times free groups.*

Proof. The solvability of the word and conjugacy problems is straightforward using the existence of normal forms for the group elements.

For the isomorphism problem, let $\langle X \mid R \rangle$ and $\langle Y \mid S \rangle$ be two arbitrary finite presentations of free-abelian times free groups G and G' (i.e., we are given two arbitrary finite presentations plus the information that both groups are free-abelian times free). So, both G and G' admit presentations of the form (1.2), say $P_{n,m}$ and $P_{n',m'}$, for some integers $m, n, m', n' \geqslant 0$, $n, n' \neq 1$ (unknown at the beginning).

It is well known that two finite presentations present the same group if and only if they are connected by a finite sequence of Tietze transformations (see [LS01]); so, there exist finite sequences of Tietze transformations, one from $\langle X \mid R \rangle$ to $P_{n,m}$, and another from $\langle Y \mid S \rangle$ to $P_{n',m'}$ (again, unknown at the beginning). Let us start two diagonal procedures exploring, respectively, the tree of all possible Tietze transformations successively applicable to $\langle X \mid R \rangle$ and $\langle Y \mid S \rangle$.

From the discussion above, it is clear that both procedures will necessarily reach presentations of the desired form in finite time. Once the parameters m, n, m', n' are known, we apply Lemma 1.1.3 and conclude that $\langle X \mid R \rangle$ and $\langle Y \mid S \rangle$ are isomorphic if and only if $n = n'$ and $m = m'$. □

2.1 Subgroup membership problem

As a first application of Proposition 1.2.6, free-abelian times free groups have solvable *membership problem*. Let us first state the problem for an arbitrary group G.

(Subgroup) membership problem, $\mathrm{MP}(G)$**.** *Given elements* $g, h_1, \ldots, h_p \in G$*, decide whether* $g \in H = \langle h_1, \ldots, h_p \rangle$ *and, in affirmative case, compute an expression of* g *as a word in the* h_i*'s.*

Proposition 2.1.1. *The membership problem for* $G = \mathbb{Z}^m \times \mathbb{F}_n$ *is solvable.*

Proof. Write $g = t^{\mathbf{a}} w$. We start by computing a basis for H following Proposition 1.2.6, say $\{t^{\mathbf{b}_1}, \ldots, t^{\mathbf{b}_{m'}}, t^{\mathbf{a}_1} u_1, \ldots, t^{\mathbf{a}_{n'}} u_{n'}\}$. Now, check whether $g\pi_{\mathbb{F}} = w \in H\pi_{\mathbb{F}} = \langle u_1, \ldots, u_{n'} \rangle$ (MP is well known to be solvable for finitely generated free groups). If the answer is negative then $g \notin H$ and we are done. Otherwise, a standard algorithm for membership in free groups gives us the (unique) expression of w as a word on the u_j's, say $w = \omega(u_1, \ldots, u_{n'})$. Finally, compute $\omega(t^{\mathbf{a}_1} u_1, \ldots, t^{\mathbf{a}_{n'}} u_{n'}) = t^{\mathbf{c}} w \in H$. It is clear that $t^{\mathbf{a}} w \in H$ if and only if $t^{\mathbf{a}-\mathbf{c}} = (t^{\mathbf{a}} w)(t^{\mathbf{c}} w)^{-1} \in H$ that is, if and only if $\mathbf{a} - \mathbf{c} \in \langle \mathbf{b}_1, \ldots, \mathbf{b}_{m'} \rangle \leqslant \mathbb{Z}^m$. This can be checked by just solving a system of linear equations; and, in the affirmative case, we can easily find an expression for g in terms of $\{t^{\mathbf{b}_1}, \ldots, t^{\mathbf{b}_{m'}}, t^{\mathbf{a}_1} u_1, \ldots, t^{\mathbf{a}_{n'}} u_{n'}\}$, like at the end of the previous proof. Finally, it only remains to convert this into an expression of g in $\{h_1, \ldots, h_p\}$ using the expressions we already have for the basis elements in terms of the h_i's. $\square$

Since any element of an arbitrary free-abelian times free group has finite support, we immediately deduce the following result.

Corollary 2.1.2. *The membership problem for arbitrary free-abelian times free groups is solvable.* $\square$

2.2 Finite index problem

In this section, the goal is to find an algorithm solving the finite index problem for a free-abelian times free group.

Finite index problem, $\mathsf{FIP}(G)$**.** *Given a finite list* $w_1, \ldots, w_s$ *of elements in* G*, decide whether the subgroup* $H = \langle w_1, \ldots, w_s \rangle$ *is of finite index in* G*; and, if so, compute the index and a system of right (or left) coset representatives for* H.

To start, recall that this same algorithmic problem is well known to be solvable both for free-abelian and for free groups.

Proposition 2.2.1. *A subgroup* $L \leqslant \mathbb{Z}^m$ *has finite index in* $\mathbb{Z}^m$ *if and only if it has (maximum) rank* m*. Moreover, given* L *(by a finite set of generators), we can algorithmically decide whether* L *is of finite index in* $\mathbb{Z}^m$*, and*

(i) *effectively compute a transversal (and so the index* $[\mathbb{Z}^m : L]$*) if the index is finite.*

(ii) *recursively enumerate a transversal for* $\mathbb{Z}^m / L$ *if the index is infinite.*

Proof. Let $\mathbf{a_1}, \ldots, \mathbf{a_s} \in \mathbb{Z}^m$ be a generating set for L. Consider the $s \times m$ integral matrix $\mathbf{A}$ whose rows are the $\mathbf{a_i}$'s, and compute its Smith normal form, i.e.

$$\mathbf{PA} = \mathbf{diag}(d_1, d_2, \ldots, d_r, 0, \ldots, 0)\, \mathbf{Q}, \qquad (2.1)$$

where $\mathbf{P} \in \mathrm{GL}_s(\mathbb{Z})$, $\mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})$, $d_1, \ldots, d_r$ are nonzero positive integers each dividing the following one ($d_1 \mid d_2 \mid \cdots \mid d_r \neq 0$), and $r \leqslant \min\{s, m\}$ is the rank of the diagonal matrix, and thus of $\mathbf{A}$ (standard algorithms are known to compute all these from $\mathbf{A}$, see [Art10] for details).

Since rows of $\mathbf{A}$ generate L, it is clear from (2.1) that L is also generated by the rows of $\mathbf{PA}$, i.e. the image under the automorphism $\mathbf{Q} \colon \mathbb{Z}^m \to \mathbb{Z}^m$, $\mathbf{v} \mapsto \mathbf{vQ}$ of the subgroup L' generated by the vectors $(d_1, 0, \ldots, 0), \ldots, (0, \ldots, 0, d_r)$.

Now, it is clear that

$$[d_1] \times \cdots \times [d_r] \times \mathbb{Z}^{m-r} \tag{2.2}$$

is a transversal for $\mathbb{Z}^m/L'$, and hence its image under $\mathbf{Q}$ is a transversal for $\mathbb{Z}^m/L$. Moreover, note that the index of L (and of L') in $\mathbb{Z}^m$ is:

$$[\mathbb{Z}^m : L] = d_1 \cdots d_r \cdot \# \mathbb{Z}^{m-r}, \tag{2.3}$$

which is finite if and only if $r = m$, i.e., if and only if L (and so L') has maximal rank in $\mathbb{Z}^m$.

Finally recall that the finite index condition (i.e., $r = m$) is algorithmically decidable by just looking at the number of nonzero diagonal elements in the Smith normal form in (2.1); and the obtained transversal for $\mathbb{Z}^m/L$ recursively enumerable just tracking forward the elements in (2.2) (diagonally if $m - r \geqslant 2$) and multiplying them by $\mathbf{Q}$. □

Corollary 2.2.2. *The finite index problem for $\mathbb{Z}^m$ is solvable.* □

On the other hand, the subgroup $H = \langle w_1, \ldots, w_s \rangle \leqslant \mathbb{F}_n$ has finite index if and only if every vertex in the core of the Schreier graph of H, denoted $\mathcal{S}(H)$, is complete (i.e., has degree $2n$); this is algorithmically checkable by means of fast algorithms. And, in this case, the labels of paths in a chosen maximal tree T from the basepoint to each vertex (resp. from each vertex to the basepoint) give a set of left (resp. right) coset representatives for H, whose index in $\mathbb{F}_n$ is then the number of vertices of $\mathcal{S}(H)$. For details, see [Sta83] for the classical reference or [KM02] for a more modern and combinatorial approach.

Hence, $\mathsf{FIP}(\mathbb{Z}^m)$ and $\mathsf{FIP}(\mathbb{F}_n)$ are solvable. In order to build an algorithm to solve the same problem in $\mathbb{Z}^m \times \mathbb{F}_n$, we shall need some well-known basic facts about indices of subgroups, recalled in the following lemmas. For a subgroup $H \leqslant G$ of an arbitrary group G, we will write $H \leqslant_{fi} G$ to denote $[G : H] < \infty$.

Lemma 2.2.3. *Let G and G' be arbitrary groups, $\rho\colon G \twoheadrightarrow G'$ an epimorphism between them, and let $H \leqslant G$ and $H' \leqslant G'$ be arbitrary subgroups. Then,*

(i) $[G : H] \geqslant [G' : H\rho]$; *in particular, if $H \leqslant_{fi} G$ then $H\rho \leqslant_{fi} G'$.*

(ii) $[G : H'\rho^{\leftarrow}] = [G' : H']$; *in particular, $H' \leqslant_{fi} G'$ if and only if $H'\rho^{\leftarrow} \leqslant_{fi} G$.*

Proof. (i) It is enough to see that the map $\widetilde{\rho}\colon G/H \to G'/H'$, $xH \mapsto (xH)\rho$ is well defined and surjective. It is well defined since the image of any subset is unique, and $(xH)\rho = x\rho\, H\rho \in G'/H\rho$. Surjectivity is immediate from that of ρ.

(ii) It is enough to see that the map $\rho^{\leftarrow}\colon G'/H' \to G/(H'\rho^{\leftarrow})\,,\; x'H' \mapsto (x'H')\rho^{\leftarrow}$ is well defined and bijective. Again, it is clear that the complete preimage of a set is unique. Moreover, from the surjectivity of ρ:

$$\begin{aligned} y \in (x'H')\rho^{\leftarrow} \;\Leftrightarrow\; y\rho \in x\rho\, H' \;\Leftrightarrow\; (x\rho)^{-1} y\rho \in H' \\ \Leftrightarrow\; (x^{-1}y)\rho \in H' \;\Leftrightarrow\; x^{-1}y \in H'\rho^{\leftarrow} \;\Leftrightarrow\; y \in x\,(H'\rho^{\leftarrow})\,, \end{aligned}$$

for certain preimage $x \in G$ of x'. Thus, the image $(x'H')\rho^{\leftarrow} = x\,(H'\rho^{\leftarrow}) \in G/H'\rho^{\leftarrow}$ and the map is well defined. This same fact (reversed) also proves surjectivity ($\forall x \in G,\; x(H'\rho^{\leftarrow}) = (x\rho\, H')\rho^{\leftarrow}$; and injectivity:

$$(x'\,H')\rho^{\leftarrow} = (y'\,H')\rho^{\leftarrow} \;\Rightarrow\; (x\rho\, H')\rho^{\leftarrow} = (y\rho\, H') \;\Rightarrow\; x\,(H'\rho^{\leftarrow}) = y\,(H'\rho^{\leftarrow})\,.$$

This completes the proof. □

Lemma 2.2.4. *Let* H, H' *be subgroups of an arbitrary group, and* $K \leqslant H$. *Then:*

$$[H:K] \;\geqslant\; [H\cap H' : K\cap H']\,, \tag{2.4}$$

in particular, $K \leqslant_{fi} H \;\Rightarrow\; K\cap H' \leqslant_{fi} H\cap H'$.

Proof. It is immediate to see that the map $H\cap H'/K\cap H' \to H/K\,,\; x(K\cap H') \mapsto xK$ is well-defined and injective. □

Lemma 2.2.5. *Let* G *and* G' *be arbitrary groups, and let* $H \leqslant G\times G'$ *be a subgroup of their direct product. Then*

$$[G\times G' : H] \;\leqslant\; [G : H\cap G]\cdot[G' : H\cap G'], \tag{2.5}$$

and

$$H \leqslant_{fi} G\times G' \;\Leftrightarrow\; H\cap G \leqslant_{fi} G \text{ and } H\cap G' \leqslant_{fi} G'. \tag{2.6}$$

Proof. It is easy to prove that the map

$$\begin{array}{rcl} G/(H\cap G) \;\times\; G'/(H\cap G') & \to & (G\times G')/H \\ \big(g\,(H\cap G)\,,\, g'\,(H\cap G')\big) & \mapsto & gg'\,H \end{array} \tag{2.7}$$

is well defined and surjective. Namely, let $g_1\,(H\cap G) = g_2\,(H\cap G)$, and $g_1'\,(H\cap G') = g_2'\,(H\cap G')$; i.e., let $g_1(g_2)^{-1} \in H\cap G$, and $g_1'(g_2')^{-1} \in H\cap G'$. Then, $g_1(g_2)^{-1}g_1'(g_2')^{-1} \in H$, and since elements in G commute with elements in G', we have that $g_1g_1'(g_2'g_2)^{-1} \in H$, that is $g_1g_1'\,H = g_2g_2'\,H$, and the map in (2.7) is well defined. Since its surjectivity is obvious, we have already proved the inequality (2.5), and hence the left implication in (2.6).

The converse implication is a particular case of Lemma 2.2.4. So, the proof is complete. □

Remark 2.2.6. Let $G = \mathbb{Z}^m \times \mathbb{F}_n$, and let H be a subgroup of G. If H has finite index in G then, applying Lemma 2.2.3.(i) to the canonical projections $\pi_{\mathbb{A}} \colon G \twoheadrightarrow \mathbb{Z}^m$ and $\pi_{\mathbb{F}} \colon G \twoheadrightarrow \mathbb{F}_n$, we have that both indices $[\mathbb{Z}^m : H\pi_{\mathbb{A}}]$, and $[\mathbb{F}_n : H\pi_{\mathbb{F}}]$ must also be finite. Since we can effectively compute generators for $H\pi_{\mathbb{F}}$ and for $H\pi_{\mathbb{A}}$, and we can decide whether $H\pi_{\mathbb{A}} \leqslant_{fi} \mathbb{Z}^m$ and $H\pi_{\mathbb{F}} \leqslant_{fi} \mathbb{F}_n$ hold, we have two effectively checkable necessary conditions for H to be of finite index in G: if either $[\mathbb{Z}^m : H\pi_{\mathbb{A}}]$ or $[\mathbb{F}_n : H\pi_{\mathbb{F}}]$ is infinite, then so is $[G : H]$.

Nevertheless, these two necessary conditions together are not sufficient to ensure finiteness of $[G : H]$, as the following easy example shows.

Example 2.2.7. Let $H = \langle sa, tb \rangle$, a subgroup of $G = \mathbb{Z}^2 \times \mathbb{F}_2 = \langle s, t \mid [s, t] \rangle \times \langle a, b \mid \rangle$. It is clear that $H\pi_{\mathbb{A}} = \mathbb{Z}^2$ and $H\pi_{\mathbb{F}} = \mathbb{F}_2$ (so, both indices are 1), but the index $[\mathbb{Z}^2 \times \mathbb{F}_2 : H]$ is infinite because no power of a belongs to H.

Note that $H \cap \mathbb{Z}^m \leqslant H\pi_{\mathbb{A}} \leqslant \mathbb{Z}^m$ and $H \cap \mathbb{F}_n \leqslant H\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$. So, according to Lemma 2.2.5, the conditions really necessary, and sufficient, for H to be of finite index in G (detailed below) are stronger than $H\pi_{\mathbb{A}} \leqslant_{fi} \mathbb{Z}^m$ and $H\pi_{\mathbb{F}} \leqslant_{fi} \mathbb{F}_n$ respectively, and none of them satisfied in the example.

Lemma 2.2.8. *Let H be a subgroup of $G = \mathbb{Z}^m \times \mathbb{F}_n$. Then,*

$$H \leqslant_{fi} G \;\Leftrightarrow\; \begin{cases} H \cap \mathbb{Z}^m \leqslant_{fi} \mathbb{Z}^m, \\ H \cap \mathbb{F}_n \leqslant_{fi} H\pi_{\mathbb{F}}, \text{ and } H\pi_{\mathbb{F}} \leqslant_{fi} \mathbb{F}_n\,. \end{cases} \tag{2.8}$$

□

Theorem 2.2.9. *The finite index problem for $\mathbb{Z}^m \times \mathbb{F}_n$ is solvable.*

Proof. From the given generators for H, we start by computing a basis of H (see Proposition 1.2.6), say

$$\{t^{\mathbf{b}_1}, \ldots, t^{\mathbf{b}_{m'}}, t^{\mathbf{a}_1}u_1, \ldots, t^{\mathbf{a}_{n'}}u_{n'}\},$$

where $0 \leqslant m' \leqslant m$, $0 \leqslant n' \leqslant p$, $\{\mathbf{b}_1, \ldots, \mathbf{b}_{m'}\}$ is a free-abelian basis for $L = H \cap \mathbb{Z}^m \simeq \mathbb{Z}^{m'}$, $\mathbf{a}_1, \ldots, \mathbf{a}_{n'} \in \mathbb{Z}^m$, and $\{u_1, \ldots, u_{n'}\}$ is a free basis for $H\pi_{\mathbb{F}} \simeq \mathbb{F}_{n'}$. As above, let us write $\mathbf{A}$ for the $n' \times m$ integral matrix whose rows are $\mathbf{a}_i \in \mathbb{Z}^m$, $i \in [n']$.

Since $L = \langle \mathbf{b}_1, \ldots, \mathbf{b}_{m'} \rangle \simeq H \cap \mathbb{Z}^m$ (with the natural isomorphism $\mathbf{b} \mapsto t^{\mathbf{b}}$, changing the notation from additive to multiplicative), the first necessary condition in (2.8) is $\operatorname{rk}(L) = m$, i.e., $m' = m$. If this is not the case, then $[G : H] = \infty$ and we are done. So, let us assume $m' = m$ and compute a set of (right) coset representatives for L in $\mathbb{Z}^m$, say $\mathbb{Z}^m = \mathbf{c}_1 L \sqcup \cdots \sqcup \mathbf{c}_r L$.

Next, check whether $H\pi_{\mathbb{F}} = \langle u_1, \ldots, u_{n'} \rangle$ has finite index in $\mathbb{F}_n$ (by computing the core of the Schreier graph of $H\pi_{\mathbb{F}}$, and checking whether is it complete or not, see Proposition 5.5.6). If this is not the case, then $[G : H] = \infty$ and we are done as well. So, let us assume $H\pi_{\mathbb{F}} \leqslant_{\mathrm{fi}} \mathbb{F}_n$, and using $\mathsf{FIP}(\mathbb{F}_n)$ (Theorem 5.5.7), compute a set of right coset representatives for $H\pi_{\mathbb{F}}$ in $\mathbb{F}_n$, say $\mathbb{F}_n = v_1(H\pi_{\mathbb{F}}) \sqcup \cdots \sqcup v_s(H\pi_{\mathbb{F}})$.

According to Lemma 2.2.8, it only remains to check whether the subgroup $H \cap \mathbb{F}_n$ has finite or infinite index in $H\pi_{\mathbb{F}}$. Call $\rho \colon \mathbb{F}_{n'} \twoheadrightarrow \mathbb{Z}^{n'}$ the abstract abelianization map for the free group of rank n' (with free basis $\{u_1, \ldots, u_{n'}\}$), and $A \colon \mathbb{Z}^{n'} \to \mathbb{Z}^m$ the linear mapping $\mathbf{v} \mapsto \mathbf{v}\mathbf{A}$ corresponding to right multiplication by the matrix $\mathbf{A}$. Note that

$$H \cap \mathbb{F}_n = \{w \in \mathbb{F}_n : \mathbf{0} \in \mathcal{C}_H(w)\} = \{w \in \mathbb{F}_n : \boldsymbol{\omega}\mathbf{A} \in L\} \leqslant H\pi_{\mathbb{F}},$$

where $\boldsymbol{\omega} = \omega\rho$ is the abelianization of the word ω which expresses w in the free basis $\{u_1, \ldots, u_{n'}\}$ of $H\pi_{\mathbb{F}}$; i.e., $\omega(u_1, \ldots, u_{n'}) = w \in \mathbb{F}_n$, see Corollary 1.2.9. Thus, $H \cap \mathbb{F}_n$ is, in terms of the free basis $\{u_1, \ldots, u_{n'}\}$, the successive full preimage of L, first by the map A and then by the map ρ, namely $H \cap \mathbb{F}_n \simeq (L)A^{-1}\rho^{-1}$ (see the diagram in Equation (2.9)).

$$\begin{array}{ccccccccc}
\mathbb{F}_n \geqslant & H\pi_{\mathbb{F}} & \simeq & \mathbb{F}_{n'} & \xrightarrow{\ \rho\ } \!\!\!\!\!\rightarrow & \mathbb{Z}^{n'} & \xrightarrow{\ A\ } & \mathbb{Z}^m \\
& \triangledown\!\!/ & & \triangledown\!\!/ & & \triangledown\!\!/ & & \triangledown\!\!/ \\
& H \cap \mathbb{F}_n & \simeq & (L)A^{\leftarrow}\rho^{\leftarrow} & \longleftarrow & (L)A^{\leftarrow} & \longleftarrow & L
\end{array} \tag{2.9}$$

Fig. 2.1: Finite index problem diagram for FATF groups

Hence, using Lemma 2.2.3.(ii), $[H\pi_{\mathbb{F}} : H \cap \mathbb{F}_n] = [\mathbb{F}_{n'} : (L)A^{\leftarrow}\rho^{\leftarrow}]$ is finite if and only if $[\mathbb{Z}^{n'} : (L)A^{\leftarrow}]$ is finite (in fact, both indices coincide). And this happens if and only if $\mathrm{rk}((L)A^{\leftarrow}) = n'$. Since $\mathrm{rk}((L)A^{\leftarrow}) = \mathrm{rk}((L \cap \mathrm{im}(A))A^{\leftarrow}) = \mathrm{rk}(L \cap \mathrm{im}(A)) + \mathrm{rk}(\ker(A))$, we can immediately check whether this rank equals n', or not. If this is not the case, then $[H\pi_{\mathbb{F}} : H \cap \mathbb{F}_n] = [\mathbb{F}_{n'} : (L)A^{\leftarrow}\rho^{\leftarrow}] = [\mathbb{Z}^{n'} : (L)A^{\leftarrow}] = \infty$ and we are done. Otherwise, $(L)A^{\leftarrow} \leqslant_{\mathrm{fi}} \mathbb{Z}^{n'}$ and so, $H \cap \mathbb{F}_n \leqslant_{\mathrm{fi}} H\pi_{\mathbb{F}}$ and $H \leqslant_{\mathrm{fi}} G$. This concludes the decision part of the problem.
For the search part, suppose $H \leqslant_{\mathrm{fi}} G$ and let us explain how to compute a set of right coset representatives for H in G (and so, the actual value of the index $[G : H]$). Having followed the algorithm described above, we have $\mathbb{Z}^m = \mathbf{c}_1 L \sqcup \cdots \sqcup \mathbf{c}_r L$, and $\mathbb{F}_n = v_1(H\pi_{\mathbb{F}}) \sqcup \cdots \sqcup v_s(H\pi_{\mathbb{F}})$. Furthermore, from the analysis in the previous paragraph, we can also compute a set of (right) coset representatives $\{\mathbf{d}_1, \ldots, \mathbf{d}_t\}$ for $(L)A^{-1}$ in $\mathbb{Z}^{n'}$, which can be biunivocally converted (taking respective ρ preimages, see Lemma 2.2.3.(ii)) into a set of right coset representatives for $H \cap \mathbb{F}_n$ in $H\pi_{\mathbb{F}}$, say $H\pi_{\mathbb{F}} = w_1(H \cap \mathbb{F}_n) \sqcup \cdots \sqcup w_t(H \cap \mathbb{F}_n)$.
Hence,

$$\mathbb{F}_n = \bigsqcup\nolimits_{j \in [s]} \bigsqcup\nolimits_{k \in [t]} v_j w_k (H \cap \mathbb{F}_n), \quad \text{and} \quad [\mathbb{F}_n : H \cap \mathbb{F}_n] = st.$$

Combining this with $\mathbb{Z}^m = \bigsqcup_{i\in[r]} t^{c_i}(H\cap\mathbb{Z}^m)$, and using the map in the proof of Lemma 2.2.5, we get

$$G = \mathbb{Z}^m \times \mathbb{F}_n = \bigcup_{i\in[r]} \bigcup_{j\in[s]} \bigcup_{k\in[t]} t^{c_i} v_j w_k H,$$

where some of the cosets $t^{c_i} v_j w_k H$ may coincide.
It only remains to perform a cleaning process in the family of rst elements

$$\{ t^{c_i} v_j w_k : i \in [r], j \in [s], k \in [t] \}$$

in order to eliminate any possible duplication in the representatives of the right cosets of H. This can be easily done by successive application of the membership problem for H, see Corollary 2.1.2. After this cleaning process, we get a genuine set of right coset representatives for H in G, and the actual value of $[G : H]$, which is at most rst. (Note that in general this cleaning process cannot be avoided as the following example shows.)
Finally, inverting all of them we will get a set of left coset representatives for H in G. □

Example 2.2.10. Let $G = \mathbb{Z}^2 \times \mathbb{F}_2 = \langle s, t \mid [s, t] \rangle \times \langle a, b \mid \rangle$ and consider the (normal) subgroups $H = \langle s, t^2, a, b^2, bab\rangle$, and $H' = \langle s, t^2, a, b^2, bab, tb\rangle = \langle s, t^2, a, tb\rangle$ of G (with bases $\{s, t^2, a, b^2, bab\}$ and $\{s, t^2, a, tb\}$, respectively). We have $H\cap\mathbb{Z}^2 = H'\cap\mathbb{Z}^2 = \langle s, t^2\rangle \leqslant_2 \mathbb{Z}^2$, and $H\cap\mathbb{F}_2 = H'\cap\mathbb{F}_2 = \langle a, b^2, bab\rangle \leqslant_2 \mathbb{F}_2$, but

$$[\mathbb{Z}^2 \times \mathbb{F}_2 : H] = 4 = [\mathbb{Z}^2 : H\cap\mathbb{Z}^2]\cdot[\mathbb{F}_2 : H\cap\mathbb{F}_2],$$

while

$$[\mathbb{Z}^2 \times \mathbb{F}_2 : H'] = 2 < 4 = [\mathbb{Z}^2 : H'\cap\mathbb{Z}^2]\cdot[\mathbb{F}_2 : H'\cap\mathbb{F}_2],$$

with (right) coset representatives $\{1, b, t, tb\}$ and $\{1, t\}$, respectively. This shows that both the equality and the strict inequality can occur in Lemma 2.2.5.

2.3 Intersection problems and Howson's property

A group G is said to have the *Howson property* if the intersection of every pair (and hence every finite family) of finitely generated subgroups $H, H' \leqslant_{fg} G$ is again finitely generated, $H\cap H' \leqslant_{fg} G$.

It is well known that $\mathbb{Z}^m$ satisfies Howson's property, since every subgroup is free-abelian of rank less than or equal to m (and so, finite). Moreover, $\mathsf{SIP}(\mathbb{Z}^m)$ and $\mathsf{CIP}(\mathbb{Z}^m)$ just reduce to solving standard systems of linear equations.

The case of free groups is more interesting. Howson himself established in 1954 that $\mathbb{F}_n$ also satisfies the Howson property, see [How54]. Since then, there have been several improvements of this result in the literature, both about lowering the upper bounds for the rank of the intersection, and about simplifying the arguments used.

The modern point of view is based on the pull-back technique for graphs: One can algorithmically represent subgroups of $\mathbb{F}_n$ by the core of their Schreier graphs (a.k.a. Stallings automata); then the graph corresponding to $H \cap H'$ is the pull-back of the graphs corresponding to H and H', easily constructible from them when they are finitely generated. This not only confirms Howson's property for $\mathbb{F}_n$ (namely, the pull-back of finite graphs is finite) but, more importantly, it provides the algorithmic aspect into the topic by solving $\mathsf{SIP}(\mathbb{F}_n)$. And, more generally, an easy variation of these arguments using pullbacks also solves $\mathsf{CIP}(\mathbb{F}_n)$, see Proposition 6.1 in [BMV10]).

As a generalization of Howson's result, B. Baumslag established in [Bau66] the conservation of Howson's property under free products (i.e., if G_1 and G_2 satisfy Howson's property then so does $G_1 * G_2$). The same statement fails dramatically if we replace the free product by a direct product. And one can find an extremely simple counterexample for this, in the family of free-abelian times free groups; the following observation is folklore (it appears in [BK98] attributed to Moldavanski, and as the solution to exercise 23.8(3) in [Bog08]).

Lemma 2.3.1. *The group $\mathbb{Z}^m \times \mathbb{F}_n$, for $m \geqslant 1$ and $n \geqslant 2$, does not satisfy Howson's property.*

Proof. The following argument is described as a solution to exercise 23.8(3) in [Bog08] (see also [DV13]). Indeed, if we write $\mathbb{Z} \times \mathbb{F}_2 = \langle t \mid - \rangle \times \langle a, b \mid - \rangle$, then the subgroups

$$H = \langle a, b \rangle = \mathbb{F}_2 \leqslant \mathbb{Z} \times \mathbb{F}_2 \text{, and}$$
$$K = \langle ta, b \rangle = \{w(ta, b) \mid w \in \mathbb{F}_2\} = \{t^{|w|_a} w(a, b) \mid w \in \mathbb{F}_2\} \leqslant \mathbb{Z} \times \mathbb{F}_2$$

are both finitely generated, but its intersection

$$H \cap K = \{t^0 w(a, b) \mid w \in \mathbb{F}_2, |w|_a = 0\} = \langle\!\langle b \rangle\!\rangle_{\mathbb{F}_2} = \langle a^{-k} b a^k, \, k \in \mathbb{Z} \rangle$$

is the normal closure of b in $\mathbb{F}_2$, which is infinitely generated, as you can immediately see from its Stallings automaton (see Corollary 5.4.33)

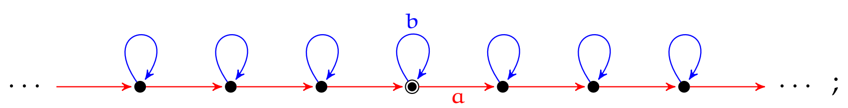

Fig. 2.2: Stallings automaton of $\langle\!\langle b \rangle\!\rangle$ in $\mathbb{F}_{\{a,b\}}$

or using this alternative argument: Suppose that $H \cap K$ is finitely generated, then there exist an $m \in \mathbb{N}$ such that $a^{m+1}ba^{-(m+1)} \in \langle a^{-k}ba^{k}, k \in [-m, m] \rangle$, and thus a^{m+1} equals the reduced form of some prefix of $w(a^{m}ba^{-m}, \ldots, b, \ldots, a^{-m}ba^{m})$, for some word w. However, the sum of the exponents of a in any such prefix must be in $[-m, m]$, which is a contradiction.

Note that both H and K are free groups of rank two whose intersection is infinitely generated. This fact, far from violating the Howson property, means that they are not simultaneously contained in any free subgroup of $\mathbb{Z} \times \mathbb{F}_2$. □

We remark that the subgroups H and H' in the previous counterexample are both isomorphic to $\mathbb{F}_2$. So, interestingly, the above is a situation where two free groups of rank 2 have a non-finitely generated (of course, free) intersection. This does not contradict the Howson property for free groups, but rather indicates that one cannot embed H and H' simultaneously into a free subgroup of $\mathbb{Z} \times \mathbb{F}_2$.

In this setting, it makes sense to consider the following two related algorithmic problems (stated for an arbitrary group G).

Subgroup intersection problem, $\mathsf{SIP}(G)$**.** *Given finitely generated subgroups H and H' of G (by finite sets of generators), decide whether the intersection $H \cap H'$ is finitely generated and, if so, compute a generating set for $H \cap H'$.*

Coset intersection problem, $\mathsf{CIP}(G)$**.** *Given finitely generated subgroups H and H' of G (by finite sets of generators), and elements $g, g' \in G$, decide whether the right cosets gH and $g'H'$ intersect trivially or not; and in the negative case (i.e., when $gH \cap g'H' = g''(H \cap H')$), compute such a $g'' \in G$.*

In the present section, we shall solve $\mathsf{SIP}(\mathbb{Z}^m \times \mathbb{F}_n)$ and $\mathsf{CIP}(\mathbb{Z}^m \times \mathbb{F}_n)$. The key point is Corollary 1.2.4: the intersection $H \cap H'$ is finitely generated if and only if its projection $(H \cap H')\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$ is finitely generated. Note that the group $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ is always finitely generated (by Howson's property of $\mathbb{F}_n$), but the inclusion is not, in general, an equality.

Lemma 2.3.2. *Let H, H' be subgroups of $\mathbb{Z}^m \times \mathbb{F}_n$. Then,*

$$(H \cap H')\pi_{\mathbb{F}} \leqslant H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}},$$

and the inclusion can be strict.

Proof. The inclusion is obvious. For the possible strictness consider the following example. □

Example 2.3.3. Consider in $\mathbb{Z} \times \mathbb{F}_2 = \langle t \mid - \rangle \times \langle a, b \mid - \rangle$ the subgroups below:

$$H = \langle ta^2, bab^{-1}, t^2 \rangle, \quad H' = \langle t^2a^3, ba, t^2 \rangle.$$

Then,

$$\begin{array}{ll} H\pi = \langle a^2, bab^{-1} \rangle, & H'\pi = \langle a^3, ba \rangle, \\ L = 2\mathbb{Z}, & L' = 2\mathbb{Z}, \\ \boldsymbol{\omega} = (3,0), & \boldsymbol{\omega}' = (2,0), \\ \mathbf{A} = \left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right], & \mathbf{A}' = \left[\begin{smallmatrix} 2 \\ 0 \end{smallmatrix}\right]. \end{array}$$

So, it is clear that $a^6 \in H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$.

We claim that $a^6 \notin (H \cap H')\pi_{\mathbb{F}}$. Indeed, $a^6 \in (H \cap H')\pi_{\mathbb{F}}$ if and only if its respective abelian completions in H, and H′ are compatible. Now, from Corollary 1.2.9:

$$\begin{aligned} \mathcal{C}_H(a^6) &= \boldsymbol{\omega}\mathbf{A} + L = (3,0) \left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right] + 2\mathbb{Z} = 1 + 2\mathbb{Z}, \\ \mathcal{C}_{H'}(a^6) &= \boldsymbol{\omega}'\mathbf{A}' + L' = (2,0) \left[\begin{smallmatrix} 2 \\ 0 \end{smallmatrix}\right] + 2\mathbb{Z} = 2\mathbb{Z}. \end{aligned}$$

Thus, $\mathcal{C}_H(a^6) \cap \mathcal{C}_{H'}(a^6) = \varnothing$, and $a^6 \notin (H \cap H')\pi_{\mathbb{F}}$, as claimed.

So, although $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$ is always finitely generated, since the inclusion $(H \cap H')\pi_{\mathbb{F}} \leqslant H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ can be strict, this opens the possibility for $(H \cap H')\pi_{\mathbb{F}}$ (and so for $H \cap H'$) to be non-finitely generated, as is the case in the example in the proof of Lemma 2.3.1.

Let us describe in detail the data involved in $\mathsf{CIP}(G)$ for $G = \mathbb{Z}^m \times \mathbb{F}_n$. By Proposition 1.2.6, we can assume that the initial finitely generated subgroups $H, H' \leqslant G$ are given by respective bases i.e., by two sets of elements

$$\begin{aligned} E &= \{t^{\mathbf{b_1}}, \ldots, t^{\mathbf{b_{m_1}}}, t^{\mathbf{a_1}}u_1, \ldots, t^{\mathbf{a_{n_1}}}u_{n_1}\}, \\ E' &= \{t^{\mathbf{b'_1}}, \ldots, t^{\mathbf{b'_{m_2}}}, t^{\mathbf{a'_1}}u'_1, \ldots, t^{\mathbf{a'_{n_2}}}u'_{n_2}\}, \end{aligned} \tag{2.10}$$

where $\{u_1, \ldots, u_{n_1}\}$ is a free basis of $H\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$, $\{u'_1, \ldots, u'_{n_2}\}$ is a free basis of $H'\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$, $\{t^{\mathbf{b_1}}, \ldots, t^{\mathbf{b_{m_1}}}\}$ is an abelian basis of $H \cap \mathbb{Z}^m$, and $\{t^{\mathbf{b'_1}}, \ldots, t^{\mathbf{b'_{m_2}}}\}$ is an abelian basis of $H' \cap \mathbb{Z}^m$. Consider the subgroups $L = \langle \mathbf{b_1}, \ldots, \mathbf{b_{m_1}} \rangle \leqslant \mathbb{Z}^m$ and $L' = \langle \mathbf{b'_1}, \ldots, \mathbf{b'_{m_2}} \rangle \leqslant \mathbb{Z}^m$, and the matrices

$$\mathbf{A} = \begin{bmatrix} \mathbf{a_1} \\ \vdots \\ \mathbf{a_{n_1}} \end{bmatrix} \in \mathcal{M}_{n_1 \times m}(\mathbb{Z}) \quad \text{and} \quad \mathbf{A}' = \begin{bmatrix} \mathbf{a'_1} \\ \vdots \\ \mathbf{a'_{n_2}} \end{bmatrix} \in \mathcal{M}_{n_2 \times m}(\mathbb{Z}).$$

We are also given two elements $g = t^{\mathbf{a}} u$ and $g' = t^{\mathbf{a}'} u'$ from G, and have to decide whether the intersection $gH \cap g'H'$ is empty or not.

Before describing the algorithm, note that $H\pi_{\mathbb{F}}$ is a free group of rank n_1. Since $\{u_1, \ldots, u_{n_1}\}$ is a free basis of $H\pi_{\mathbb{F}}$, every element $w \in H\pi_{\mathbb{F}}$ can be written in a unique way as a word on the u_i's, say $w = \omega(u_1, \ldots, u_{n_1})$. Abelianizing this word, we get the abelianization map $\rho_1 \colon H\pi_{\mathbb{F}} \twoheadrightarrow \mathbb{Z}^{n_1}$, $w \mapsto \boldsymbol{\omega}$ (not to be confused with the restriction to $H\pi_{\mathbb{F}}$ of the ambient abelianization $\mathbb{F}_n \twoheadrightarrow \mathbb{Z}^n$, which will have no role in this proof). Similarly, we define the morphism $\rho_2 \colon H'\pi_{\mathbb{F}} \twoheadrightarrow \mathbb{Z}^{n_2}$.

With all this data given, note that $gH \cap g'H'$ is empty if and only if its projection to the free component is empty,

$$gH \cap g'H' = \varnothing \;\Leftrightarrow\; (gH \cap g'H')\pi_{\mathbb{F}} = \varnothing;$$

so, it will be enough to study this last projection. And, since this projection contains precisely those elements from $(gH)\pi_{\mathbb{F}} \cap (g'H')\pi_{\mathbb{F}} = (u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})$ having compatible abelian completions in $gH \cap g'H'$, a direct application of Lemma 1.2.7 gives the following result.

Lemma 2.3.4. *With the above notation, the projection $(gH \cap g'H')\pi_{\mathbb{F}}$ consists precisely on those elements $v \in (u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})$ such that*

$$N_v \;=\; (\mathbf{a} + \boldsymbol{\omega}\mathbf{A} + L) \cap \big(\mathbf{a}' + \boldsymbol{\omega}'\mathbf{A}' + L'\big) \;\neq\; \varnothing\,, \tag{2.11}$$

where $\boldsymbol{\omega} = w\rho_1$ and $\boldsymbol{\omega}' = w'\rho_2$ are, respectively, the abelianizations of the abstract words $\omega \in \mathbb{F}_{n_1}$ and $\omega' \in \mathbb{F}_{n_2}$ expressing $w = u^{-1}v \in H\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$ and $w' = u'^{-1}v \in H'\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$ in terms of the free bases $\{u_1, \ldots, u_{n_1}\}$ and $\{u'_1, \ldots, u'_{n_2}\}$ (i.e., $u \cdot \omega(u_1, \ldots, u_{n_1}) = v = u' \cdot \omega'(u'_1, \ldots, u'_{n_2})$). That is,

$$(gH \cap g'H')\pi_{\mathbb{F}} \;=\; \{v \in (u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}}) \,:\, N_v \neq \varnothing\} \subseteq (u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})$$

□

Theorem 2.3.5. *The coset intersection problem for $\mathbb{Z}^m \times \mathbb{F}_n$ is solvable.*

Proof. Let $G = \mathbb{Z}^m \times \mathbb{F}_n$ be a finitely generated free-abelian times free group. Using the solution to $\mathsf{CIP}(\mathbb{F}_n)$, we start by checking whether $(u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})$ is empty or not. In the first case $(gH \cap g'H')\pi_{\mathbb{F}}$, and so $gH \cap g'H'$, will also be empty and we are done. Otherwise, we can compute $v_0 \in \mathbb{F}_n$ such that

$$(u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}}) \;=\; v_0 \cdot (H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}})\,, \tag{2.12}$$

compute words $\omega_0 \in \mathbb{F}_{n_1}$ and $\omega'_0 \in \mathbb{F}_{n_2}$ such that $u \cdot \omega_0(u_1, \ldots, u_{n_1}) = v_0 = u' \cdot \omega'_0(u'_1, \ldots, u'_{n_2})$, and compute a free basis, $\{v_1, \ldots, v_{n_3}\}$, for $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ together with expressions of the ν_i's in terms of the free bases for $H\pi_{\mathbb{F}}$ and $H'\pi_{\mathbb{F}}$, $\nu_i = \nu_i(u_1, \ldots, u_{n_1}) = \nu_i'(u_1', \ldots, u_{n_2}')$, $i \in [n_3]$.

Let $\rho_3 \colon H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} \twoheadrightarrow \mathbb{Z}^{n_3}$ be the corresponding abelianization map. Abelianizing the words ν_i and ν_i', we can compute the rows of the matrices $\mathbf{P}$ and $\mathbf{P}'$ (of sizes $n_3 \times n_1$ and $n_3 \times n_2$, respectively) describing the abelianizations of the inclusion maps $H\pi_{\mathbb{F}} \xleftarrow{\iota} H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} \xrightarrow{\iota'} H'\pi_{\mathbb{F}}$, see the central part of the diagram (2.13) below.

By (2.12), $u^{-1}v_0 \in H\pi_{\mathbb{F}}$ and $u'^{-1}v_0 \in H'\pi_{\mathbb{F}}$. So, left translation by $w_0 = u^{-1}v_0$ is a permutation of $H\pi_{\mathbb{F}}$ (not a homomorphism, unless $w_0 = 1$), say $\lambda_{w_0} \colon H\pi_{\mathbb{F}} \to H\pi_{\mathbb{F}}$, $x \mapsto w_0 x = u^{-1}v_0 x$. Analogously, we have the left translation by $w_0' = u'^{-1}v_0$, say $\lambda_{w_0'} \colon H'\pi_{\mathbb{F}} \to H'\pi_{\mathbb{F}}$, $x \mapsto w_0' x = u'^{-1}v_0 x$. We include these translations in our diagram:

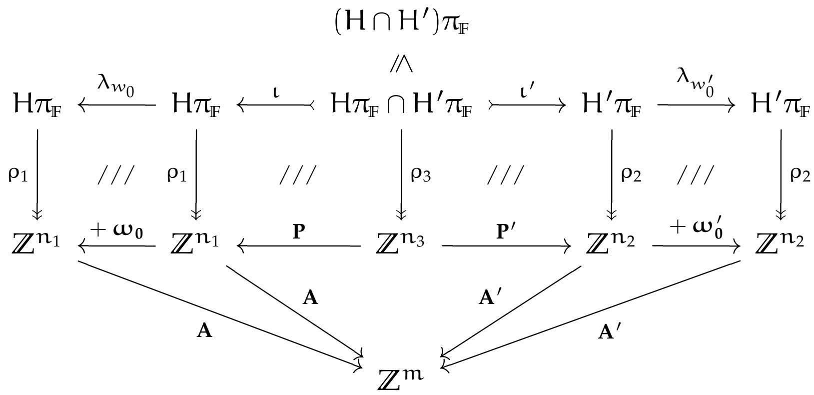

(2.13)

where $\boldsymbol{\omega_0} = w_0\rho_1 \in \mathbb{Z}^{n_1}$ and $\boldsymbol{\omega_0'} = w_0'\rho_2 \in \mathbb{Z}^{n_2}$ are the abelianizations of w_0 and w_0' with respect to the free bases $\{u_1, \ldots, u_{n_1}\}$ and $\{u_1', \ldots, u_{n_2}'\}$, respectively.

Now, for every $v \in (u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})$, using Lemma 2.3.4 and the commutativity of the upper part of the above diagram, we have

$$\begin{aligned} N_v &= \big(\mathbf{a} + (u^{-1}v)\rho_1\mathbf{A} + L\big) \cap \big(\mathbf{a}' + (u'^{-1}v)\rho_2\mathbf{A}' + L'\big) \\ &= \big(\mathbf{a} + (v_0^{-1}v)\iota\lambda_{w_0}\rho_1\mathbf{A} + L\big) \cap \Big(\mathbf{a}' + (v_0^{-1}v)\iota'\lambda_{w_0'}\rho_2\mathbf{A}' + L'\Big) \\ &= \big(\mathbf{a} + (\boldsymbol{\omega_0} + (v_0^{-1}v)\rho_3\mathbf{P})\mathbf{A} + L\big) \cap \big(\mathbf{a}' + (\boldsymbol{\omega_0'} + (v_0^{-1}v)\rho_3\mathbf{P}')\mathbf{A}' + L'\big) \\ &= \big(\mathbf{a} + \boldsymbol{\omega_0}\mathbf{A} + (v_0^{-1}v)\rho_3\mathbf{P}\mathbf{A} + L\big) \cap \big(\mathbf{a}' + \boldsymbol{\omega_0'}\mathbf{A}' + (v_0^{-1}v)\rho_3\mathbf{P}'\mathbf{A}' + L'\big)\,. \end{aligned}$$

With this expression, we can characterize, in a computable way, which elements from $(u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})$ do belong to $(gH \cap g'H')\pi_{\mathbb{F}}$.

Lemma 2.3.6. *With the current notation we have*

$$(gH \cap g'H')\pi_{\mathbb{F}} \;=\; M\rho_3^{\leftarrow}\lambda_{v_0} \;\subseteq\; (u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})\,, \tag{2.14}$$

where $M \subseteq \mathbb{Z}^{n_3}$ *is the preimage by the linear mapping* $\mathbf{PA} - \mathbf{P}'\mathbf{A}' \colon \mathbb{Z}^{n_3} \to \mathbb{Z}^m$ *of the linear variety*

$$N = \mathbf{a}' - \mathbf{a} + \omega_0'\mathbf{A}' - \omega_0\mathbf{A} + (L + L') \subseteq \mathbb{Z}^m. \tag{2.15}$$

Proof. By Lemma 2.3.4, an element $v \in (u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}})$ belongs to $(gH \cap g'H')\pi_{\mathbb{F}}$ if and only if $N_v \neq \varnothing$. That is, if and only if the vector $\mathbf{x} = (v_0^{-1}v)\rho_3 \in \mathbb{Z}^{n_3}$ satisfies that the two varieties $\mathbf{a} + \omega_0\mathbf{A} + \mathbf{xPA} + L$ and $\mathbf{a}' + \omega_0'\mathbf{A}' + \mathbf{xP}'\mathbf{A}' + L'$ do intersect. But this happens if and only if the vector

$$\left(\mathbf{a} + \omega_0\mathbf{A} + \mathbf{xPA}\right) - \left(\mathbf{a}' + \omega_0'\mathbf{A}' + \mathbf{xP}'\mathbf{A}'\right) = \mathbf{a} - \mathbf{a}' + \omega_0\mathbf{A} - \omega_0'\mathbf{A}' + \mathbf{x}(\mathbf{PA} - \mathbf{P}'\mathbf{A}')$$

belongs to $L + L'$. That is, if and only if $\mathbf{x}(\mathbf{PA} - \mathbf{P}'\mathbf{A}')$ belongs to N. Hence, v belongs to $(gH \cap g'H')\pi_{\mathbb{F}}$ if and only if $\mathbf{x} = (v_0^{-1}v)\rho_3 \in M$, i.e., if and only if $v \in M\rho_3^{\leftarrow}\lambda_{v_0}$. □

With all the data already computed, we explicitly have the variety N and, using standard linear algebra, we can compute M (which could be empty, because N may possibly be disjoint with the image of $\mathbf{PA} - \mathbf{P}'\mathbf{A}'$). In this situation, the algorithmic decision on whether $gH \cap g'H'$ is empty or not is straightforward.

Lemma 2.3.7. *With the current notation, and assuming that* $(u \cdot H\pi_{\mathbb{F}}) \cap (u' \cdot H'\pi_{\mathbb{F}}) \neq \varnothing$, *the following are equivalent:*

(a) $gH \cap g'H' = \varnothing$,

(b) $(gH \cap g'H')\pi_{\mathbb{F}} = \varnothing$,

(c) $M\rho_3^{\leftarrow} = \varnothing$,

(d) $M = \varnothing$,

(e) $N \cap \operatorname{im}(\mathbf{PA} - \mathbf{P}'\mathbf{A}') = \varnothing$. □

If $gH \cap g'H' = \varnothing$, we are done. Otherwise, $N \cap \operatorname{im}(\mathbf{PA} - \mathbf{P}'\mathbf{A}') \neq \varnothing$ and we can compute a vector $\mathbf{x} \in \mathbb{Z}^{n_3}$ such that $\mathbf{x}(\mathbf{PA} - \mathbf{P}'\mathbf{A}') \in N$. Take now any preimage of $\mathbf{x}$ by ρ_3, for example $v_1^{x_1} \cdots v_{n_3}^{x_{n_3}}$ if $\mathbf{x} = (x_1, \ldots, x_{n_3})$, and by (2.14), $u'' = v_0 v_1^{x_1} \cdots v_{n_3}^{x_{n_3}} \in (gH \cap g'H')\pi_{\mathbb{F}}$.

It only remains to find $\mathbf{a}'' \in \mathbb{Z}^m$ such that $g'' = t^{\mathbf{a}''}u'' \in gH \cap g'H'$. To do this, observe that $u'' \in (gH \cap g'H')\pi_{\mathbb{F}}$ implies the existence of a vector $\mathbf{a}''$ such that $t^{\mathbf{a}''}u'' \in t^{\mathbf{a}}uH \cap t^{\mathbf{a}'}u'H'$, i.e., such that $t^{\mathbf{a}''-\mathbf{a}}u^{-1}u'' \in H$ and $\mathbf{t}^{\mathbf{a}''-\mathbf{a}'}u'^{-1}u'' \in H'$. In other words, there exists a vector $\mathbf{a}'' \in \mathbb{Z}^m$ such that $\mathbf{a}'' - \mathbf{a} \in \mathcal{C}_H(u^{-1}u'')$ and $\mathbf{a}'' - \mathbf{a}' \in \mathcal{C}_{H'}(u'^{-1}u'')$. That is, the affine varieties $\mathbf{a} + \mathcal{C}_H(u^{-1}u'')$ and $\mathbf{a}' + \mathcal{C}_{H'}(u'^{-1}u'')$ do intersect. By Corollary 1.2.9, we can compute equations for these two varieties, and compute a vector in its intersection. This is the $\mathbf{a}'' \in \mathbb{Z}^m$ we are looking for. $\square$

The above argument applied to the case where $g = g' = 1$ is giving us valuable information about the subgroup intersection $H \cap H'$; this will allow us to solve $\mathsf{SIP}(\mathbb{Z}^m \times \mathbb{F}_n)$ as well. Note that, in this case, $\mathbf{a} = \mathbf{a}' = \mathbf{0}$, $u = u' = 1$ and so, $v_0 = 1$, $w_0 = w_0' = 1$, and $\boldsymbol{\omega_0} = \boldsymbol{\omega_0'} = \mathbf{0}$.

Theorem 2.3.8. *The subgroup intersection problem for $\mathbb{Z}^m \times \mathbb{F}_n$ is solvable.*

Proof. Let $G = \mathbb{Z}^m \times \mathbb{F}_n$ be a finitely generated free-abelian times free group. As in the proof of Theorem 2.3.5, we can assume that the initial finitely generated subgroups $H, H' \leqslant G$ are given by respective bases, i.e., by two sets of elements like in (2.10), $E = \{t^{\mathbf{b}_1}, \ldots, t^{\mathbf{b}_{m_1}}, t^{\mathbf{a}_1}u_1, \ldots, t^{\mathbf{a}_{n_1}}u_{n_1}\}$ and $E' = \{t^{\mathbf{b}_1'}, \ldots, t^{\mathbf{b}_{m_2}'}, t^{\mathbf{a}_1'}u_1', \ldots, t^{\mathbf{a}_{n_2}'}u_{n_2}'\}$. Consider the subgroups $L, L' \leqslant \mathbb{Z}^m$ and the matrices $\mathbf{A} \in \mathcal{M}_{n_1 \times m}(\mathbb{Z})$ and $\mathbf{A}' \in \mathcal{M}_{n_2 \times m}(\mathbb{Z})$ as above. We shall algorithmically decide whether the intersection $H \cap H'$ is finitely generated or not and, in the affirmative case, shall compute a basis for $H \cap H'$.

Let us apply the algorithm from the proof of Theorem 2.3.5 to the cosets $1 \cdot H$ and $1 \cdot H'$; that is, take $g = g' = 1$, i.e., $u = u' = 1$ and $\mathbf{a} = \mathbf{a}' = \mathbf{0}$. Of course, $H \cap H'$ is not empty, and $v_0 = 1$ serves as an element in the intersection, $v_0 \in H \cap H'$. With this choice, the algorithm works with $w_0 = w_0' = 1$ and $\boldsymbol{\omega_0} = \boldsymbol{\omega_0'} = \mathbf{0}$ (so, we can forget the two translation parts in diagram (2.13)). Lemma 2.3.6 tells us that $(H \cap H')\pi_{\mathbb{F}} = M\rho_3^{\leftarrow} \leqslant H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$, where M is the preimage by the linear mapping $\mathbf{PA} - \mathbf{P}'\mathbf{A}' \colon \mathbb{Z}^{n_3} \to \mathbb{Z}^{n_1}$ of the subspace $N = L + L' \leqslant \mathbb{Z}^m$. In this situation, the following lemma decides whether $H \cap H'$ is finitely generated or not.

Lemma 2.3.9. *With the current notation, the following are equivalent:*

(a) $H \cap H'$ *is finitely generated,*

(b) $(H \cap H')\pi_{\mathbb{F}}$ *is finitely generated,*

(c) $M\rho_3^{\leftarrow}$ *is either trivial or of finite index in* $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$*,*

(d) *either* $n_3 = 1$ *and* $M = \{\mathbf{0}\}$*, or* M *is of finite index in* $\mathbb{Z}^{n_3}$*,*

(e) *either* $n_3 = 1$ *and* $M = \{\mathbf{0}\}$*, or* $\mathrm{rk}(M) = n_3$.

Proof. The equivalence (a) $\Leftrightarrow$ (b) is proved in Corollary 1.2.4. The equivalence (b) $\Leftrightarrow$ (c) comes from the well-known fact (see, for example, [LS01] pags. 16-18) that, in the finitely generated free group $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$, the subgroup $(H \cap H')\pi_{\mathbb{F}} = M\rho_3^{\leftarrow}$ is normal and so, finitely generated if and only if it is either trivial or of finite index. But, by Lemma 2.2.3.(ii), the index $[H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} : M\rho_3^{\leftarrow}]$ is finite if and only if

$[\mathbb{Z}^{n_3} : M]$ is finite; this gives (c) $\Leftrightarrow$ (d). The last equivalence is a basic fact in linear algebra. $\square$

We have computed n_3 and an abelian basis for M. If $n_3 = 0$ we immediately deduce that $H \cap H'$ is finitely generated. If $n_3 = 1$ and $M = \{\mathbf{0}\}$ we also deduce that $H \cap H'$ is finitely generated. Otherwise, we check whether $\operatorname{rk}(M)$ equals n_3; if this is the case then again $H \cap H'$ is finitely generated; if not, $H \cap H'$ is infinitely generated.

It only remains to compute a basis for $H \cap H'$, in case it is finitely generated. We know from (1.6) that

$$H \cap H' = \big((H \cap H') \cap \mathbb{Z}^m\big) \times (H \cap H')\pi_{\mathbb{F}}\sigma,$$

where σ is any splitting for $\pi_{\mathbb{F}|H\cap H'} \colon H \cap H' \twoheadrightarrow (H \cap H')\pi_{\mathbb{F}}$; then we can easily get a basis of $H \cap H'$ by putting together a basis of each part. The strategy will be the following: first, we compute an abelian basis for

$$(H \cap H') \cap \mathbb{Z}^m = (H \cap \mathbb{Z}^m) \cap (H' \cap \mathbb{Z}^m) = L \cap L'$$

by just solving a system of linear equations. Second, we shall compute a free basis for $(H \cap H')\pi_{\mathbb{F}}$. And finally, we will construct an explicit splitting σ and will use it to get a free basis for $(H \cap H')\pi_{\mathbb{F}}\sigma$. Putting together these two parts, we shall be done.

To compute a free basis for $(H \cap H')\pi_{\mathbb{F}}$ note that, if $n_3 = 0$, or $n_3 = 1$ and $M = \{\mathbf{0}\}$, then $(H \cap H')\pi_{\mathbb{F}} = 1$ and there is nothing to do. In the remaining case, $\operatorname{rk}(M) = n_3 \geqslant 1$, $M\rho_3^{\leftarrow} = (H \cap H')\pi_{\mathbb{F}}$ has finite index in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$, and so it is finitely generated. We give two alternative options to compute a free basis for it.

The subgroup M has finite index in $\mathbb{Z}^{n_3}$, and we can compute a system of coset representatives of $\mathbb{Z}^{n_3}$ modulo M,

$$\mathbb{Z}^{n_3} = M\mathbf{c_1} \sqcup \cdots \sqcup M\mathbf{c_d}$$

Now, being ρ_3 surjective, and according to Lemma 2.2.3.(ii), we can transfer the previous partition via ρ_3 to obtain a system of right coset representatives of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ modulo $M\rho_3^{\leftarrow}$:

$$H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} = (M\rho_3^{\leftarrow})z_1 \sqcup \cdots \sqcup (M\rho_3^{\leftarrow})z_d, \tag{2.16}$$

where we can take, for example, $z_i = v_1^{c_{i,1}} v_2^{c_{i,2}} \cdots v_{n_3}^{c_{i,n_3}} \in H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$, for each vector $\mathbf{c_i} = (c_{i,1}, c_{i,2}, \ldots, c_{i,n_3}) \in \mathbb{Z}^{n_3}$, $i \in [d]$.

Now let us construct the core of the Schreier graph for $M\rho_3^{\leftarrow} = (H \cap H')\pi_{\mathbb{F}}$ (with respect to $\{v_1, \ldots, v_{n_3}\}$, a free basis for $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$), $\mathcal{S}(M\rho_3^{\leftarrow})$, in the following way: consider the graph with the cosets of (2.16) as vertices, and with no edge. Then, for every vertex $(M\rho_3^{\leftarrow})z_i$ and every letter v_j, add an edge labeled v_j from $(M\rho_3^{\leftarrow})z_i$ to $(M\rho_3^{\leftarrow})z_i v_j$, algorithmically identified among the available vertices by repeatedly using the solvability of the membership problem for $M\rho_3^{\leftarrow}$ (note that we can do this by abelianizing the candidate and checking the defining equations for M). Once we have run over all i, j, we shall get the full graph $\mathcal{S}(M\rho_3^{\leftarrow})$, from which we can easily obtain a free basis for $(H \cap H')\pi_{\mathbb{F}}$ in terms of $\{v_1, \ldots, v_{n_3}\}$.

Alternatively, let $\{\mathbf{m_1}, \ldots, \mathbf{m_{n_3}}\}$ be an abelian basis for M (which we already have from the previous construction), say $\mathbf{m_i} = (m_{i,1}, m_{i,2}, \ldots, m_{i,n_3}) \in \mathbb{Z}^{n_3}$, $i = 1, \ldots, n_3$, and consider the elements $x_i = v_1^{m_{i,1}} v_2^{m_{i,2}} \cdots v_{n_3}^{m_{i,n_3}} \in H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$. It is clear that $M\rho_3^{\leftarrow}$ is the subgroup of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ generated by $x_1, \ldots, x_{n_3}$ and all the infinitely many commutators from elements in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$. But $M\rho_3^{\leftarrow}$ is finitely generated so, finitely many of those commutators will be enough. Enumerate all of them, $y_1, y_2, \ldots$ and keep computing the core $\mathcal{S}_j$ of the Schreier graph for the subgroup $\langle x_1, \ldots, x_{n_3}, y_1, \ldots, y_j \rangle$ for increasing j's until obtaining a complete graph with d vertices (i.e., until reaching a subgroup of index d).

When this happens, we shall have computed the core of the Schreier graph for $M\rho_3^{\leftarrow} = (H \cap H')\pi_{\mathbb{F}}$ (with respect to $\{v_1, \ldots, v_{n_3}\}$, a free basis of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$), from which we can easily find a free basis for $(H \cap H')\pi_{\mathbb{F}}$, in terms of $\{v_1, \ldots, v_{n_3}\}$.

Finally, it remains to compute an explicit splitting σ for $\pi_{\mathbb{F}|H \cap H'} \colon H \cap H' \twoheadrightarrow (H \cap H')\pi_{\mathbb{F}}$. We have a free basis $\{z_1, \ldots, z_d\}$ for $(H \cap H')\pi_{\mathbb{F}}$, in terms of $\{v_1, \ldots, v_{n_3}\}$; so, using the expressions $v_i = v_i(u_1, \ldots, u_{n_1})$ that we have from the beginning of the proof, we can get expressions $z_i = \eta_i(u_1, \ldots, u_{n_1})$. From here,

$$\eta_i(t^{\mathbf{a_1}} u_1, \ldots, t^{\mathbf{a_{n_1}}} u_{n_1}) = t^{\mathbf{e_i}} z_i \in H,$$

and projects to z_i, so $\mathcal{C}_{z_i, H} = \mathbf{e_i} + L$ (see Corollary 1.2.9), $i \in [d]$. Similarly, we can get vectors $\mathbf{e}_i' \in \mathbb{Z}^m$ such that $\mathcal{C}_{z_i, H'} = \mathbf{e}_i' + L'$. Since, by construction, $\mathcal{C}_{z_i, H \cap H'} = \mathcal{C}_{z_i, H} \cap \mathcal{C}_{z_i, H'}$ is a non-empty affine variety in $\mathbb{Z}^m$ with direction $L \cap L'$, we can compute vectors $\mathbf{e}_i'' \in \mathbb{Z}^m$ on it by just solving the corresponding systems of linear equations, $i \in [d]$. Now, $z_i \mapsto t^{\mathbf{e}_i''} z_i$ is the desired splitting $H \cap H' \overset{\sigma}{\leftarrow} (H \cap H')\pi_{\mathbb{F}}$, and $\{t^{\mathbf{e}_1''} z_1, \ldots, t^{\mathbf{e}_d''} z_d\}$ is the free basis for $(H \cap H')\pi_{\mathbb{F}}\sigma$ we were looking for.

As mentioned above, putting together this free basis with the abelian basis we already have for $L \cap L'$, we get a basis for $H \cap H'$, concluding the proof. □

Corollary 2.3.10. *Let H, H' be two free non-abelian subgroups of finite rank in $\mathbb{Z}^m \times \mathbb{F}_n$. With the previous notation, the intersection $H \cap H'$ is finitely generated if and only if either $H \cap H' = 1$, or $\mathbf{PA} = \mathbf{P'A'}$.*

Proof. Under the conditions of the statement, we have $L = L' = \{\mathbf{0}\}$. Hence, $N = L + L' = \{\mathbf{0}\}$ and its preimage by $\mathbf{PA} - \mathbf{P'A'}$ is $M = \ker(\mathbf{PA} - \mathbf{P'A'}) \leqslant \mathbb{Z}^{n_3}$. Now, by Lemma 2.3.9, $H \cap H'$ is finitely generated if and only if either

$$(H \cap H')\pi_{\mathbb{F}} = M\rho_3^{\leftarrow} = 1\,, \quad \text{or} \quad n_3 - \operatorname{rk}(\operatorname{im}(\mathbf{PA} - \mathbf{P'A'})) = \operatorname{rk}(M) = n_3\,;$$

that is, if and only if either

$$(H \cap H')\pi_{\mathbb{F}} = 1\,, \quad \text{or} \quad \mathbf{PA} = \mathbf{P'A'}\,.$$

But, since $L = L' = \{\mathbf{0}\}$, $(H \cap H')\pi_{\mathbb{F}} = 1$ if and only if $H \cap H' = 1$. □

We consider now two examples to illustrate the preceding algorithm.

Example 2.3.11. Let us analyze again the example given in the proof of Lemma 2.3.1, under the light of the previous corollary. We considered in $\mathbb{Z} \times \mathbb{F}_2 = \langle t \mid \rangle \times \langle a, b \mid \rangle$ the subgroups $H = \langle a, b \rangle$ and $H' = \langle ta, b \rangle$, both free non-abelian of rank 2. It is clear that $\mathbf{A} = \left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right]$ and $\mathbf{A'} = \left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right]$, while $H\pi_{\mathbb{F}} = H'\pi_{\mathbb{F}} = H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} = \mathbb{F}_2$; in particular, $n_3 = 2$ and $H \cap H' \neq 1$. In these circumstances, both inclusions $H\pi_{\mathbb{F}} \hookleftarrow H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} \hookrightarrow H'\pi_{\mathbb{F}}$ are the identity maps, so $\mathbf{P} = \mathbf{P'} = \mathbf{1}$ is the 2×2 identity matrix and hence, $\mathbf{PA} = \left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right] \neq \left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right] = \mathbf{P'A'}$. According to Corollary 2.3.10, this means that $H \cap H'$ is not finitely generated, as we had seen before.

Example 2.3.12. Consider two finitely generated subgroups $H, H' \leqslant \mathbb{F}_n \leqslant \mathbb{Z}^m \times \mathbb{F}_n$. In this case we have $\mathbf{A} = (\mathbf{0}) \in \mathcal{M}_{n_1,m}$ and $\mathbf{A'} = (\mathbf{0}) \in \mathcal{M}_{n_2,m}$ and so, $\mathbf{PA} = (\mathbf{0}) = \mathbf{P'A'}$. Thus, Corollary 2.3.10 just corroborates Howson's property for $\mathbb{F}_n$.

To finish this section, we present an application of Theorem 2.3.8 to a nice geometric problem. In [Sah15], Sahattchieve studies quasi-convexity of subgroups of $\mathbb{Z}^m \times \mathbb{F}_n$ with respect to the natural component-wise action of $\mathbb{Z}^m \times \mathbb{F}_n$ on the product space, $\mathbb{R}^m \times T_n$, of the m-dimensional euclidean space and the regular $(2n)$-valent infinite tree T_n: a subgroup $H \leqslant \mathbb{Z}^m \times \mathbb{F}_n$ is *quasi-convex* if the orbit Hp of some (and hence every) point $p \in \mathbb{R}^m \times T_n$ is a quasi-convex subset of $\mathbb{R}^m \times T_n$ (see [Sah15] for more details). One of the results obtained is the following characterization.

Theorem 2.3.13 (Sahattchieve)**.** *Let H be a subgroup of $\mathbb{Z}^m \times \mathbb{F}_n$. Then, H is quasi-convex if and only if H is either cyclic or virtually of the form $B \times A$, for some finitely generated $B \leqslant_{fg} \mathbb{Z}^m$ and $A \leqslant_{fg} \mathbb{F}_n$. (In particular, quasi-convex subgroups are finitely generated.)*

Combining this with our Theorem 2.3.8, we can easily establish an algorithm to decide whether a given finitely generated subgroup of $\mathbb{Z}^m \times \mathbb{F}_n$ is quasi-convex or not (with respect to the above mentioned action).

Corollary 2.3.14. *There is an algorithm which, given a finite list* $h_1, \ldots, h_s$ *of elements in* $\mathbb{Z}^m \times \mathbb{F}_n$*, decides whether the subgroup* $H = \langle h_1, \ldots, h_s \rangle$ *is quasi-convex or not.*

Proof. First, apply Proposition 1.2.6 to compute a basis for H. If it contains only one element, then H is cyclic and we are done.

Otherwise (H is not cyclic) we can easily compute a free-abelian basis and a free basis for the respective projections $H\pi_{\mathbb{A}} \leqslant \mathbb{Z}^m$ and $H\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$. From the basis for H we can immediately extract a free-abelian basis for $\mathbb{Z}^m \cap H = H\pi_{\mathbb{A}} \cap H$. And, using Theorem 2.3.8, we can decide whether $\mathbb{F}_n \cap H = H\pi_{\mathbb{F}} \cap H$ is finitely generated or not and, in the affirmative case, compute a free basis for it.

Finally, we can decide whether $H\pi_{\mathbb{A}} \cap H \leqslant_{fi} H\pi_{\mathbb{A}}$ and $H\pi_{\mathbb{F}} \cap H \leqslant_{fi} H\pi_{\mathbb{F}}$ hold or not (applying the well-known solutions to $\mathsf{FIP}(\mathbb{Z}^m)$ and $\mathsf{FIP}(\mathbb{F}_{n'})$ or, alternatively, using the more general Theorem 2.2.9 above); note that if we detected that $H\pi_{\mathbb{F}} \cap H$ is infinitely generated then it must automatically be of infinite index in $H\pi_{\mathbb{F}}$ (which, of course, is finitely generated).

We claim that H is quasi-convex if and only if $H\pi_{\mathbb{A}} \cap H \leqslant_{fi} H\pi_{\mathbb{A}}$ and $H\pi_{\mathbb{F}} \cap H \leqslant_{fi} H\pi_{\mathbb{F}}$; this will conclude the proof.

For the implication to the right (and applying Theorem 2.3.13), assume that $B \times A \leqslant_{fi} H$ for some $B \leqslant \mathbb{Z}^m$ and $A \leqslant \mathbb{F}_n$ being finitely generated. Applying $\pi_{\mathbb{A}}$ and $\pi_{\mathbb{F}}$ we get $B \leqslant_{fi} H\pi_{\mathbb{A}}$ and $A \leqslant_{fi} H\pi_{\mathbb{F}}$, respectively (see Lemma 2.2.3.(i)). But $B \leqslant H\pi_{\mathbb{A}} \cap H \leqslant H\pi_{\mathbb{A}}$ and $A \leqslant H\pi_{\mathbb{F}} \cap H \leqslant H\pi_{\mathbb{F}}$ hence, $H\pi_{\mathbb{A}} \cap H \leqslant_{fi} H\pi_{\mathbb{A}}$ and $H\pi_{\mathbb{F}} \cap H \leqslant_{fi} H\pi_{\mathbb{F}}$.

For the implication to the left, assume $H\pi_{\mathbb{A}} \cap H \leqslant_{fi} H\pi_{\mathbb{A}}$ and $H\pi_{\mathbb{F}} \cap H \leqslant_{fi} H\pi_{\mathbb{F}}$ (and, in particular, $H\pi_{\mathbb{F}} \cap H$ finitely generated). Take $B = H\pi_{\mathbb{A}} \cap H \leqslant_{fi} H\pi_{\mathbb{A}} \leqslant \mathbb{Z}^m$ and $A = H\pi_{\mathbb{F}} \cap H \leqslant_{fi} H\pi_{\mathbb{F}} \leqslant \mathbb{F}_n$, and we get $B \times A \leqslant_{fi} H\pi_{\mathbb{A}} \times H\pi_{\mathbb{F}}$ (see Lemma 2.2.5). But H is in between, $B \times A \leqslant H \leqslant H\pi_{\mathbb{A}} \times H\pi_{\mathbb{F}}$, hence $B \times A \leqslant_{fi} H$ and, by Theorem 2.3.13, H is quasi-convex. □

2.4 Fixed subgroups

In this section we shall study when the subgroup fixed by an endomorphism or an automorphism of $\mathbb{Z}^m \times \mathbb{F}_n$ is finitely generated and, in this case, consider the problem of computing a basis for it.

Fixed Points Problems, $\mathsf{FPP_A}(G)$**,** $\mathsf{FPP_E}(G)$**.** *Given an automorphism (resp. endomorphism)* Ψ *of* G *by images of generators, decide whether* $\operatorname{Fix}\Psi$ *is finitely generated and, if so, compute a generating set.*

Of course, the fixed point subgroup of an arbitrary endomorphism of $\mathbb{Z}^m$ is finitely generated, and the problems $\mathsf{FPP_A}(\mathbb{Z}^m)$ and $\mathsf{FPP_E}(\mathbb{Z}^m)$ are clearly solvable, just reducing them to solve the corresponding systems of linear equations.

Again, the case of free groups is much more complicated, and has a rich and interesting history (see [Ven02] for a historical survey). After several conjectures about the finite generability of $\operatorname{Fix}\alpha$, for different kinds of automorphisms $\alpha \in \operatorname{Aut}(\mathbb{F}_n)$, Gersten finally proved in [Ger87] that $\operatorname{rk}(\operatorname{Fix}\alpha) < \infty$ for every automorphism $\alpha \in \operatorname{Aut}(\mathbb{F}_n)$; and shortly after, Goldstein and Turner [GT86] extended this result to arbitrary endomorphisms of $\mathbb{F}_n$.

After these results, the natural quest for possible bounds of this rank ended up with the publication in 1992 of the seminal paper [BH92] by M. Bestvina, and M. Handel introducing the powerful technique of train tracks, and providing the first (and furthermore tight) uniform upper bound for the rank of the subgroup of fixed points of an automorphism $\alpha \in \operatorname{Aut}(\mathbb{F}_n)$, namely $\operatorname{rk}(\operatorname{Fix}\alpha) \leqslant n$.

Regarding computability, some partial results were obtained (see, for example, Cohen and Lustig [CL89] for positive automorphisms, Turner [Tur95] for special irreducible automorphisms, and Bogopolski [Bog00] for the case $n = 2$) before O. Bogopolski and O. Maslakova, making strong use of the theory of train tracks, finally published in [BM15] an algorithm to compute a free basis for $\operatorname{Fix}\alpha$, for a general automorphism $\alpha \in \operatorname{Aut}\mathbb{F}_n$ (see also the preprint [FH14] for an alternative proof).

Theorem 2.4.1 (Bogopolski and Maslakova, 2015, [BM15])**.** *There exists an algorithm which, given an automorphism* α *of* $\mathbb{F}_n$ *finds a basis of its fixed point subgroup* $\operatorname{Fix}(\alpha) = \{w \in \mathbb{F}_n \,:\, w\alpha = w\}$. □

On the other hand, as far as we know, the problem $\mathsf{FPP_E}(\mathbb{F}_n)$ remains still open in general.

When one moves to free-abelian times free groups, the situation is even more involved. Similar to what happens with respect to the Howson property, $\operatorname{Fix}\Psi$ is not necessarily finitely generated for $\Psi \in \operatorname{Aut}(\mathbb{Z} \times \mathbb{F}_2)$, and essentially the same example from Lemma 2.3.1 can be recycled here.

Example 2.4.2. Consider the type I automorphism Ψ given by $a \mapsto ta$, $b \mapsto b$, $t \mapsto t$; clearly, $t^r\, w(a,b) \mapsto t^{r+|w|_a}\, w(a,b)$ and so,

$$\operatorname{Fix}\Psi = \{\, t^r w(a,b) \mid |w|_a = 0 \,\} = \langle\!\langle t,\, b \rangle\!\rangle = \langle t, a^{-k} b a^k \ \ (k \in \mathbb{Z}) \rangle$$

is not finitely generated.

In the present section we shall analyze the fixed point subgroup of an endomorphism of a free-abelian times free group, and give an explicit characterization of when it is finitely generated. In this case, we shall also consider the computability of a finite basis, and solve $\mathsf{FPP_A}(\mathbb{Z}^m \times \mathbb{F}_n)$ and $\mathsf{FPP_E}(\mathbb{Z}^m \times \mathbb{F}_n)$ modulo the corresponding problems for free groups, $\mathsf{FPP_A}(\mathbb{F}_n)$ and $\mathsf{FPP_E}(\mathbb{F}_n)$. (Our arguments descend directly from $\operatorname{End}(\mathbb{Z}^m \times \mathbb{F}_n)$ to $\operatorname{End}(\mathbb{F}_n)$ in such a way that any partial solution to the free problems can be used to give the corresponding partial solution to the free-abelian times free problems, see Proposition 2.4.6 below.)

Let us distinguish the two types of endomorphisms according to Proposition 1.3.1 (and starting with the easier type II ones).

Proposition 2.4.3. *Let $G = \mathbb{Z}^m \times \mathbb{F}_n$ with $n \neq 1$, and consider a type II endomorphism Ψ, namely*

$$\Psi = \Psi_{z,\mathbf{l},\mathbf{h},\mathbf{Q},\mathbf{P}} : t^{\mathbf{a}}u \mapsto t^{\mathbf{aQ}+\mathbf{uP}} z^{\mathbf{a}\mathbf{l}^\top + \mathbf{u}\mathbf{h}^\top},$$

where $1 \neq z \in \mathbb{F}_n$ is not a proper power, $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$, $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$, $\mathbf{0} \neq \mathbf{l} \in \mathbb{Z}^m$, and $\mathbf{h} \in \mathbb{Z}^n$. Then, $\operatorname{Fix}\Psi$ is finitely generated, and a basis for $\operatorname{Fix}\Psi$ is computable.

Proof. First note that $\operatorname{im}\Psi$ is an abelian subgroup of $\mathbb{Z}^m \times \mathbb{F}_n$. Then, by Corollary 1.2.3, it must be isomorphic to $\mathbb{Z}^{m'}$ for a certain $m' \leqslant m+1$. Therefore, $\operatorname{Fix}\Psi \leqslant \operatorname{im}(\Psi)$ is isomorphic to a subgroup of $\mathbb{Z}^{m'}$, and thus finitely generated.

According to the definition, an element $t^{\mathbf{a}}u$ is fixed by Ψ if and only if $t^{\mathbf{aQ}+\mathbf{uP}} z^{\mathbf{a}\mathbf{l}^\top + \mathbf{u}\mathbf{h}^\top} = t^{\mathbf{a}}u$. For this to be satisfied, u must be a power of z, say $u = z^r$ for certain $r \in \mathbb{Z}$, and abelianizing we get $\mathbf{u} = r\mathbf{z}$, and the system of equations

$$\left.\begin{aligned} \mathbf{a}\mathbf{l}^\top + r\mathbf{z}\mathbf{h}^\top &= r \\ \mathbf{a}(\mathbf{I}_m - \mathbf{Q}) &= r\mathbf{z}\mathbf{P} \end{aligned}\right\} \tag{2.17}$$

whose set $\mathcal{S}$ of integral solutions $(\mathbf{a}, r) \in \mathbb{Z}^{m+1}$ describe precisely the subgroup of fixed points by Ψ:

$$\operatorname{Fix}\Psi = \{ t^{\mathbf{a}} z^r : (\mathbf{a}, r) \in \mathcal{S} \}. \qquad \square$$

Theorem 2.4.4. *Let $G = \mathbb{Z}^m \times \mathbb{F}_n$ with $n \neq 1$, and consider a type I endomorphism Ψ, namely*

$$\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}} : t^{\mathbf{a}}\, u \mapsto t^{\mathbf{aQ}+\mathbf{uP}}\, u\phi,$$

where $\phi \in \operatorname{End}(\mathbb{F}_n)$, $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$, and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$. Let $N = \operatorname{im}(\mathbf{I}_m - \mathbf{Q}) \cap \operatorname{im}\mathbf{P_o}$, where $\mathbf{P_o}$ is the restriction of $\mathbf{P}\colon \mathbb{Z}^n \to \mathbb{Z}^m$ to $(\operatorname{Fix}\phi)\rho$, the image of $\operatorname{Fix}\phi \leqslant \mathbb{F}_n$ under the global abelianization $\rho\colon \mathbb{F}_n \twoheadrightarrow \mathbb{Z}^n$.

Then, $\operatorname{Fix}\Psi$ is finitely generated if and only if one of the following happens:

(a) $\operatorname{Fix}\phi = 1$;

(b) $\operatorname{Fix}\phi$ *is cyclic,* $(\operatorname{Fix}\phi)\rho \neq \{\mathbf{0}\}$, *and* $\mathrm{N}\mathbf{P}_{\mathbf{o}}^{\leftarrow} = \{\mathbf{0}\}$; *or*

(c) $\operatorname{rk}(\mathrm{N}) = \operatorname{rk}(\operatorname{im}\mathbf{P}_{\mathbf{o}})$.

Proof. An element $t^{\mathbf{a}}u$ is fixed by Ψ if and only if $t^{\mathbf{aQ}+\mathbf{uP}}u\phi = t^{\mathbf{a}}u$, i.e., if and only if

$$\left.\begin{aligned} u\phi &= u \\ \mathbf{a}(\mathbf{I_m} - \mathbf{Q}) &= \mathbf{uP} \end{aligned}\right\}$$

That is,

$$\operatorname{Fix}\Psi \;=\; \{t^{\mathbf{a}}u \in G \;:\; u \in \operatorname{Fix}\phi \;\text{ and }\; \mathbf{a}(\mathbf{I_m} - \mathbf{Q}) = \mathbf{uP}\}, \tag{2.18}$$

where $\mathbf{u} = u\rho$, and $\rho\colon \mathbb{F}_n \twoheadrightarrow \mathbb{Z}^n$ is the abelianization map. As we have seen in Corollary 1.2.4, $\operatorname{Fix}\Psi$ is finitely generated if and only if so is its projection to the free part

$$(\operatorname{Fix}\Psi)\pi_{\mathbb{F}} \;=\; \operatorname{Fix}\phi \,\cap\, \{u \in \mathbb{F}_n \;:\; \mathbf{uP} \in \operatorname{im}(\mathbf{I_m} - \mathbf{Q})\}. \tag{2.19}$$

Now (identifying integral matrices $\mathbf{A}$ with the corresponding linear mappings $\mathbf{v} \mapsto \mathbf{vA}$, as usual), let M be the image of $\mathbf{I_m} - \mathbf{Q}$, and consider its preimage first by $\mathbf{P}$ and then by ρ, see the following diagram:

$$\begin{array}{ccccccc} \operatorname{Fix}\phi \leqslant & \mathbb{F}_n & \xrightarrow{\;\rho\;} & \mathbb{Z}^n & \xrightarrow{\;\mathbf{P}\;} & \mathbb{Z}^m \circlearrowleft^{\mathbf{I_m}-\mathbf{Q}} & \\ & \triangledown\!\!/ & & \triangledown\!\!/ & & \triangledown\!\!/ & \\ & \mathrm{M}\mathbf{P}^{\leftarrow}\rho^{\leftarrow} & \longleftarrow & \mathrm{M}\mathbf{P}^{\leftarrow} & \longleftarrow & \mathrm{M} = \operatorname{im}(\mathbf{I_m} - \mathbf{Q}) & \end{array} \tag{2.20}$$

Equation (2.19) can be rewritten as

$$(\operatorname{Fix}\Psi)\pi_{\mathbb{F}} \;=\; \operatorname{Fix}\phi \,\cap\, \mathrm{M}\mathbf{P}^{\leftarrow}\rho^{\leftarrow}. \tag{2.21}$$

However, this description does not show whether $\operatorname{Fix}\Psi$ is finitely generated because $\operatorname{Fix}\phi$ is in fact finitely generated, but $\mathrm{M}\mathbf{P}^{\leftarrow}\rho^{\leftarrow}$ is not in general. We shall avoid the intersection with the whole fixed subgroup $\operatorname{Fix}\phi$ by reducing M to a certain subgroup. Let ρ_{o} be the restriction of ρ to $\operatorname{Fix}\phi$ (not to be confused with the abelianization map of the subgroup $\operatorname{Fix}\phi$ itself), let $\mathbf{P_o}$ be the restriction of $\mathbf{P}$ to $\operatorname{im}\rho_o$, and let $N = M \cap \operatorname{im}\mathbf{P_o}$, see the following diagram:

$$
\begin{array}{ccccccc}
 & & & & \overset{\mathbf{I_m}-\mathbf{Q}}{\circlearrowright} & & \\
\mathbb{F}_n & \xrightarrow{\ \rho\ } & \mathbb{Z}^n & \xrightarrow{\ \mathbf{P}\ } & \mathbb{Z}^m & \geqslant & M = \operatorname{im}(\mathbf{I_m}-\mathbf{Q}) \\
\rotatebox{90}{$\leqslant$} & & \rotatebox{90}{$\leqslant$} & & \rotatebox{90}{$\leqslant$} & & \\
\operatorname{Fix}\phi & \xrightarrow{\ \rho_o\ } & \operatorname{im}\rho_o & \xrightarrow{\ \mathbf{P_o}\ } & \operatorname{im}\mathbf{P_o} & & \\
\rotatebox{90}{$\leqslant$} & & \rotatebox{90}{$\leqslant$} & & \rotatebox{90}{$\leqslant$} & & \\
(\operatorname{Fix}\Psi)\pi_{\mathbb{F}} = N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow} & \longleftarrow & N\mathbf{P_o}^{\leftarrow} & \longleftarrow & N & = & M \cap \operatorname{im}\mathbf{P_o}\,.
\end{array}
\tag{2.22}
$$

Fig. 2.3: Fixed points subgroup diagram for FATF groups

Then, Equation (2.21) rewrites into

$$(\operatorname{Fix}\Psi)\pi_{\mathbb{F}} = N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow}.$$

Now, since $N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow}$ is a normal subgroup of $\operatorname{Fix}\phi$ (not, in general, of $\mathbb{F}_n$), it is finitely generated if and only if it is either trivial, or of finite index in $\operatorname{Fix}\phi$.
Note that the restricted abelianization ρ_o is injective (and thus bijective) if and only if $\operatorname{Fix}\phi$ is either trivial, or cyclic not abelianizing to zero (indeed, for this to be the case we cannot have two freely independent elements in $\operatorname{Fix}\phi$ and so, $\operatorname{rk}(\operatorname{Fix}\phi) \leqslant 1$).
Thus, $(\operatorname{Fix}\Psi)\pi_{\mathbb{F}} = N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow} = 1$ if and only if $\operatorname{Fix}\phi$ is trivial; or $\operatorname{Fix}\phi$ is cyclic not abelianizing to zero, and $N\mathbf{P_o}^{\leftarrow} = \{\mathbf{0}\}$.
On the other side, by Lemma 2.2.3.(ii), the preimage $N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow}$ has finite index in $\operatorname{Fix}\phi$ if and only if N has finite index in $\operatorname{im}\mathbf{P_o}$ i.e., if and only if $\operatorname{rk}(N) = \operatorname{rk}(\operatorname{im}\mathbf{P_o})$. □

Example 2.4.5. Let us analyze again Example 2.4.2, under the light of the Theorem 2.4.4. We considered the automorphism Ψ of $\mathbb{Z} \times \mathbb{F}_2 = \langle t \mid \rangle \times \langle a, b \mid \rangle$ given by $a \mapsto ta$, $b \mapsto b$ and $t \mapsto t$. That is, $\Psi = \Psi_{\mathrm{id}_{\mathbb{F}_2}, \mathbf{I_1}, \mathbf{P}}$, where $\mathbf{P}$ is the 2×1 matrix $\mathbf{P} = (\mathbf{1}, \mathbf{0})^{\mathsf{T}}$. Now, it is clear that $\operatorname{Fix}\mathrm{id}_{\mathbb{F}_2} = \mathbb{F}_2$ and so, conditions (a) and (b) from Theorem 2.4.4 do not hold. Furthermore, $\rho_o = \rho$, $\mathbf{P_o} = \mathbf{P}$, $M = \operatorname{im}(\mathbf{0}) = \{\mathbf{0}\}$, $N = \{\mathbf{0}\}$, while $\operatorname{im}\mathbf{P_o} = \mathbb{Z}$. Hence, condition (c) from Theorem 2.4.4 does not hold either, according to the fact that $\operatorname{Fix}\Psi$ is not finitely generated.

Finally, the proof of Theorem 2.4.4 is explicit enough to allow us to make the whole argument algorithmic: given a type I endomorphism $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}} \in \operatorname{End}(\mathbb{Z}^m \times \mathbb{F}_n)$, the decision on whether $\operatorname{Fix}\Psi$ is finitely generated or not, and the computation of a basis for it affirmative case, can be made effective assuming we have a procedure to compute a (free) basis for $\operatorname{Fix}\phi$.

Proposition 2.4.6. *Let* $G = \mathbb{Z}^m \times \mathbb{F}_n$ *with* $n \neq 1$, *and let* $\Psi = \Psi_{\phi,\mathbf{Q},\mathbf{P}}$ *be a type I endomorphism of* G. *Assuming a (finite and free) basis for* $\operatorname{Fix}\phi$ *is given to us, we can* *algorithmically decide whether* $\operatorname{Fix}\Psi$ *is finitely generated or not and, in case it is, compute a basis for it.*

Proof. Let $\{v_1, \ldots, v_p\}$ be the (finite and free) basis for $\operatorname{Fix}\phi \leqslant \mathbb{F}_n$ given to us in the hypothesis.

Theorem 2.4.4 describes $\operatorname{Fix}\Psi$ and when is it finitely generated. Assuming the notation from the proof there, we can compute abelian bases for $N \leqslant \operatorname{im}\mathbf{P_o} \leqslant \mathbb{Z}^m$ and $N\mathbf{P_o}^{\leftarrow} \leqslant \operatorname{im}\rho_o \leqslant \mathbb{Z}^n$. Then, we can easily check whether any of the following three conditions hold:

(a) $\operatorname{Fix}\phi$ is trivial,

(b) $\operatorname{Fix}\phi = \langle z \rangle$, $z\rho \neq \mathbf{0}$ and $N\mathbf{P_o}^{\leftarrow} = \{\mathbf{0}\}$, for certain $z \in \mathbb{F}_n$,

(c) $\operatorname{rk}(N) = \operatorname{rk}(\operatorname{im}\mathbf{P_o})$.

If all three conditions fail, then $\operatorname{Fix}\Psi$ is not finitely generated and we are done. Otherwise, $\operatorname{Fix}\Psi$ is finitely generated and it remains to compute a basis. From (1.6), we have

$$\operatorname{Fix}\Psi = \big((\operatorname{Fix}\Psi) \cap \mathbb{Z}^m\big) \times (\operatorname{Fix}\Psi)\pi_{\mathbb{F}}\sigma,$$

where $\operatorname{Fix}\Psi \xleftarrow{\sigma} (\operatorname{Fix}\Psi)\pi_{\mathbb{F}}$ is any splitting of $\pi_{\mathbb{F}|\operatorname{Fix}\Psi}\colon \operatorname{Fix}\Psi \twoheadrightarrow (\operatorname{Fix}\Psi)\pi_{\mathbb{F}}$. We just have to compute a basis for each part and put them together (after computing some splitting σ). Regarding the abelian part, equation (2.18) tells us that

$$(\operatorname{Fix}\Psi) \cap \mathbb{Z}^m = \{\, t^{\mathbf{a}} \,:\, \mathbf{a}(\mathbf{I_m} - \mathbf{Q}) = \mathbf{0} \,\},$$

and we can easily find an abelian basis for it by just computing $\ker(\mathbf{I_m} - \mathbf{Q})$.

Consider now the free part. In cases (a) and (b), $(\operatorname{Fix}\Psi)\pi_{\mathbb{F}} = 1$ and there is nothing to compute. Note that, in these cases, $\operatorname{Fix}\Psi$ is an abelian subgroup of $\mathbb{Z}^m \times \mathbb{F}_n$.

Assume case (c), i.e., $\operatorname{rk}(N) = \operatorname{rk}(\operatorname{im}\mathbf{P_o})$. In this situation, N has finite index in $\operatorname{im}\mathbf{P_o}$ and so, $N\mathbf{P_o}^{\leftarrow}$ has finite index in $\operatorname{im}\rho_o$. So, we can effectively compute a set of coset representatives of $\operatorname{im}\rho_o$ modulo $N\mathbf{P_o}^{\leftarrow}$, say

$$\operatorname{im}\rho_o = (N\mathbf{P_o}^{\leftarrow})\mathbf{c_1} \sqcup \cdots \sqcup (N\mathbf{P_o}^{\leftarrow})\mathbf{c_q}.$$

Now, according to Lemma 2.2.3 (b), we can transfer this partition via ρ_o to obtain a system of right coset representatives of $\operatorname{Fix}\phi$ modulo $(\operatorname{Fix}\Psi)\pi_{\mathbb{F}} = N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow}$,

$$\operatorname{Fix}\phi = (N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow})z_1 \sqcup \cdots \sqcup (N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow})z_q. \tag{2.23}$$

To compute the z_i's, note that $\mathbf{v_1} = v_1\rho_o, \ldots, \mathbf{v_p} = v_p\rho_o$ generate $\operatorname{im}\rho_o$, write each $\mathbf{c_i} \in \operatorname{im}\rho_o$ as a (not necessarily unique) linear combination of them, say $\mathbf{c_i} = c_{i,1}\mathbf{v_1} + \cdots + c_{i,p}\mathbf{v_p}$, $i \in [q]$, and take $z_i = v_1^{c_{i,1}} v_2^{c_{i,1}} \cdots v_p^{c_{i,p}} \in \operatorname{Fix}\phi$.

Now, construct a free basis for $N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow} = (\operatorname{Fix}\Psi)\pi_{\mathbb{F}}$ following the first of the two alternatives at the end of the proof of Theorem 2.3.8 (the second one does not work here because ρ_o is not the abelianization of the subgroup $\operatorname{Fix}\phi$, but the restriction there of the abelianization of $\mathbb{F}_n$):

Build the Schreier graph $\mathsf{Sch}\left(N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow}\right)$ for $N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow} \leqslant \operatorname{Fix}\phi$ with respect to $\{v_1, \ldots, v_p\}$, in the following way: consider the graph with the cosets of (2.23) as vertices, and with no edge. Then, for every vertex $(N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow})z_i$ and every letter v_j, add an edge labeled v_j from $(N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow})z_i$ to $(N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow})z_i v_j$, algorithmically identified among the available vertices by repeatedly using the membership problem for $N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow}$ (note that we can easily do this by abelianizing the candidate and checking whether it belongs to $N\mathbf{P_o}^{\leftarrow}$). Once we have run over all i, j, we shall get the full graph $\mathsf{Sch}\left(N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow}\right)$, from which we can easily obtain a free basis for $N\mathbf{P_o}^{\leftarrow}\rho_o^{\leftarrow} = (\operatorname{Fix}\Psi)\pi_{\mathbb{F}}$.

Finally, having a free basis for $(\operatorname{Fix}\Psi)\pi_{\mathbb{F}}$, we can easily construct a splitting for $\pi_{\mathbb{F}|\operatorname{Fix}\Psi}\colon \operatorname{Fix}\Psi \twoheadrightarrow (\operatorname{Fix}\Psi)\pi_{\mathbb{F}}$ by just computing, for each generator $u \in (\operatorname{Fix}\Psi)\pi_{\mathbb{F}}$, a preimage $t^{\mathbf{a}}u \in \operatorname{Fix}\Psi$, where $\mathbf{a} \in \mathbb{Z}^m$ is a completion found by solving the system of equations $\mathbf{a}(\mathbf{I_m} - \mathbf{Q}) = \mathbf{uP}$ (see (2.18)).

This completes the proof. □

Bringing together Propositions 2.4.3 and 2.4.6, and Theorem 2.4.4, we get the following result.

Theorem 2.4.7. *The fixed points problem* $\mathsf{FPP_A}(\mathbb{Z}^m \times \mathbb{F}_n)$ *is solvable.* □

Since $\mathsf{FPP_E}(\mathbb{F}_n)$ is still open, the corresponding result about endomorphisms of $\mathbb{Z}^m \times \mathbb{F}_n$, can only be stated in conditional form at this point.

Theorem 2.4.8. *If the fixed points problem* $\mathsf{FPP_E}(\mathbb{F}_n)$ *is solvable then* $\mathsf{FPP_E}(\mathbb{Z}^m \times \mathbb{F}_n)$ *is also solvable.* □

To close this section, we point the reader to some very recent results related to fixed subgroups of endomorphisms of partially commutative groups. In [RSS13], E. Rodaro, P.V. Silva and M. Sykiotis characterize which partially commutative groups G satisfy that $\operatorname{Fix}\Psi$ is finitely generated for every $\Psi \in \operatorname{End}(G)$ (and, of course, free-abelian times free groups are included there); they also provide similar results concerning automorphisms.

2.5 Whitehead problems

J.H.C. Whitehead, back in the 1930's, gave an algorithm [Whi36] to decide, given two elements u and v from a finitely generated free group $\mathbb{F}_n$, whether there exists an automorphism $\phi \in \operatorname{Aut}(\mathbb{F}_n)$ sending one to the other, $v = u\phi$. Whitehead's algorithm uses a (today) very classical piece of combinatorial group theory technique called 'peak reduction', see also [LS01]. Several variations of this problem (like replacing u and v by tuples of words, relaxing equality to equality up to conjugacy, adding conditions on the conjugators, replacing words by subgroups, replacing automorphisms to monomorphisms or endomorphisms, etc.), as well as extensions of all these problems to other families of groups, can be found in the literature, all of them generically known as *Whitehead problems*. Here, we state the standard Whitehead problems for an arbitrary group G.

Whitehead problems, $\mathsf{WhP}_\mathsf{A}(G)$**,** $\mathsf{WhP}_\mathsf{M}(G)$**,** $\mathsf{WhP}_\mathsf{E}(G)$**.** *Given two elements* $u, v \in G$*, decide whether there exists an automorphism (resp. monomorphism, endomorphism)* ϕ *of* G *such that* $u\phi = v$*; and if so, find one by giving images of generators.*

In this section we shall solve these three problems for free-abelian times free groups. We note that a new version of the classical peak-reduction theorem has been developed by M. Day [Day14] for an arbitrary partially commutative group, see also [Day09]. These techniques allow the author to solve the Whitehead problem for partially commutative groups, in its form relative to automorphisms and tuples of conjugacy classes. In particular $\mathsf{WhP}_\mathsf{A}(G)$ (which was conjectured in [Day09]) is solved in [Day14] for any partially commutative group G. As far as we know, $\mathsf{WhP}_\mathsf{M}(G)$ and $\mathsf{WhP}_\mathsf{E}(G)$ remain unsolved in general for a PC-group G. Our Theorem 2.5.5 below is a small contribution into this direction, solving these problems for free-abelian times free groups.

Let us begin by reminding the situation of the Whitehead problems for free-abelian and for free groups. We firstly recall some well-known facts about free-abelian groups that, in particular, solve $\mathsf{WhP}_\mathsf{A}(\mathbb{Z}^m)$, $\mathsf{WhP}_\mathsf{M}(\mathbb{Z}^m)$ and $\mathsf{WhP}_\mathsf{E}(\mathbb{Z}^m)$. Here, for a vector $\mathbf{a} = (a_1, \ldots, a_m) \in \mathbb{Z}^m$, we write $\gcd(\mathbf{a})$ to denote the greatest common divisor of the a_i's (with the convention that $\gcd(\mathbf{0}) = 0$).

Lemma 2.5.1. *An element* $\mathbf{a} \in \mathbb{Z}^m$ *is primitive (member of a free-abelian basis of* $\mathbb{Z}^m$*) if and only if* $\gcd(\mathbf{a}) = 1$.

Proof. [$\Rightarrow$] By contrapositive, suppose that $d = \gcd(\mathbf{a}) \neq 1$. Then $\mathbf{a} = d \cdot \hat{\mathbf{a}}$, where $\gcd(\hat{\mathbf{a}}) = 1$. Thus, for every $\mathbf{a_2}, \ldots \mathbf{a_m} \in \mathbb{Z}^m$,

$$\det(\mathbf{a}, \mathbf{a_2}, \ldots, \mathbf{a_m}) \;=\; d \cdot \det(\hat{\mathbf{a}}, \mathbf{a_2}, \ldots, \mathbf{a_m}) \;\neq\; \pm 1\,,$$

and so $\{\mathbf{a}, \mathbf{a_2}, \ldots, \mathbf{a_m}\}$ is not a free-abelian basis of $\mathbb{Z}^m$.

[$\Leftarrow$] If $\gcd(\mathbf{a}) = 1$ we know (see for example [Art10, Section 14.4]) that there exists a matrix $\mathbf{P} \in \mathrm{GL}_m(\mathbb{Z})$ such that $\mathbf{aP} = (1, 0, \ldots, 0)$. That is, $\mathbf{a}$ is the first row of $\mathbf{P}^{-1}$, and therefore a primitive element of $\mathbb{Z}^m$. □

Lemma 2.5.2. *Let* $\mathbf{a} \in \mathbb{Z}^m$. *Then,*

(i) *An element* $\mathbf{y} \in \mathbb{Z}^n$ *is a homomorphic image of* $\mathbf{a}$ *if and only if* $\gcd(\mathbf{a}) \mid \gcd(\mathbf{y})$. *That is, the* $\mathbb{Z}^n$*-homomorphic orbit of* $\mathbf{a}$ *is:*

$$\{\, \mathbf{aM} \,:\, \mathbf{M} \in \mathcal{M}_{m\times n}(\mathbb{Z})\,\} = \left\{\, \mathbf{y} \in \mathbb{Z}^n \,:\, \gcd(\mathbf{a}) \mid \gcd(\mathbf{y}) \,\right\}.$$

(ii) *An element* $\mathbf{x} \in \mathbb{Z}^m$ *is an automorphic image of* $\mathbf{a}$ *if and only if* $\gcd(\mathbf{x}) = \gcd(\mathbf{a})$. *That is, the automorphic orbit of* $\mathbf{a}$ *is:*

$$\{\, \mathbf{aP} \,:\, \mathbf{P} \in \mathrm{GL}_m(\mathbb{Z})\,\} = \{\, \mathbf{x} \in \mathbb{Z}^m \,:\, \gcd(\mathbf{a}) = \gcd(\mathbf{x})\,\}.$$

(iii) *An element* $\mathbf{x} \in \mathbb{Z}^m$ *is a monomorphic image of* $\mathbf{a} \neq 0$ *if and only if* $\gcd(\mathbf{a}) \mid \gcd(\mathbf{x}) \neq 0$. *That is, the monomorphic orbit of* $\mathbf{a} \neq 0$ *is:*

$$\{\, \mathbf{aQ} \,:\, \mathbf{Q} \in \mathcal{M}_m(\mathbb{Z}) \text{ with } \det(\mathbf{Q}) \neq 0\,\} = \left\{\, \mathbf{x} \in \mathbb{Z}^m \,:\, \gcd(\mathbf{a}) \mid \gcd(\mathbf{x}) \,\right\} \setminus \{\,\mathbf{0}\,\},$$

Proof. The case $\mathbf{a} = \mathbf{0}$ holds trivially. So, suppose $\mathbf{a} \neq \mathbf{0}$, let $d = \gcd(\mathbf{a})$, and let $\mathbf{a} = d \cdot \hat{\mathbf{a}}$, where $\gcd(\tilde{\mathbf{a}}) = 1$.

(i) It is clear that $\mathbf{aM} = d \cdot \hat{\mathbf{a}}\mathbf{M}$, and so $d \mid \gcd(\mathbf{aM})$. Conversely, given $\mathbf{y} \in \mathbb{Z}^n$ such that $d \mid \gcd(\mathbf{y})$, let $\mathbf{y} = d \cdot \hat{\mathbf{y}}$. Since $\hat{\mathbf{a}}$ is primitive, there exist a matrix $\mathbf{M} \in \mathcal{M}_{m\times n}(\mathbb{Z})$ such that $\hat{\mathbf{a}}\mathbf{M} = \hat{\mathbf{y}}$, and hence $\mathbf{aM} = \mathbf{y}$.

(ii) If $\mathbf{x} = \mathbf{aP}$, then $\mathbf{a} = \mathbf{xP}^{-1}$, and the inclusion to the right follows from (i). The inclusion to the left is immediate from Lemma 2.5.1.

(iii) The inclusion to the right is proved in the same way as in (i). However, note that now, the extra assumptions that $\mathbf{a} \neq 0$, and $\mathbf{Q}$ is injective exclude $\mathbf{0}$ from the orbit. Conversely, given $\mathbf{x} \in \mathbb{Z}^n \setminus \{\mathbf{0}\}$ such that $d \mid \gcd(\mathbf{x})$, we can write $\mathbf{x} = c \cdot d \cdot \hat{\mathbf{x}}$, where $c \neq 0$, and $\gcd(\hat{\mathbf{x}}) = 1$. Since both $\hat{\mathbf{a}}$ and $\hat{\mathbf{x}}$ are primitive (by (ii)) there exists a matrix $\mathbf{P} \in \mathrm{GL}_m(\mathbb{Z})$ such that $\hat{\mathbf{x}} = \hat{\mathbf{a}}\mathbf{P}$. Then,

$$\mathbf{x} = c \cdot d \cdot \hat{\mathbf{x}} = c \cdot d \cdot \hat{\mathbf{a}}\mathbf{P} = c \cdot \mathbf{aP} = \mathbf{a}[c \cdot \mathbf{P}],$$

and it is enough to take $\mathbf{Q} = c\mathbf{P}$. □

Since the previous characterizations are fully algorithmic, and the corresponding search problems reduce to standard linear algebra over $\mathbb{Z}$, we already have the solvability of all three problems for finitely generated free-abelian groups.

Corollary 2.5.3. $\mathsf{WhP_A}(\mathbb{Z}^m)$, $\mathsf{WhP_M}(\mathbb{Z}^m)$, *and* $\mathsf{WhP_E}(\mathbb{Z}^m)$ *are solvable.* □

As expected, the same problems for the free group $\mathbb{F}_n$ are much more complicated. As mentioned above, the case of automorphisms was solved by Whitehead back in the 1930's. The case of endomorphisms can be solved by writing a system of equations over $\mathbb{F}_n$ (with unknowns being the images of a given free basis for $\mathbb{F}_n$), and then solving it by the powerful Makanin's algorithm. Finally, the case of monomorphisms was recently solved by Ciobanu and Houcine.

Theorem 2.5.4. *Let $\mathbb{F}_n$ denote the free group of rank n. Then,*

(i) $\mathsf{WhP}_{\mathsf{A}}(\mathbb{F}_n)$ *is solvable* (Whitehead, 1936, [Whi36]).

(ii) $\mathsf{WhP}_{\mathsf{M}}(\mathbb{F}_n)$ *is solvable* (Ciobanu and Houcine, 2010, [CH10]).

(iii) $\mathsf{WhP}_{\mathsf{E}}(\mathbb{F}_n)$ *is solvable* (Makanin, 1982, [Mak82]). □

Our goal is to make use of the results above, together with our description of the different kinds of endomorphisms of $\mathbb{Z}^m \times \mathbb{F}_n$, in order to prove the solvability of all three problem in this new family.

Theorem 2.5.5. *The Whitehead problems for automorphisms, monomorphisms, and endomorphisms, are solvable for finitely generated free-abelian times free groups, namely:*

(i) $\mathsf{WhP}_{\mathsf{A}}(\mathbb{Z}^m \times \mathbb{F}_n)$ *is solvable.*

(ii) $\mathsf{WhP}_{\mathsf{M}}(\mathbb{Z}^m \times \mathbb{F}_n)$ *is solvable.*

(iii) $\mathsf{WhP}_{\mathsf{E}}(\mathbb{Z}^m \times \mathbb{F}_n)$ *is solvable.*

Proof. We are given two elements $t^{\mathbf{a}}\,u,\ t^{\mathbf{b}}\,v \in G = \mathbb{Z}^m \times \mathbb{F}_n$, and have to decide whether there exists an automorphism (resp. monomorphism, endomorphism) of G sending one to the other. And in the affirmative case, find one of them. For convenience, we shall prove (ii), (i), and (iii) in this order.

(ii) Since all monomorphisms of G are of type I, we have to decide whether there exist a monomorphism ϕ of $\mathbb{F}_n$, and matrices $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$ and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$, with $\det \mathbf{Q} \neq 0$, such that $(t^{\mathbf{a}}u)\Psi_{\phi,\mathbf{Q},\mathbf{P}} = t^{\mathbf{b}}\,v$. Separating the free and free-abelian parts, we get two independent problems:

$$\left.\begin{aligned} u\phi &= v \\ \mathbf{aQ} + \mathbf{uP} &= \mathbf{b} \end{aligned}\right\} \tag{2.24}$$

On one hand, we can use Theorem 2.5.4.(ii) to decide whether there exists a monomorphism ϕ of $\mathbb{F}_n$ such that $u\phi = v$. If not then our problem has no solution either, and we are done; otherwise, $\mathsf{WhP}_{\mathsf{M}}(\mathbb{F}_n)$ gives us such a ϕ.

On the other hand, we need to know whether there exist matrices $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$ and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$, with $\det \mathbf{Q} \neq 0$ and such that $\mathbf{aQ} + \mathbf{uP} = \mathbf{b}$, where $\mathbf{u} \in \mathbb{Z}^n$

is the abelianization of $u \in \mathbb{F}_n$ (given from the beginning). If $\mathbf{a} = \mathbf{0}$ or $\mathbf{u} = \mathbf{0}$, this is already solved in Lemma 2.5.2.(i) or Lemma 2.5.2.(iii) Otherwise, write $0 \neq \alpha = \gcd(\mathbf{a})$ and $0 \neq \mu = \gcd(\mathbf{u})$; and, according to Lemma 2.5.2, we have to decide whether there exist $\mathbf{a}' \in \mathbb{Z}^m$ and $\mathbf{u}' \in \mathbb{Z}^m$, with $\mathbf{a}' \neq \mathbf{0}$, $\alpha \mid \gcd(\mathbf{a}')$, and $\mu \mid \gcd(\mathbf{u}')$, such that $\mathbf{a}' + \mathbf{u}' = \mathbf{b}$. Writing $\mathbf{a}' = \alpha\,\mathbf{x}$ and $\mathbf{u}' = \mu\,\mathbf{y}$, the problem reduces to test whether the following linear system of equations

$$\left.\begin{array}{cccc} \alpha\,x_1 & + \;\mu\,y_1 & = & b_1 \\ & \vdots & \vdots & \\ \alpha\,x_m & + \;\mu\,y_m & = & b_m \end{array}\right\} \tag{2.25}$$

has any integral solution $x_1,\ldots,x_m,y_1,\ldots,y_m \in \mathbb{Z}$ such that $(x_1,\ldots,x_m) \neq \mathbf{0}$. A necessary and sufficient condition for the system (2.25) to have a solution is $\gcd(\alpha,\mu) \mid b_j$, for every $j \in [m]$. And note that, if (x_1,y_1) is a solution to the first equation, then $(x_1+\mu, y_1-\alpha)$ is another one; since $\mu \neq 0$, the condition $(x_1,\ldots,x_m) \neq \mathbf{0}$ is then superfluous. Therefore, the answer is affirmative if and only if $\gcd(\alpha,\mu) \mid b_j$, for every $j \in [m]$; and, in this case, we can easily reconstruct a monomorphism Ψ of G such that $(t^{\mathbf{a}}\,u)\Psi = t^{\mathbf{b}}\,v$.

(i) The argument for automorphisms is completely parallel to the previous discussion replacing the conditions ϕ monomorphism and $\det \mathbf{Q} \neq 0$, to ϕ automorphism and $\det \mathbf{Q} = \pm 1$. We manage the first change by using Theorem 2.5.4.(i) instead of Theorem 2.5.4.(ii) The second change forces us to look for solutions of the linear system (2.25) with the extra requirement $\gcd(\mathbf{x}) = 1$ (because now $\gcd(\mathbf{a}')$ should be equal to, and not just a multiple, of α).

So, if any of the conditions $\gcd(\alpha,\mu) \mid b_j$ fails, the answer is negative and we are done. Otherwise, write $\rho = \gcd(\alpha,\mu)$, $\alpha = \rho\alpha'$ and $\mu = \rho\mu'$, and the general solution for the j-th equation in (2.25) is

$$(x_j, y_j) \;=\; (x_j^0, y_j^0) + \lambda_j(\mu', -\alpha'), \quad \lambda_j \in \mathbb{Z},$$

where (x_j^0, y_j^0) is a particular solution, which can be easily computed. Thus, it only remains to decide whether there exist $\lambda_1,\ldots,\lambda_m \in \mathbb{Z}$ such that

$$\gcd(x_1^0 + \lambda_1\mu', \ldots, x_m^0 + \lambda_m\mu') \;=\; 1. \tag{2.26}$$

We claim that this happens if and only if

$$\gcd(x_1^0, \ldots, x_m^0, \mu') \;=\; 1, \tag{2.27}$$

which is clearly a decidable condition.

Reorganizing a Bezout identity for (2.26) we can obtain a Bezout identity for (2.27). Hence (2.26) implies (2.27). For the converse, assume the integers $x_1^0,\ldots,x_m^0,\mu'$

are coprime, and we can fulfill equation (2.26) by taking $\lambda_1 = \cdots = \lambda_{m-1} = 0$ and λ_m equal to the product of the primes dividing $x_1^0, \ldots, x_{m-1}^0$ but not x_m^0 (take $\lambda_m = 1$ if there is no such prime). Indeed, let us see that any prime p dividing $x_1^0, \ldots, x_{m-1}^0$ is not a divisor of $x_m^0 + \lambda_m \mu'$. If p divides x_m^0, then p does not divide neither μ' nor λ_m and therefore $x_m^0 + \lambda_m \mu'$ either. If p does not divide x_m^0, then p divides λ_m by construction, hence p does not divide $x_m^0 + \lambda_m \mu'$. This completes the proof of the claim, and of the theorem for automorphisms.

(iii) In our discussion now, we should take into account endomorphisms of types I and II.

Again, the argument to decide whether there exists an endomorphism of type I sending $t^{\mathbf{a}} u$ to $t^{\mathbf{b}} v$, is completely parallel to the above proof (ii), replacing the condition ϕ monomorphism by ϕ endomorphism, and deleting the condition $\det \mathbf{Q} \neq 0$ (allowing here an arbitrary matrix $\mathbf{Q}$). We manage the first change by using Theorem 2.5.4.(iii) instead of Theorem 2.5.4.(ii) The second change simply leads us to solve the system (2.25) with no extra condition on the variables; so, the answer is affirmative if and only if $\gcd(\alpha, \mu) \mid b_j$, for every $j \in [m]$.

It remains to consider endomorphisms of type II, $\Psi_{z,\mathbf{l},\mathbf{h},\mathbf{Q},\mathbf{P}}$. So, given our elements $t^{\mathbf{a}} u$ and $t^{\mathbf{b}} v$, and separating the free and free-abelian parts, we have to decide whether there exist $z \in \mathbb{F}_n$, $\mathbf{l} \in \mathbb{Z}^m$, $\mathbf{h} \in \mathbb{Z}^n$, $\mathbf{Q} \in \mathcal{M}_m(\mathbb{Z})$, and $\mathbf{P} \in \mathcal{M}_{n \times m}(\mathbb{Z})$ such that

$$\left.\begin{aligned} z^{\mathbf{a}\mathbf{l}^T + \mathbf{u}\mathbf{h}^T} &= v \\ \mathbf{a}\mathbf{Q} + \mathbf{u}\mathbf{P} &= \mathbf{b} \end{aligned}\right\} \tag{2.28}$$

(note that we can ignore the condition $\mathbf{l} \neq \mathbf{0}$ because if $\mathbf{l} = \mathbf{0}$ then the endomorphism becomes of type I as well, and this case is already considered before). Again the two equations are independent. About the free part, note that the integers $\mathbf{a}\mathbf{l}^T + \mathbf{u}\mathbf{h}^T$ with $\mathbf{l} \in \mathbb{Z}^m$ and $\mathbf{h} \in \mathbb{Z}^n$ are precisely the multiples of $d = \gcd(\mathbf{a}, \mathbf{u})$; so, it has a solution if and only if v is a d-th power in $\mathbb{F}_n$, a very easy condition to check. And about the second equation, it is exactly the same as when considering endomorphisms of type I, so its solvability is already discussed. □

2.6 Twisted conjugacy problem

Recall that given G a group, and α an automorphism of G, we say that two elements $g, h \in G$ are α(twisted)-conjugate, denoted by $g \sim_\alpha h$, if there exists an element $k \in G$ such that $(k\alpha)^{-1} g k = h$. Then, we also say that k is an α(-twisted) conjugator of g into h. It is routine to check that α-conjugacy is an equivalence relation. Note also, that id_G-conjugacy corresponds to standard conjugacy in G.

Recently , O. Bogopolski, A. Martino, and E. Ventura rediscovered twisted conjugacy as an important ingredient relating conjugacy with orbit decidability in certain extensions of groups (see Theorem 2.6.9). To this end, they consider in [BMV10] the algorithmic problem below. As usual we will assume that a finite presentation for G is given, and that the input of the problem is given in terms of this presentation.

Twisted conjugacy problem, $\mathsf{TCP}(G)$**.** *Given two elements* $u, v \in G$*, and an automorphism* $\alpha \in \operatorname{Aut} G$*, decide whether there exists an element* $w \in G$ *such that* $(w\alpha)^{-1}uw = v$*; and if so, find one such* α*-conjugate* w*.*

The problem obtained by restricting the family $\mathcal{F} \subseteq \operatorname{Aut} G$ where we can pick the automorphism α, is called the $\mathcal{F}$-(twisted) conjugacy problem for G, denoted by $\mathsf{TCP}_{\mathcal{F}}(G)$. Accordingly, if $\mathcal{F} = \{\alpha\}$ (i.e., if the automorphism α is fixed) we have the α-(twisted) conjugacy problem for G, denoted by $\mathsf{TCP}_{\alpha}(G)$.

Remark 2.6.1. Note that $\mathsf{TCP}_{id}(G) = \mathsf{CP}(G)$ but, in general, $\mathsf{TCP}(G)$ is a strictly stronger algorithmic problem that standard $\mathsf{CP}(G)$ (see [BMV10, Corollary 4.9]) for an example of a group with solvable CP but unsolvable TCP.

Our goal in this section is to prove the solvability of $\mathsf{TCP}(\mathbb{Z}^m \times \mathbb{F}_n)$, and derive some consequences using Theorem 2.6.9.

We emphasize that TCP is well known to be solvable for each of the factors in $\mathbb{Z}^m \times \mathbb{F}_n$. Concretely, the solvability of $\mathsf{TCP}(\mathbb{Z}^m)$ reduces to solving a system of linear diophantine equations, whereas that of $\mathsf{TCP}(\mathbb{F}_n)$ is proved in [Bog+06] as a prelude to that of $\mathsf{CP}(\mathbb{F}_n \rtimes \mathbb{Z})$.

Theorem 2.6.2 (Bogopolski, Martino, Maslakova, and Ventura, 2006, [Bog+06])**.** $\mathsf{TCP}(\mathbb{F}_n)$ *is solvable.* □

Note, however, that it is not enough to independently consider the solvability of TCP for the factors ($\mathbb{F}_n$ and $\mathbb{Z}^m$) to deduce the solvability of $\mathsf{TCP}(\mathbb{Z}^m \times \mathbb{F}_n)$. So, calling $\mathcal{A} = \{(\phi, \mathbf{I_m}, \mathbf{0}) : \phi \in \operatorname{Aut}(\mathbb{F}_n)\}$, and $\mathcal{B} = \{(\mathrm{id}_{\mathbb{F}_n}, \mathbf{Q}, \mathbf{0}) : \mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})\}$, it is clear from the characterization in Figure 7.4 that, taken separately, the solvabilities of $\mathsf{TCP}(\mathbb{F}_n)$, and $\mathsf{TCP}(\mathbb{Z}^m)$ translate into those of $\mathsf{TCP}_{\mathcal{A}}(\mathbb{Z}^m \times \mathbb{F}_n)$ and $\mathsf{TCP}_{\mathcal{B}}(\mathbb{Z}^m \times \mathbb{F}_n)$, respectively; subcases that obviously do not cover the whole automorphism group $\operatorname{Aut}(\mathbb{Z}^m \times \mathbb{F}_n)$. In fact, as you can deduce from Figure 7.4, they not even constitute a family of generators. We will need, therefore, a less coarse approach to achieve our result.

Remark 2.6.3. Note that any TCP search problem is guaranteed to be algorithmically solvable (by a brute force argument) once the decision problem is known to answer YES. Namely, given two elements $g, h \in G$ we can always recursively enumerate the elements in $k \in G$ checking whether the twisted-conjugacy condition $(k\alpha)^{-1}gk = h$ holds. Since such an element is known to exist, this search procedure will certainly terminate.

The following characterization of the twisted conjugators between two elements is straightforward.

Lemma 2.6.4. *Let* $g, h \in G$, $\alpha \in \operatorname{Aut} G$, *and* k_0 *an* α*-conjugator of* g *into* h*. Then, for all* $k \in G$,

$$(k\alpha)^{-1}gk = h \quad \Leftrightarrow \quad k \in \operatorname{Fix}(\alpha\gamma_g)\, k_0\,,$$

where $\gamma_g \colon x \mapsto g^{-1}xg$ *is the standard (right) conjugation in* G.

In particular, the set of α*-conjugators of one element into another is either empty, or a coset of* $\operatorname{Fix}(\alpha\gamma_g)$.

Proof. Since, by assumption, $(k_0\alpha)^{-1}gk_0 = h$, then

$$\begin{aligned} (k^{-1})\alpha\, g\, k = h \;&\Leftrightarrow\; (k_0^{-1})\alpha\, g\, k_0 = (k^{-1})\alpha\, g\, k \\ &\Leftrightarrow\; g^{-1}\,(kk_0^{-1})\alpha\, g = kk_0^{-1} \\ &\Leftrightarrow\; kk_0^{-1} \in \operatorname{Fix}(\alpha\gamma_g) \\ &\Leftrightarrow\; k \in \operatorname{Fix}(\alpha\gamma_g)\, k_0\,. \end{aligned}$$ □

The lemma below is, as we will see in the proof of Theorem 2.6.6, essentially a compact version of the decision problem obtained from $\mathsf{TCP}(\mathbb{Z}^m \times \mathbb{F}_n)$ after using our description for the involved automorphism, and separating the free and free abelian parts.

Lemma 2.6.5. *The problem consisting in, given* $u, v \in \mathbb{F}_n$, $\alpha \in \operatorname{Aut}(\mathbb{F}_n)$, $\mathbf{P} \in\in \mathcal{M}_{n\times m}(\mathbb{Z})$, $\mathbf{d} \in \mathbb{Z}^m$, *and* $L \leqslant \mathbb{Z}^m$, *whether there exist an* α*-conjugator* w *of* u *into* v, *such that* $\mathbf{wP} + \mathbf{d} \in H$ *is algorithmically solvable.*

That is, we can algorithmically decide whether

$$\exists w \in \mathbb{F}_n \text{ such that } \begin{cases} (w\alpha)^{-1}uw = v\,, \text{ and} \\ \mathbf{wP} + \mathbf{d} \in L\,. \end{cases} \tag{2.29}$$

Proof. Using Theorem 2.6.2, we can algorithmically decide the existence of α-conjugators and, in affirmative case, find one such conjugator (that we will call w_0). If the algorith to decide $\mathsf{TCP}(\mathbb{F}_n)$ answers NO, then we also answer NO, and we are done. Otherwise, according Lemma 2.6.4, our problem becomes that of deciding whether:

$$\exists w \in \mathbb{F}_n \text{ such that } \begin{cases} w \in \operatorname{Fix}(\alpha\gamma_u)w_0\,, \text{ and} \\ \mathbf{wP} + \mathbf{d} \in L\,. \end{cases}$$

Now, since a basis for the fixed point subgroup of any automorphism $\alpha \in \operatorname{Aut}(\mathbb{F}_n)$ is algorithmically computable [BM15], we can find a generating set for $\operatorname{Fix}(\alpha\gamma_u)$

and, after taking the image by the ambient abelianization $\mathbb{F}_n \twoheadrightarrow \mathbb{Z}^n$, compute a basis for the subgroup $M = (\mathrm{Fix}(\alpha\gamma_u))^{ab} \leqslant \mathbb{Z}^n$. In this way we reduce our initial problem to deciding whether:

$$\exists \mathbf{w} \in \mathbb{Z}^n \text{ such that } \begin{cases} \mathbf{w} \in M + \mathbf{w_0}\,, \text{ and} \\ \mathbf{wP} + \mathbf{d} \in L\,, \end{cases}$$

that is, to solving a linear system of diophantine equations, which is well known to be algorithmically decidable. □

Theorem 2.6.6. $\mathsf{TCP}(\mathbb{Z}^m \times \mathbb{F}_n)$ *is solvable.*

Proof. Given elements $t^{\mathbf{a}}\, u, t^{\mathbf{b}}\, v \in \mathbb{Z}^m \times \mathbb{F}_n$, and an automorphism $\Psi \in \mathrm{Aut}(\mathbb{Z}^m \times \mathbb{F}_n)$, we want to decide whether there exists an element $t^{\mathbf{c}}\, w \in \mathbb{Z}^m \times \mathbb{F}_n$ such that:

$$t^{\mathbf{c}}\, w)^{-1}\Psi\; t^{\mathbf{a}}\, u\; t^{\mathbf{c}}\, w \;=\; t^{\mathbf{b}}\, v\,.$$

Since every automorphism of $\mathbb{Z}^m \times \mathbb{F}_n$ is of the form $\Psi\colon t^{\mathbf{a}}\, u \mapsto t^{\mathbf{aQ}+\mathbf{uP}}\, u\alpha$, where $\alpha \in (\mathrm{Aut}\, \mathbb{F}_n)$, $\mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})$, and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$, we obtain the following equivalent equation:

$$t^{\,\mathbf{a}\,+\,\mathbf{c}(\mathbf{I_m}-\mathbf{Q})\,-\,\mathbf{wP}}\, (w\alpha)^{-1} u w \;=\; t^{\mathbf{b}}\, v\,, \tag{2.30}$$

where, $\mathbf{w}$ denotes the abelianization of $w \in \mathbb{F}_n$.

Now, separating the free, and free-abelian parts in (2.30), $\mathsf{TCP}(\mathbb{Z}^m \times \mathbb{F}_n)$ is reduced to the following problem: given $\mathbf{a}, \mathbf{b} \in \mathbf{Z}^m$, $u, v \in \mathbb{F}_n$, $\alpha \in \mathrm{Aut}(\mathbb{F}_n)$, $\mathbf{Q} \in \mathrm{GL}_m(\mathbb{Z})$, and $\mathbf{P} \in \mathcal{M}_{n\times m}(\mathbb{Z})$, decide whether there exist $\mathbf{c} \in \mathbb{Z}^m$, and $w \in \mathbb{F}_n$ such that:

$$\begin{cases} (w\phi)^{-1} u w \;=\; v \\ \mathbf{wP} + \mathbf{b} - \mathbf{a} \;=\; \mathbf{c}(\mathbf{I_m} - \mathbf{Q})\,, \end{cases} \tag{2.31}$$

which is a particular case of Lemma 2.6.5, where $\mathbf{d} = \mathbf{b} - \mathbf{a}$, and $L = \mathrm{im}(\mathbf{I_m} - \mathbf{Q})$. □

Let us now recall the aforementioned relation involving TCP, CP, and orbit decidability (OP) in order to derive some consequences from Theorem 2.6.6. To this end, we introduce below some new terminology.

Orbit problems, $\mathsf{OP}_{\mathcal{A}}(G)$**.** *For a subset* $\mathcal{A} \leqslant \mathrm{End}(G)$*, decide, given* $g, h \in G$*, whether there exist a homomorphism* $\alpha \in \mathcal{A}$*, such that* $g\alpha = h$*; and if so, find one such endomorphism.*

When $\mathsf{OP}_{\mathcal{A}}(G)$ *is algorithmically decidable, we equivalently say that* $\mathcal{A}$ *is* orbit decidable *(*OD*, for short).*

Remark 2.6.7. Recall that if $\mathcal{A}$ is recursively enumerable (e.g. given by a images of a finite set of generators), then the orbit search problem is solvable by brute force.

Remark 2.6.8. Note that many classical algorithmic problems match the previous pattern. For example, the conjugacy problem of a group G is nothing more that the orbit decidability of the subgroup of its inner automorphisms, i.e., $\mathsf{CP}(G) = \mathsf{OD}(\operatorname{Inn} G)$; note also the Whitehead problems are precisely the orbit problems of the corresponding subgroups.

Note also that, given a short exact sequence of groups

$$1 \longrightarrow F \xrightarrow{\iota} G \xrightarrow{\rho} H \longrightarrow 1\,,$$

since $F\iota$ is normal in G, then for every $g \in G$ the restriction to $F\iota$ of the conjugation γ_g in G is an automorphism, $x\iota \mapsto g^{-1}x\iota\, g$ of $F\iota$ that, through ι, induces an automorphism α_g a F; concretely $\alpha_g \coloneqq \iota\gamma_{g|_{F\iota}}\iota^{-1}$, i.e.,

$$\alpha_g \colon x \overset{\iota}{\longmapsto} x\iota \overset{\gamma_{g|_{F\iota}}}{\longmapsto} g^{-1}(x\iota)\, g \overset{\iota^{-1}}{\longmapsto} (g^{-1}(x\iota)\, g)\iota^{-1}.$$

Identifying, as usual, x and $x\iota$, we have that for every $g \in G$, $\alpha_g : x \mapsto g^{-1}xg$ is an automorphism (not necessarily inner) of F. It is clear that the set of such automorphisms,

$$\mathcal{A}_G = \{\alpha_g \,:\, g \in G\},$$

is a subgroup of $\operatorname{Aut}(F)$ containing $\operatorname{Inn}(F)$. We call it the *action subgroup* of the given short exact sequence.

Assuming certain hypothesis on the short exact sequence, and the groups involved, the theorem below shows that the solvability of the conjugacy problem for G is equivalent to the orbit decidability of the action subgroup $\mathcal{A}_G \leqslant \operatorname{Aut} F$.

Theorem 2.6.9 (Bogopolski, Martino, and Ventura, 2010, [BMV10])**.** *Let*

$$1 \longrightarrow F \longrightarrow G \longrightarrow H \longrightarrow 1 \tag{2.32}$$

be an algorithmic short exact sequence of groups such that:

(i) $\mathsf{TCP}(F)$ *is solvable.*

(ii) $\mathsf{CP}(H)$ *is solvable.*

(iii) *for every* $1 \neq h \in H$, *the subgroup* $\langle h \rangle$ *has finite index in its centralizer* $C_H(h)$, *and there is an algorithm which computes a finite set of coset representatives,* $z_{h,1}, \ldots, z_{h,t_h} \in H$, *i.e.,*

$$C_H(h) = \langle h \rangle z_{h,1} \sqcup \cdots \sqcup \langle h \rangle z_{h,t_h}\,.$$

Then,

$$\mathsf{CP}(G) \textit{ is decidable} \quad \Leftrightarrow \quad \mathcal{A}_G \textit{ is orbit decidable},$$

where $\mathcal{A}_G$ is the action subgroup of (2.32). ☐

Remark 2.6.10. Note that Theorem 2.6.9 provides a kind of machinery to prove the conjugacy problem for group extensions: ar far as you are able to prove conditions (i),(ii),(iii) in Theorem 2.6.9, and orbit decidability for the action subgroup, you automatically have the solvability of the conjugacy problem for the central subgroup of the extension.

Applying these considerations to some of our previous results we immediately obtain the following corollary.

Corollary 2.6.11. *If $\Phi_1, \ldots, \Phi_r$ is a generating set for* $\mathrm{Aut}(F_n \times \mathbb{Z}^m)$*, then the group* $(\mathbb{Z}^m \times \mathbb{F}_n) \rtimes_{\Phi_1,\ldots,\Phi_r} \mathbb{F}_r$ *has solvable conjugacy problem.*

Proof. It is enough to apply Theorem 2.6.9 on the short exact sequence:

$$1 \longrightarrow \mathbb{Z}^m \times \mathbb{F}_n \longrightarrow (\mathbb{Z}^m \times \mathbb{F}_n) \rtimes_{\Phi_1,\ldots,\Phi_r} \mathbb{F}_r \longrightarrow \mathbb{F}_r \longrightarrow 1\,.$$

Note that in this case the action subgroup is $\mathrm{Aut}(F_n \times \mathbb{Z}^m)$ all the requisites in Remark 2.6.10 are satisfied. Namely:

- $\mathsf{TCP}(\mathbb{Z}^m \times \mathbb{F}_n)$ is solvable: Theorem 2.6.6.
- el $\mathsf{CP}(\mathbb{F}_r)$ is solvable: this is a well-known fact, see [LS01, Proposition I.2.14].
- for all $1 \neq w \in \mathbb{F}_r$, $\langle C_{F_r}(w) = \langle \mathrm{root}(w) \rangle$, and condition (iii) in Theorem 2.6.9 holds trivially.
- $\mathrm{Aut}(F_n \times \mathbb{Z}^m)$ is orbit decidable: Theorem 2.5.5.(i).

We conclude, therefore, that $\mathsf{CP}((F_n \times \mathbb{Z}^m) \rtimes_{\alpha_1,\ldots,\alpha_r} F_r)$ is solvable, as claimed. ☐

Part II

Semidirect products

Group extensions and semidirect products

3

Semidirect products (see Section 3.2) are natural generalizations of direct products, and thus a natural target in order to extend the results in Part I. In turn, they constitute a particular case of much more general constructions called group extensions. It is often convenient to see semidirect products within this broader framework; so we introduce semidirect products from both points of view.

3.1 Group extensions

Definition 3.1.1. A sequence of groups and homomorphisms

$$\cdots \xrightarrow{\phi_{i-1}} G_i \xrightarrow{\phi_i} G_{i+1} \xrightarrow{\phi_{i+1}} \cdots$$

is said to be an *exact sequence* if the image of every homomorphism is the kernel of the next one; i.e., if for all i, $\operatorname{im}(\phi_i) = \ker(\phi_{i+1})$. A *short exact sequence* is an exact sequence of the form: $1 \to G_1 \to G_2 \to G_3 \to 1$.

Definition 3.1.2. Let N and Q be arbitrary groups. A *group extension of* N *by* Q is a group G having a normal subgroup (isomorphic to) N, such that the quotient G/N is isomorphic to Q.

Note that we can compactly summarize the previous definition saying that a group G is an extension of N by Q if it admits a short exact sequence of the form:

$$1 \longrightarrow N \longrightarrow \iota G \longrightarrow \pi Q \longrightarrow 1\,. \tag{3.1}$$

Indeed, from the definition of exact sequence:

- The homomorphism ι is injective, and we usually identify the group N with its (isomorphic) image $N\iota$, which is normal in G.
- The homomorphism π is onto; ie $G\pi = Q$.
- The quotient G/N is isomorphic to Q.

Then, we say that N is a *base group* for G, with *quotient group* Q (and vice-versa).

In particular, a group extension G of N (in the sense of definition 3.1.2) is always a group extension of N in the broad sense; i.e., the base group N is always (isomorphic to) a subgroup of G. However, it is *not true* that the quotient group Q is always (isomorphic to) a subgroup of G. Indeed, we will see that the case when this happens corresponds exactly to semidirect products.

Definition 3.1.3. A short exact sequence like (3.1) is said to split if there exists a homomorphism $\sigma\colon Q \to G$ such that $\sigma\pi = \mathrm{id}_Q$. Then, we say that σ is a *section* (or *split*) of π, and we write

$$1 \longrightarrow N \xrightarrow{\iota} G \underset{\sigma}{\overset{\pi}{\rightleftarrows}} Q \longrightarrow 1 \tag{3.2}$$

Remark 3.1.4. Note that a section is always injective. Hence, in a split extension, the quotient group Q is always (isomorphic to) a subgroup of G.

3.2 Semidirect products

We summarize in this section the standard definition(s) of semidirect product, its elementary properties, and its relation with group extensions (see for example [Rob96] for details).

Throughout this section N and Q will be arbitrary groups, and

$$\begin{aligned} \alpha\colon Q &\to \mathrm{Aut}(N) \\ q &\mapsto \alpha_q \end{aligned} \tag{3.3}$$

a group homomorphism (with autos acting on the right).

Definition 3.2.1. The *(external) semidirect product* of N by Q with action α is the group defined on the cartesian product $Q \times N$ by the operation:

$$(q_1, n_1) \cdot (q_2, n_2) = (q_1 \cdot q_2, (n_1)\alpha_{q_2} \cdot n_2) \ . \tag{3.4}$$

We will interchangeably denote the semidirect product of N by Q with action α by $Q \ltimes_\alpha N$ or $N \rtimes_\alpha Q$.

We will write $G = Q \ltimes N$ (or or $G = N \rtimes Q$) to abbreviate that there exists a homomorphism $\alpha\colon Q \to \mathrm{Aut}(N)$ such that $G = Q \ltimes_\alpha N$.

Remark 3.2.2. It is straightforward to check that (3.4) defines an associative law, with neutral element $(1_Q, 1_N)$, and that every element $(q, n) \in Q \ltimes_\alpha N$ has an inverse $(q, n)^{-1} = (q^{-1}, (n^{-1})\alpha_q^{-1})$.

Remark 3.2.3. Note that if $\alpha\colon q \mapsto \mathrm{id}_N$, then $Q \ltimes_\alpha N = Q \times N$.

The results stated hereinafter in this section, whose proof we omit, are straightforward and absolutely standard (see, for example [Rob96]).

Proposition 3.2.4. *Let* $G = Q \ltimes N$. *Then:*

1. *The map* $\iota_N \colon n \mapsto (1_q, n)$ *is a monomorphism* $N \rightarrowtail Q \ltimes N$.

   (Hereinafter, we identify the group N *with* $N \times 1_Q$*.)*

2. *The map* $\pi_Q \colon (q, n) \mapsto q$ *is an epimorphism* $Q \ltimes N \twoheadrightarrow Q$ *with kernel* $N \times 1_Q \simeq N$.

3. *The map* $\mathrm{id}_Q \colon q \mapsto (q, 1_N)$ *is a section of* π_Q.

   (Hereinafter, we identify the group Q *with* $Q \times 1_N$*.)*

That is, $G = Q \ltimes_\alpha N$ *admits the short exact sequence*

$$1 \longrightarrow N \xrightarrow{\iota_N} G \xrightarrow{\pi_Q} Q \longrightarrow 1\,, \tag{3.5}$$

which splits via id_Q. □

Corollary 3.2.5. *Let* $G = Q \ltimes N$. *Then:*

[SD1] *The group* N *is a normal subgroup of* $Q \ltimes N$, *and the quotient* $(Q \ltimes N)/N$ *is isomorphic to* Q.

[SD2] *The group* Q *is a subgroup of* $Q \ltimes N$.

[SD3] $G = QN$, *and* $Q \cap N = 1$.

(i.e., every element in G *can be written in a unique way as a product of an element in* Q*, and an element in* N*.)* □

Indeed, it turns out that the three properties in Corollary 3.2.5 — or the short exact sequence (3.5) — fully characterize semidirect products.

Theorem 3.2.6. *Let* G *be a group. Then, the following statements are equivalent:*

1. *The group* G *admits a split short exact sequence (say* $1 \to N \to G \to Q \to 1$*).*

2. *There exist subgroups* $Q, N \leqslant G$ *satisfying the properties [SD1-SD3].*

3. *The group* G *is isomorphic to the semidirect product* $Q \ltimes_\alpha N$, *with action* $\alpha \colon q \mapsto \gamma_{q|N}\ (n \mapsto q^{-1}nq)$. □

Remark 3.2.7. Note that from property [SD3], every element $g \in G = Q \ltimes_\alpha N$ admits a *normal form*:

$$g = qn, \ \text{where } q \in Q, \ \text{and } n \in N\,; \tag{3.6}$$

and then, from (3.4), we have the following *multiplication rules* for he semidirect product: for every $q \in Q$, and every $n \in N$,

$$n \cdot q = q \cdot (n)\alpha_q \quad \text{and} \quad q \cdot n = (n)\alpha_q^{-1} \cdot q\,. \tag{3.7}$$

So, we have the following mnemonic criteria for a semidirect product $Q \ltimes_\alpha N$:

1. every element $q \in Q$ "jumps to the left (resp. right)" of an element $n \in N$ at the price of applying α_q (resp. α_q^{-1}) to the jumped element n.

2. every element $n \in N$ "jumps to the right (resp. left)" of an element $q \in Q$ at the price of applying α_q (resp. α_q^{-1}) to the jumping element n.

Observe that, in particular, we have that $1_G = 1_Q 1_N$, and for every $g = qn$ in $Q \ltimes_\alpha N$,

$$g^{-1} = (qn)^{-1} = n^{-1}q^{-1} = q^{-1}(n^{-1})\alpha_q^{-1}. \tag{3.8}$$

in concordance with Remark 3.2.2.

Corollary 3.2.8. *Let* $Q = \langle X \mid R \rangle$, $N = \langle Y \mid S \rangle$ *be arbitrary groups, and let* $\alpha \colon Q \to \mathrm{Aut}(N)$ *be a homomorphism of groups. Then, the semidirect product* $Q \ltimes_\alpha N$ *admits the presentation*

$$\big\langle X, Y \,\big|\, R, S, x_i^{-1} y_j x_i = (y_j)\alpha_{x_i} \ \forall x_i \in X, \forall y_j \in Y \big\rangle, \tag{3.9}$$

which is usually abbreviated $\langle Q, N \mid x_i^{-1} y_j x_i = (y_j)\alpha_{x_i} \ \forall x_i \in X, \forall y_j \in Y \rangle$, *and called a* standard presentation *for the semidirect product* $Q \ltimes_\alpha N$ *with respect to the given presentations.* □

3.2.1 Abelianization of semidirect products

One straightforward consequence of the form of standard presentations is the following description for the abelianization of a semidirect product.

Lemma 3.2.9. *Let* $Q = \langle X \mid R \rangle$ *and* $N = \langle Y \mid S \rangle$ *be arbitrary groups, and let* $\alpha \colon Q \to \mathrm{Aut}(N)$ *be a homomorphism of groups. Then, the abelianization of the semidirect product* $Q \ltimes_\alpha N$ *is*

$$(Q \ltimes_\alpha N)^{ab} \simeq Q^{ab} \oplus \frac{N^{ab}}{\langle \mathrm{im}(\alpha_x^{ab} - \mathrm{id}),\ x \in X \rangle}. \tag{3.10}$$

Moreover, if Q, N *are given by finite presentations, say with* $\#X = n$, *and* $\#Y = m$; *then the abelianization is computable. Namely:*

$$(Q \ltimes_\alpha N)^{ab} \simeq \frac{\mathbb{Z}^n}{\langle R^{ab} \rangle} \oplus \frac{\mathbb{Z}^m}{\langle S^{ab} \rangle + \sum_{i=1}^{n} \mathrm{im}(\mathbf{A_i} - \mathbf{I}m)}, \tag{3.11}$$

where, $\mathbf{A_i} := (\alpha_{x_i})^{ab} \colon \mathbb{Z}^m \to \mathbb{Z}^m$, *for* $i = 1, \ldots, n$.

Proof. Recall that $Q = \langle X \mid R \rangle$, and $N = \langle Y \mid S \rangle$. After adding the commutators between generators as relators in the standard presentation for $Q \ltimes_\alpha N$ we obtain:

$$\begin{aligned}(Q \ltimes_\alpha N)^{ab} &= \left\langle X, Y \;\middle|\; \begin{smallmatrix} R,S, & \\ x^{-1}yx = (y)\alpha_x, & (x\in X, y\in Y) \\ [x,x'],[y,y'],[x,y] & (x,x'\in X, y,y'\in Y) \end{smallmatrix} \right\rangle \\ &\simeq \left\langle X \;\middle|\; R,\ [x,x']\ \ (x,x' \in X) \right\rangle \times \left\langle Y \;\middle|\; \begin{smallmatrix} S,[y,y'] & (y,y'\in Y) \\ y=(y)\alpha_x & (x\in X, y\in Y) \end{smallmatrix} \right\rangle \\ &\simeq Q^{ab} \oplus \frac{N^{ab}}{\langle \operatorname{im}(\alpha_x^{ab} - \operatorname{id}),\ x \in X \rangle},\end{aligned}$$

which is exactly the claim in (3.10). The second claim follows easily from (3.11) (whose denominators are computable), and the classification theorem of finitely generated abelian groups. □

Corollary 3.2.10. *The abelianization of* $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_n \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$ *is computable. Namely:*

$$\mathbb{G}_{\mathbf{A}}^{ab} = \mathbb{Z}^n \oplus \frac{\mathbb{Z}^m}{\sum_{j=1}^{n} \operatorname{im}(\mathbf{A_j} - \mathbf{I_m})} = \mathbb{Z}^n \oplus \frac{\mathbb{Z}^m}{\operatorname{im}\widetilde{\mathbf{A}}}, \tag{3.12}$$

where $\widetilde{\mathbf{A}} \coloneqq (\mathbf{A_1} - \mathbf{I_m} \mid \mathbf{A_2} - \mathbf{I_m} \mid \cdots \mid \mathbf{A_n} - \mathbf{I_m})^{\mathsf{T}}$. *In particular,*

$$b(\mathbb{G}_{\mathbf{A}})^{ab} = n + k, \tag{3.13}$$

$$\operatorname{rk}\mathbb{G}_{\mathbf{A}}^{ab} = n + k + t, \tag{3.14}$$

where k, t *are the free-abelian, and torsion ranks of the quotient* $\mathbb{Z}^m / \operatorname{im}\widetilde{\mathbf{A}}$, *respectively.*

□

4 Algorithmic recognition of infinite cyclic extensions

In the present chapter, we study algorithmic problems about recognition of certain algebraic properties among some families of group extensions. Indeed, we see that yet for the relatively easy family of $\mathbb{Z}$-extensions one can find positive and negative results, i.e., both solvable and unsolvable "recognition problems".

For example, we prove that one cannot algorithmically decide whether a finitely presented $\mathbb{Z}$-extension admits a finitely generated base group. Even when the extension has a unique possible base group, it is not decidable in general whether this particular base group is finitely generated or not. As a consequence, we prove general undecidability for the Bieri–Neumann–Strebel invariant: there is no algorithm which, given a finite presentation for a group G, and a character $\chi \colon G \to \mathbb{R}$ as input, decides whether $[\chi]$ belongs to the BNS invariant of G, $[\chi] \in \Sigma(G)$, or not. Although this result seems quite natural, since this geometric invariant has long been agreed to be hard to compute in general (see for example [MV95; PS10; KMM15; KP14]), as far as we know, its undecidability does not seem to appear in the literature. Following our study of recognition properties, we finally consider the isomorphism problem in certain classes of unique $\mathbb{Z}$-extensions, and prove that it is equivalent to the semi-conjugacy problem for the corresponding deranged outer automorphisms (see details in Section 4.10).

The structure of the chapter is as follows. In Section 4.1 we state the recognition problems we are interested in. In Section 4.2 we introduce the general framework for our study: (finitely presented) $\mathbb{Z}^r$-extensions (denoted $*$-by-$\mathbb{Z}^r$), unique $\mathbb{Z}^r$-extensions (denoted $!$-by-$\mathbb{Z}^r$), as well as the subfamily of fg-by-$\mathbb{Z}^r$ groups, and will investigate the above problems for them. In Sections 4.4 and 4.5 we focus on case $r = 1$ (i.e., infinite cyclic extensions) which will be the main target of the chapter. The central result in Section 4.6 is Theorem 4.6.3, showing that the membership problem for fg-by-$\mathbb{Z}$ (among other similar families) is undecidable, even within the class $!$-by-$\mathbb{Z}$. As an application, Section 4.7 contains the undecidability of the BNS invariant (Theorem 4.7.4). In Section 4.8 we search for "standard presentations" of fg-by-$\mathbb{Z}$ groups (Proposition 4.8.1). Finally, in Section 4.10 we characterize the isomorphism problem in the subclass of unique $\mathbb{Z}$-extensions by means of the so-called semi-conjugacy problem (a weakened version of the standard conjugacy problem) for deranged outer automorphisms (Theorem 4.10.4).

4.1 Algorithmic recognition of groups

Algorithmic behavior of groups has been a very fundamental concern in Combinatorial and Geometric Group Theory since the very beginning of this branch of Mathematics in the early 1900s. The famous three problems stated by Max Dehn in 1911 are prototypical examples of this fact: the Word, Conjugacy, and Isomorphism problems have been very influential in the literature along these last hundred years. Today, these problems (together with a great and growing collection of variations) are the center of what is known as Algorithmic Group Theory.

Dehn's Isomorphism Problem is probably the paradigmatic example of what is popularly understood as "algorithmic recognition of groups". Namely, let $\mathcal{G}_{\mathsf{fp}}$ be the family of finite presentations of groups. Then (with the usual abuse of notation of denoting in the same way a presentation and the presented group):

- *Isomorphism problem* [IP]: given two finite presentations $G_1, G_2 \in \mathcal{G}_{\mathsf{fp}}$, decide whether they present isomorphic groups, $G_1 \simeq G_2$, or not.

It is well known that, in this full generality, Dehn's Isomorphism Problem is unsolvable, see for example [Mil92]. So, a natural next step is to study what happens when we restrict the inputs to a certain subfamily $\mathcal{H} \subseteq \mathcal{G}_{\mathsf{fp}}$:

- *Isomorphism problem within* $\mathcal{H}$ [$\mathsf{IP}(\mathcal{H})$]: given two finite presentations $H_1, H_2 \in \mathcal{H}$, decide whether they present isomorphic groups, $H_1 \simeq H_2$, or not.

Since we are interested in groups, we will only consider families of presentations closed by isomorphism; in this way, the problems considered are actually about groups (although represented by finite presentations). The literature is full of results solving the isomorphism problem for more and more such subfamilies $\mathcal{H}$ of $\mathcal{G}_{\mathsf{fp}}$, or showing its unsolvability even when restricted to smaller and smaller subfamilies $\mathcal{H}$.

Another recognition aspect is that of deciding whether a given group satisfies certain property, i.e., whether it belongs to a certain previously defined family. For two arbitrary subfamilies $\mathcal{H}, \mathcal{G} \subseteq \mathcal{G}_{\mathsf{fp}}$, we define the:

- *(Family) Membership problem for* $\mathcal{H}$ *within* $\mathcal{G}$ [$\mathsf{MP}_{\mathcal{G}}(\mathcal{H})$]: given a finite presentation $G \in \mathcal{G}$, decide whether $G \in \mathcal{H}$ or not.

If $\mathcal{H} \subseteq \mathcal{G}$ and $\mathsf{MP}_{\mathcal{G}}(\mathcal{H})$ is decidable we will also say that the inclusion $\mathcal{H} \subseteq \mathcal{G}$ is decidable. When the considered ambient family is the whole family of finitely presented groups (i.e., $\mathcal{G} = \mathcal{G}_{\mathsf{fp}}$) we will usually omit any reference to it and simply talk about the membership problem for $\mathcal{H}$, denoted $\mathsf{MP}(\mathcal{H})$.

A classic undecidability result due to Adian [Adi57b; Adi57a] and Rabin [Rab58] (see also [Mil92]) falls into this scheme. Namely, when $\mathcal{H}$ is a *Markov* subfamily (i.e., a nonempty subfamily $\emptyset \neq \mathcal{H} \subseteq \mathcal{G}_{fp}$ such that the subgroups of groups in $\mathcal{H}$ do not completely cover $\mathcal{G}_{fp}$), then $\mathsf{MP}(\mathcal{H})$ is not decidable. This turns out to include the impossibility of deciding membership for countless well known families of finitely presented groups (e.g. trivial, finite, abelian, nilpotent, solvable, free, torsion-free, simple, automatic, etc.).

Note that $\mathsf{MP}(\mathcal{H})$ being decidable is the same as saying that $\mathcal{H}$ is a recursive set of finite presentations. And, even when $\mathsf{MP}(\mathcal{H})$ is not decidable, we can still ask for a recursive enumeration of the elements in $\mathcal{H}$:

- *Enumeration problem for* $\mathcal{H}$ [$\mathsf{EP}(\mathcal{H})$]: enumerate all the elements in $\mathcal{H}$.

In many cases the considered subfamily $\mathcal{H} \subseteq \mathcal{G}_{fp}$ entails a concept of "good" or "standard" presentation $\mathcal{S} \subseteq \mathcal{H}$, for the groups presented. For example, if $\mathcal{H}$ is the family of (finite presentations for) braid groups $\{B_n \mid n \geqslant 2\}$, we can define the set of *standard* presentations $\mathcal{S} \subseteq \mathcal{H}$ to be those of the form

$$\left\langle \sigma_1, \ldots, \sigma_{n-1} \;\middle|\; \begin{array}{ll} \sigma_i\sigma_j = \sigma_j\sigma_i, & |i-j| > 1 \\ \sigma_i\sigma_j\sigma_i = \sigma_j\sigma_i\sigma_j, & |i-j| = 1 \end{array} \right\rangle ;$$

in this case, the family enumeration problem for $\mathcal{S}$ consists, on input an arbitrary finite presentation presenting a braid group, to compute its (unique) standard one for it, i.e., to recognize the number of strands n.

4.2 Group extensions

Let G and Q be arbitrary groups. We say that G is a group extension by Q (or a Q-*extension*) if G can be homomorphically mapped onto Q, i.e., if there exists a normal subgroup $H \trianglelefteqslant G$ such that the quotient G/H is isomorphic to Q.

Of course, this situation gives rise to the short exact sequence

$$1 \longrightarrow H \longrightarrow G \longrightarrow Q \longrightarrow 1$$

for some group H, and we will write that G is $*$-by-Q. One also says that such an H is a *base group* for the extension, and that G is an extension of H by Q; accordingly, if we want to specify the base group we will say that G is H-by-Q.

We remark that a given group extension by Q may admit many, even non-isomorphic, different base groups (see Corollary 4.9.2).

If H is the only (as subset) normal subgroup of G with quotient G/H isomorphic to Q, then we say that the Q-extension is *unique*, or that G is a *unique extension* by Q; in the same vein as before, we will write that G is $!$-by-Q (or that G is $!H$-by-Q, if we want to specify who is the unique normal subgroup).

It will be convenient to extend this notation allowing to replace the groups H and Q by any group property (which we will usually write in sans typeface). Concretely, given two properties of groups, $\mathcal{P}_1$, $\mathcal{P}_2$, we say that a group G is $\mathcal{P}_1$ - by - $\mathcal{P}_2$ (resp. $!\mathcal{P}_1$ - by - $\mathcal{P}_2$) if it is H - by - Q (resp. $!H$ - by - Q) for certain groups H satisfying $\mathcal{P}_1$, and Q satisfying $\mathcal{P}_2$. In this way we can easily refer to families of group extensions in terms of the behavior of their base and quotient groups. So, for example, a group G is fg - by - $\mathbb{Z}^r$ if it is H - by - $\mathbb{Z}^r$ for some finitely generated group H. And it is !fg - by - $\mathbb{Z}^r$ if it is $!H$ - by - $\mathbb{Z}^r$ for some finitely generated group H; i.e., if it has a unique normal subgroup with quotient isomorphic to $\mathbb{Z}^r$, which happens to be finitely generated (not to be confused with G having a unique finitely generated normal subgroup whose quotient is isomorphic to $\mathbb{Z}^r$ — and possibly some others infinitely generated as well).

When we need to add extra assumptions (i.e., satisfying some property $\mathcal{P}$) on the elements of certain family $\mathcal{G}$, we will denote the new family $[\mathcal{G}]_{\mathcal{P}}$. For example, $[\mathsf{abelian}]_{\mathsf{fg}}$ denotes the family of finitely generated abelian groups, while $[* \text{ - by - } \mathbb{Z}]_{\mathsf{fp}}$ denotes the family of finitely presented extensions by $\mathbb{Z}$.

4.3 Finitely generated groups (as free-abelian extensions)

The family of finitely generated groups (seen as extensions by free-abelian groups) provides both a neat framework to show some of the introduced terminology, and a first approximation to the infinite-cyclic extensions studied in the following sections. In particular, the families $*$ - by - $\mathbb{Z}^r$ and ! - by - $\mathbb{Z}^r$ turn out to be very easy to recognize algorithmically.

If G is a finitely generated group, then it is well known that its abelianization $G^{ab} = G/[G, G]$ is of the form $G^{ab} = \mathbb{Z}^r \oplus T$, where r is a nonnegative integer, and T is a finite abelian group, both canonically determined by G. The integer $r \geqslant 0$ is usually known as the first Betti number of G, denoted $b(G)$, while the finite group T is the torsion subgroup of G^{ab}, and is called the *abelian torsion* of G. Let us denote:

- π^{ab} the abelianization map of G (i.e., $\pi^{ab} \colon G \twoheadrightarrow G^{ab}$),
- π_* the projection map from G^{ab} onto its free-abelian part $\mathbb{Z}^r$ (i.e., $\pi_* \colon G^{ab} \twoheadrightarrow G^{ab} / \operatorname{Tor} G^{ab}$),
- π^{ab}_* the composition of π^{ab} followed by π_* (i.e., $\pi^{ab}_* := \pi^{ab}\pi_*$), and

- G_*^{ab} the image of G under π_*^{ab} (i.e., $G_*^{ab} := (G)\pi_*^{ab}$).

Then, we can construct the following diagram of groups and homomorphisms:

$$
\begin{array}{ccccccc}
 & & & & \mathbb{Z}^r \oplus T & & \mathbb{Z}^r \\
 & & & & \wr\! | & & \wr\! | \\
\ker \pi_*^{ab} = (\pi^{ab})^{-1}(T) & \rightarrowtail & G & \overset{\pi^{ab}}{\twoheadrightarrow} & G^{ab} & \overset{\pi_*}{\twoheadrightarrow} & G_*^{ab} \\
 & & & \multicolumn{3}{c}{\underbrace{\qquad\qquad\qquad}_{\pi_*^{ab}}} &
\end{array}
\tag{4.1}
$$

Note that both π^{ab} and π_* (and thus π_*^{ab}) in (4.1) are quotients by fully characteristic subgroups, namely: the commutator subgroup $[G, G] \leqslant G$, the torsion subgroup $T = \operatorname{Tor} G \leqslant G^{ab}$ (and $(\pi^{ab})^{-1}(T) \leqslant G$) respectively. Therefore, every endomorphism (resp. automorphism) ϕ of G factors to respective unique endomorphisms (resp. automorphisms) ϕ^{ab} and ϕ_*^{ab} of G^{ab} and G_*^{ab} respectively:

$$
\begin{array}{ccccc}
 & & \mathbb{Z}^r \oplus T & & \mathbb{Z}^r \\
 & & \wr\! | & & \wr\! | \\
G & \overset{\pi^{ab}}{\twoheadrightarrow} & G^{ab} & \overset{\pi_*}{\twoheadrightarrow} & G_*^{ab} \\
\downarrow \phi & & \downarrow \phi^{ab} & & \downarrow \phi_*^{ab} \\
G & \overset{\pi^{ab}}{\twoheadrightarrow} & G^{ab} & \overset{\pi_*}{\twoheadrightarrow} & G_*^{ab}
\end{array}
\tag{4.2}
$$

Moreover, the following chains are homomorphisms between the corresponding groups of transformations:

$$
\begin{array}{ccccc}
\operatorname{End}(G) & \to & \operatorname{End}(G^{ab}) & \to & \operatorname{End}(G_*^{ab}) \\
\phi & \mapsto & \phi^{ab} & \mapsto & \phi_*^{ab},
\end{array}
$$

and

$$
\begin{array}{ccccc}
\operatorname{Aut}(G) & \to & \operatorname{Aut}(G^{ab}) & \to & \operatorname{Aut}(G_*^{ab}) \\
\phi & \mapsto & \phi^{ab} & \mapsto & \phi_*^{ab}.
\end{array}
$$

Of course, ϕ_*^{ab} can be thought as a square $r \times r$ matrix over $\mathbb{Z}$ (with determinant ± 1 if ϕ is an automorphism). In Section 4.5 we will relate certain properties of a $\mathbb{Z}$-extension $H \rtimes_\phi \mathbb{Z}$ with the automorphism ϕ_*^{ab} associated to ϕ.

Definition 4.3.1. The *first Betti number* of a finitely generated group G, denoted $b(G)$, is the rank of the free-abelian part of its abelianization; with the previous notation,

$$b(G) := \operatorname{rk} G_*^{ab}.$$

We collect here some elementary properties of the first Betti number which will be useful later.

Lemma 4.3.2. *Let G be a finitely generated group. Then,*

1. $b(G) = b(G^{ab}) = b(G^{ab}_*) = \operatorname{rk} G^{ab}_* \leqslant \operatorname{rk}(G^{ab})$, *with equality if and only if* G^{ab} *is free-abelian;*

2. *for every subgroup* $H \leqslant G^{ab}$, $b(G^{ab}/H) = b(G^{ab}) - b(H)$;

3. *if* $G^{ab} = H_1 \oplus \cdots \oplus H_k$, *then* $b(G^{ab}) = b(H_1) + \cdots + b(H_k)$. □

From diagram (4.1) it is clear that any finitely generated group G is an extension by $\mathbb{Z}^k$, for every $k \leqslant r = b(G)$. Indeed, it is not difficult to see that $b(G)$ is the maximum rank for a free-abelian quotient of G. Hence, we have the following straightforward characterizations.

Lemma 4.3.3. *Let* G *be a finitely generated group, and* k *a nonnegative integer. Then,* G *is an extension by* $\mathbb{Z}^k$ *if and only if* $k \leqslant b(G)$, *i.e.,*

$$G \text{ is } *\text{-by-}\mathbb{Z}^k \quad \Leftrightarrow \quad k \leqslant b(G). \tag{4.3}$$

And, if so, every possible base group of G *by* $\mathbb{Z}^k$ *must contain* $\ker \pi^{ab}_* = (\pi^{ab})^{-1}(T)$.

Proof. Since $\mathbb{Z}^k$ is abelian and torsion-free, every epimorphism $\rho\colon G \twoheadrightarrow \mathbb{Z}^k$ should factor through π^{ab}_* (abelianizing and killing the torsion)

$$\begin{array}{ccc} G & \xrightarrow{\ \rho\ } & \mathbb{Z}^k \\ {\scriptstyle \pi^{ab}_*}\downarrow & \nearrow_{\tilde{\rho}} & \\ \mathbb{Z}^{b(G)} & & \end{array} \tag{4.4}$$

providing an epimorphism $\mathbb{Z}^{b(G)} \twoheadrightarrow \mathbb{Z}^k$, which implies $k \leqslant b(G)$. The converse is clear from diagram (4.1). Finally, the inclusion $\ker \pi^{ab}_* \leqslant \ker \rho$ is immediate from the factorization (4.4). □

A characterization of unique extensions by free-abelian groups follows easily.

Lemma 4.3.4. *Let* G *be a finitely generated group, and* k *a nonnegative integer. Then,* G *is a unique extension by* $\mathbb{Z}^k$ *if and only if* $k = b(G)$, *i.e.,*

$$G \text{ is } !\text{-by-}\mathbb{Z}^k \quad \Leftrightarrow \quad k = b(G). \tag{4.5}$$

If so, the (unique) base group of G *by* $\mathbb{Z}^{b(G)}$ *is* $\ker \pi^{ab}_* = (\pi^{ab})^{-1}(T)$; *i.e., every finitely generated group* G *is* $!\big((\pi^{ab})^{-1}(T)\big)\text{-by-}\mathbb{Z}^{b(G)}$.

Proof. It is enough to realize that the inclusion $\ker \pi^{ab}_* \leqslant \ker \rho$ in (4.4) is an equality — and thus the base group is unique and equal to $(\pi^{ab})^{-1}(T)$ — if and only if the epimorphism $\tilde{\rho}\colon \mathbb{Z}^{b(G)} \twoheadrightarrow \mathbb{Z}^k$ is bijective, i.e., if and only if $k = b(G)$. □

Note that one can easily compute the Betti number of any group given by a finite presentation: just abelianize it (i.e., add as relators the commutators of any pair of generators in the presentation) and then apply the Classification Theorem for finitely generated abelian groups, which is clearly algorithmic. Thus, Lemmas 4.3.3 and 4.3.4 immediately imply the decidability of the membership problem for these families of groups.

Corollary 4.3.5. *For every* $k \geqslant 0$*, the membership problem for the families* $*\text{-by-}\mathbb{Z}^k$ *and* $!\text{-by-}\mathbb{Z}^k$ *is decidable; i.e., there exists an algorithm which takes any finite presentation as input and decides whether the presented group is* $*\text{-by-}\mathbb{Z}^k$ *(resp.* $!\text{-by-}\mathbb{Z}^k$*) or not.* □

4.4 Infinite-cyclic extensions

We will concentrate now on infinite cyclic extensions, concretely in the family $*\text{-by-}\mathbb{Z}$ and its subfamily $!\text{-by-}\mathbb{Z}$. Let us describe them in a different way: since $\mathbb{Z}$ is a free group, every short exact sequence of the form $1 \to H \to G \to \mathbb{Z} \to 1$ splits, and so G is a semidirect product of H by $\mathbb{Z}$; namely $G \simeq H \rtimes_\alpha \mathbb{Z}$, for some $\alpha \in \operatorname{Aut}(H)$. Let us recall this well known construction in order to fix our notation.

Given an arbitrary group H and an automorphism $\alpha \in \operatorname{Aut}(H)$, define the *semidirect product of* H *by* $\mathbb{Z}$ *determined by* α as the group $H \rtimes_\alpha \mathbb{Z}$ with underlying set $H \times \mathbb{Z}$ and operation given by

$$(h, m) \cdot (k, n) = (h\,\alpha^m(k), m + n), \tag{4.6}$$

for all $h, k \in H$, and $m, n \in \mathbb{Z}$. Of course, $h \mapsto (h, 0)$ is a natural embedding of H in $H \rtimes_\alpha \mathbb{Z}$, and we then have the natural short exact sequence

$$1 \to H \to H \rtimes_\alpha \mathbb{Z} \to \mathbb{Z} \to 1. \tag{4.7}$$

Therefore, $H \rtimes_\alpha \mathbb{Z}$ belongs to the family $*\text{-by-}\mathbb{Z}$. Recall that we can have $H \rtimes_\alpha \mathbb{Z} \simeq K \rtimes_\beta \mathbb{Z}$, with $H = K$ but $\alpha \neq \beta$; or even with $H \not\simeq K$. We discuss this phenomena in Section 4.9 (see Lemma 4.9.3, and Corollary 4.9.2, respectively).

In the opposite direction, assume that G is in the family $*\text{-by-}\mathbb{Z}$. Choose a homomorphism onto $\mathbb{Z}$, say $\rho\colon G \twoheadrightarrow \mathbb{Z}$, and consider the short exact sequence given by

$$1 \to H \to G \overset{\rho}{\twoheadrightarrow} \mathbb{Z} \to 1,$$

where $H = \ker \rho \trianglelefteqslant G$. Choose and denote by t a preimage in G of any of the two generators of $\mathbb{Z}$ (note that choosing such t is equivalent to choosing a split homomorphism for ρ). Now consider the conjugation by t in G, say $\gamma_t\colon G \to G$, $g \mapsto tgt^{-1}$, and denote by $\alpha \in \operatorname{Aut}(H)$ its restriction to H (note that γ_t is an inner automorphism of G, but α may very well not be inner as an automorphism of H). By construction, we have

$$th = \alpha(h)t, \tag{4.8}$$

for every $h \in H$. At this point, it is clear that every element from G can be written in a unique way as ht^k, for some $h \in H$ and $k \in \mathbb{Z}$. And—from (4.8)—the operation in G can be easily understood by thinking that t (respectively, t^{-1}) jumps to the right of elements from H at the price of applying α (respectively, α^{-1}):

$$ht^m \cdot kt^n = h\,\alpha^m(k)\,t^{m+n}. \tag{4.9}$$

This is, precisely, the multiplicative version of (4.6). Hence, $G \simeq H \rtimes_\alpha \mathbb{Z}$, the semidirect product of H by $\mathbb{Z}$ determined by α.

From this discussion it follows easily that, for any presentation of H, say $H = \langle X \mid R\rangle$, and any $\alpha \in \mathrm{Aut}(H)$, the semidirect product $G = H \rtimes_\alpha \mathbb{Z}$ admits a presentation of the form

$$\left\langle X, t \;\middle|\; R\,,\, txt^{-1} = \alpha(x)\ (x \in X) \right\rangle. \tag{4.10}$$

Note that (4.10) is a finite presentation if and only if the initial presentation for H was finite. So, a group G admits a finite presentation of type (4.10) if and only if G is fp-by-$\mathbb{Z}$. This provides the notion of standard presentation in this context.

Definition 4.4.1. A *standard presentation* for a fp-by-$\mathbb{Z}$ group G is a *finite* presentation of the form (4.10).

The previous discussion provides the following alternative descriptions for the family of finitely presented $\mathbb{Z}$-extensions. For any group H, we have

$$[H\text{-by-}\mathbb{Z}]_{\mathsf{fp}} = \{H \rtimes_\alpha \mathbb{Z} \text{ f.p.} \mid \alpha \in \mathrm{Aut}(H)\},$$

and then,

$$\begin{aligned}[*\text{-by-}\mathbb{Z}]_{\mathsf{fp}} &= \{\, G \text{ fp} \mid b(G) \geqslant 1 \,\} \\ &= \{\, H \rtimes_\alpha \mathbb{Z} \text{ fp} \mid H \text{ group, and } \alpha \in \mathrm{Aut}(H) \,\}.\end{aligned}$$

Remark 4.4.2. Note that we have made no assumptions on the base group H. Imposing natural conditions on it, we get the inclusions

$$\mathsf{fp}\text{-by-}\mathbb{Z} \subseteq [\mathsf{fg}\text{-by-}\mathbb{Z}]_{\mathsf{fp}} \subseteq [*\text{-by-}\mathbb{Z}]_{\mathsf{fp}}\,, \tag{4.11}$$

which will be seen throughout the paper to be both strict.

The strictness of the second inclusion in (4.11) is a direct consequence of Corollary 4.6.2, while the strictness of the first one is proved below (we thank Conchita Martínez for pointing out the candidate group (4.12) in the subsequent proof).

Proposition 4.4.3. *The inclusion* $\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z} \subsetneq [\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}]_{\mathsf{fp}}$ *is strict. That is, there exist finitely presented* $\mathbb{Z}$*-extensions of finitely generated groups, which are not* $\mathbb{Z}$*-extensions of any finitely presented group.*

Proof. Let $p < q < r$ be three different prime numbers, and consider the additive group A of the ring $\mathbb{Z}[\frac{1}{p}, \frac{1}{q}, \frac{1}{r}]$, which is well known to be generated by $X = \{1/p^n\}_{n\in\mathbb{N}} \cup \{1/q^n\}_{n\in\mathbb{N}} \cup \{1/r^n\}_{n\in\mathbb{N}}$, but not finitely generated (for any given finite set of elements in A, let k be the biggest p-exponent in the denominators, and it is easy to see that $1/p^{k+1} \in A$ is not in the subgroup generated by them). Finally consider the two commuting automorphisms $\alpha, \beta \colon A \to A$ given by $\alpha \colon a \mapsto \frac{p}{r}a$, and $\beta \colon a \mapsto \frac{q}{r}a$.

Our candidate G is the (metabelian) semidirect product of A by $\mathbb{Z}^2 = \langle t, s \mid [t, s] \rangle$, with action $t \mapsto \alpha$, $s \mapsto \beta$, namely,

$$\begin{aligned} G &= A \rtimes_{\alpha,\beta} \mathbb{Z}^2 \\ &= \left\langle A, t, s \;\middle|\; ts = st,\ tat^{-1} = \frac{p}{r}a,\ sas^{-1} = \frac{q}{r}a \right\rangle. \end{aligned} \tag{4.12}$$

(One has to be careful here with the notation: it is typically multiplicative for the nonabelian group G, but additive for the abelian group A, while α and β are defined using products of rational numbers; beware, in particular, of the element $1 \in \mathbb{Z} \subseteq A$ which is additive and, of course, nontrivial.)

It is easy to see that G is generated by $1 \in A$, and $t, s \in \mathbb{Z}^2$: indeed, conjugating 1 by all powers of t and s we obtain, respectively, p^n/r^n and q^n/r^n, and then $\lambda_n p^n/r^n + \mu_n q^n/r^n = 1/r^n$ for appropriate integers λ_n, μ_n, by Bezout's identity; with the same trick and having r^n/p^n and $1/r^n$, we get $1/p^n r^n$ and so, $1/p^n$; and similarly, one gets $1/q^n$. Note that, in order to obtain all of A, it is enough to get $1/p^n$, $1/q^n$ and $1/r^n$ for n big enough; this will be used later.

To see that G is finitely presented, it is enough to use Theorem A(ii) in [BS80], which provides a precise condition for a finitely generated metabelian group to be finitely presented. This is a result, due to Bieri and Strebel, that later lead to the development of the so called Bieri–Neumann–Strebel theory (see Section 4.7).

Note that the group G is finitely generated and metabelian, having A as an abelian normal subgroup with quotient $\mathbb{Z}^2$. We know that A is not finitely generated as group; however, with $\mathbb{Z}^2$ acting by conjugation, A becomes a $\mathbb{Z}^2$-module, which is finitely generated by the exact same argument as in the previous paragraph. But even more: for all nontrivial valuation $v \colon \mathbb{Z}^2 \to \mathbb{R}$, A is also finitely generated over at least one of the two monoids $\{(n,m) \in \mathbb{Z}^2 \mid \nu(n,m) \geqslant 0\}$, or $\{(n,m) \in \mathbb{Z}^2 \mid \nu(n,m) \leqslant 0\}$. This is because any such valuation has the form $(n,m) \mapsto \alpha n + \beta m$ for some $(0,0) \neq (\alpha,\beta) \in \mathbb{Z}^2$, and then it is routine to show that, starting with $1 \in A$, conjugating only either by those $t^n s^m$ with $\alpha n + \beta m \geqslant 0$, or those with $\alpha n + \beta m \leqslant 0$, and adding, we can get all of A (we leave the details to the reader). By Theorem A(ii) from [BS80], this implies that the group G is finitely presented.

Now consider the subgroup $H = \langle ts, A \rangle \leqslant G$, which is clearly normal and produces a quotient $G/H = \mathbb{Z} = \langle z \mid - \rangle$. Since

$$H \simeq A \rtimes_{\alpha\circ\beta} \mathbb{Z} = \left\langle A, z \;\middle|\; zaz^{-1} = \frac{pq}{r^2} a \right\rangle$$

is generated by $1, z$ (by the same reason as above), we deduce that G is both a $\mathbb{Z}$-extension of its finitely generated subgroup H, and finitely presented; i.e., $G \in [\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}]_{\mathsf{fp}}$.

It remains to see that $G \notin \mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ (i.e., G is not a $\mathbb{Z}$-extension of any finitely presented subgroup). We do not know whether this is true for every p, q, r, but we shall prove it for particular values of the parameters; concretely for $(p,q,r) = (2,3,5)$.

It is easy to see that the derived subgroup G' is contained in A. We shall prove that, when this inclusion is indeed an equality—for example, when $(p,q,r) = (2,3,5)$, as it is straightforward to see—then G is not a $\mathbb{Z}$-extension of any finitely presented subgroup. That is, no normal subgroup $N \trianglelefteqslant G$ with $G/N \simeq \mathbb{Z}$ can be finitely presented. In fact, let $N \trianglelefteqslant G$ be such a subgroup. Then $A = G' \trianglelefteqslant N \trianglelefteqslant G$ and, taking quotients by A, we obtain $1 \trianglelefteqslant N/A \trianglelefteqslant G/A = \mathbb{Z}^2 = \langle t, s \rangle$. But $\mathbb{Z} \simeq G/N \simeq \frac{G/A}{N/A} \simeq \frac{\mathbb{Z}^2}{N/A}$. So, it must be $N/A \simeq \mathbb{Z}$.

Now, choose $(n,m) \neq (0,0)$ such that N/A is generated by $t^n s^m A$ (we can clearly assume $n > 0$); and deduce that $N \simeq A \ltimes_\varphi \mathbb{Z}$ with action $\varphi\colon a \mapsto \frac{p^n q^m}{r^{n+m}} a$. In particular, N is finitely generated by an argument as above. Note also that the action of φ is by multiplication by a simplified fraction, say λ/μ, with λ and μ both different from ± 1 (if $m \geqslant 0$ it is multiplication by $\frac{p^n q^m}{r^{n+m}}$; and if $m < 0$, it is multiplication by $\frac{p^n r^{|m|}}{r^n q^{|m|}}$).

Finally, let us apply again Theorem A(ii) in [BS80], now to the short exact sequence $1 \to A \to N \to \mathbb{Z} \to 1$. The only nontrivial valuations $\mathbb{Z} \to \mathbb{R}$ are $1 \mapsto \pm 1$, and it is easy to see that A is not finitely generated neither as a $\mathbb{Z}^+$-module (with finitely many elements one cannot obtain $1/\lambda^n$ for $n \gg 0$), nor as a $\mathbb{Z}^-$-module (with finitely many elements one cannot obtain $1/\mu^n$ for $n \gg 0$). Therefore, N is not finitely presented, and the group G is not a $\mathbb{Z}$-extension of any finitely presented subgroup, as we wanted to prove. □

4.5 Unique infinite-cyclic extensions

Recall that the family of unique $\mathbb{Z}$-extensions (i.e., groups having a unique normal subgroup with quotient $\mathbb{Z}$) which are finitely presented is denoted $[!\text{-by-}\mathbb{Z}]_{\mathrm{fp}}$.

As seen in Lemma 4.3.4, the family $[!\text{-by-}\mathbb{Z}]_{\mathrm{fp}}$ consists precisely of those groups G in $[*\text{-by-}\mathbb{Z}]_{\mathrm{fp}}$ such that $b(G) = 1$. For finitely generated base groups H, Proposition 4.5.5 provides a quite simple characterization of this unicity condition in terms of the defining automorphism α.

We first obtain a convenient description of the center and abelianization of an infinite-cyclic extension $H \rtimes_\alpha \mathbb{Z} = \langle X, t \mid R,\ txt^{-1} = \alpha(x)\ (x \in X) \rangle$ in terms of the defining automorphism $\alpha \in \operatorname{Aut} H$.

4.5.1 Center and abelianization.

Observe that, if $\alpha \neq \mathrm{id}_H$ then t does not commute with some element of H (namely, those $h \in H$ such that $\alpha(h) \neq h$); in particular, $H \rtimes_\alpha \mathbb{Z}$ is not abelian in this cases. In fact, it is straightforward to obtain a complete characterization of the elements in the center of $H \rtimes_\alpha \mathbb{Z}$ just reinterpreting the commutativity conditions using.

Lemma 4.5.1. *Let* $\alpha \in \operatorname{Aut}(H)$, *and* $h\,t^k \in H \rtimes_\alpha \mathbb{Z}$. *Then,*

1. $h\,t^k$ *commutes with every element in* H *if and only if* $\alpha^k = \gamma_{h^{-1}} \in \operatorname{Inn}(H)$,
2. $h\,t^k$ *commutes with* $t \in \mathbb{Z}$ *if and only if* $\alpha(h) = h$.

Therefore, the center of $H \rtimes_\alpha \mathbb{Z}$ *is*

$$Z(H \rtimes_\alpha \mathbb{Z}) = \{h\,t^k \,:\, h \in \operatorname{Fix}(\alpha) \text{ and } \alpha^k = \gamma_{h^{-1}}\}. \tag{4.13}$$

□

A characterization of the triviality of the center of $H \rtimes_\alpha \mathbb{Z}$ follows immediately.

Corollary 4.5.2. *The group* $H \rtimes_\alpha \mathbb{Z}$ *has nontrivial center if and only if either* α *is of finite order in* $\operatorname{Aut}(H)$, *or* α^k *is conjugation by a nontrivial* α*-fixed element, for some* $k \in \mathbb{Z}$. □

In particular we get quite neat characterizations in the following two degenerated cases:

Corollary 4.5.3. *For* $\alpha \in \operatorname{Aut}(H)$,

1. *if* $\operatorname{Fix}(\alpha)$ *is trivial, then:* $Z(H \rtimes_\alpha \mathbb{Z}) = 1 \Leftrightarrow \alpha$ *is of infinite order in* $\operatorname{Aut}(H)$,
2. *if* $Z(H)$ *is trivial, then:* $Z(H \rtimes_\alpha \mathbb{Z}) = 1 \Leftrightarrow [\alpha]$ *is of infinite order in* $\operatorname{Out}(H)$. □

Lemma 4.5.4. *Let* H *be an arbitrary group, and let* $\alpha \in \operatorname{Aut}(H)$*. Then,*

$$(H \rtimes_\alpha \mathbb{Z})^{ab} \simeq \frac{H^{ab}}{\operatorname{im}(\alpha^{ab} - \mathrm{id})} \oplus \mathbb{Z}. \tag{4.14}$$

Moreover, if H *is finitely generated, then so is* $H \rtimes_\alpha \mathbb{Z}$*, and*

$$(H \rtimes_\alpha \mathbb{Z})^{ab} \simeq \mathbb{Z}^{k+1} \oplus T, \tag{4.15}$$

where k *is the rank of* $\ker(\alpha_*^{ab} - \mathrm{id})$*, and* T *is a finite abelian group.*

Proof. Let $H = \langle X \mid R \rangle$. Abelianizing $H \rtimes_\alpha \mathbb{Z} = \langle X, t \mid R,\ tx_i t^{-1} = \alpha(x_i)\ \ (x_i \in X) \rangle$, we get

$$\begin{aligned}
(H \rtimes_\alpha \mathbb{Z})^{ab} &= \left\langle X, t \;\middle|\; \begin{array}{ll} R, & \\ tx_i t^{-1} = \alpha(x_i) & (x_i \in X), \\ x_i x_j = x_j x_i & (x_i, x_j \in X), \\ tx_i = x_i t & (x_i \in X) \end{array} \right\rangle \\
&\simeq \left\langle X \;\middle|\; \begin{array}{ll} R, & \\ x_i = \alpha(x_i) & (x_i \in X), \\ x_i x_j = x_j x_i & (x_i, x_j \in X) \end{array} \right\rangle \times \langle t \mid - \rangle \\
&\simeq \frac{H^{ab}}{\operatorname{im}(\alpha^{ab} - \mathrm{id})} \oplus \langle t \mid - \rangle.
\end{aligned} \tag{4.16}$$

For the second part, suppose that H is finitely generated. Then so is H^{ab} and thus, using Lemma 4.3.2, we have

$$\begin{aligned}
b(H \rtimes_\alpha \mathbb{Z}) &= b_1\left(\frac{H^{ab}}{\operatorname{im}(\alpha^{ab} - \mathrm{id})} \right) + 1 \\
&= b(H^{ab}) - b(\operatorname{im}(\alpha^{ab} - \mathrm{id})) + 1 \\
&= b(H_*^{ab}) - b(\operatorname{im}(\alpha_*^{ab} - \mathrm{id})) + 1 \\
&= b(\ker(\alpha_*^{ab} - \mathrm{id})) + 1 \\
&= \operatorname{rk}(\ker(\alpha_*^{ab} - \mathrm{id})) + 1,
\end{aligned} \tag{4.17}$$

which is, precisely, what we wanted to prove. □

4.5.2 Deranged automorphisms

This last result, combined with Lemma 4.3.4, provides a computable characterization for automorphisms defining !fg-by-$\mathbb{Z}$ groups. We state the result in a more general setting, the algorithmic part corresponding to the finitely generated base group case.

Proposition 4.5.5. *Let* H *be an arbitrary group, and* $\alpha \in \operatorname{Aut}(H)$ *such that the semidirect product* $H \rtimes_\alpha \mathbb{Z}$ *is finitely generated. Then, the following conditions are equivalent:*

a. $H \rtimes_\alpha \mathbb{Z}$ *is* $!$-by-$\mathbb{Z}$*;*

b. $b(H \rtimes_\alpha \mathbb{Z}) = 1$*;*

c. $H = \ker(\pi_*^{ab}) = (\pi^{ab})^{-1}(T)$

(i.e., H *is the full preabelianization of the torsion subgroup* $T \leqslant (H \rtimes_\alpha \mathbb{Z})^{ab}$*);*

d. H *is a fully characteristic subgroup of* $H \rtimes_\alpha \mathbb{Z}$*.*

Moreover, if H *is finitely generated, then the following additional condition is also equivalent:*

e. α_*^{ab} *has no nontrivial fixed points (equivalently, 1 is not an eigenvalue of* α_*^{ab}*,* $\ker(\alpha_*^{ab} - \mathrm{id}) = 0$*, or* $\det(\alpha_*^{ab} - \mathrm{id}) \neq 0$*).*

Proof. [a. $\Leftrightarrow$ b. $\Leftrightarrow$ c.]. This is precisely the content of Lemma 4.3.4, for $k = 1$.

[c. $\Rightarrow$ d.]. This is clear, since $\ker(\pi_*^{ab}) = (\pi^{ab})^{-1}(T)$, the full preabelianization of the torsion subgroup T of the abelian group $(H \rtimes_\alpha \mathbb{Z})^{ab}$, which is fully characteristic.

[d. $\Rightarrow$ b.]. By contradiction, suppose that $H \rtimes_\alpha \langle t \mid - \rangle$ has Betti number at least 2. Then, there exists an epimorphism $\rho \colon H \rtimes_\alpha \langle t \rangle \twoheadrightarrow \mathbb{Z}^2$, and thus an element $h \in H$ such that $\rho(h) \neq 0$. Take a primitive element $v \in \mathbb{Z}^2$ such that $\rho(h) = \lambda v$ for some $\lambda \in \mathbb{Z}$, $\lambda \neq 0$. The subgroup $\langle v \rangle \simeq \mathbb{Z}$ is a direct summand of $\mathbb{Z}^2$, and the composition

$$\begin{array}{ccccccc} H \rtimes_\alpha \langle t \rangle & \overset{\rho}{\twoheadrightarrow} & \mathbb{Z}^2 & \twoheadrightarrow & \langle v \rangle & \rightarrowtail & H \rtimes_\alpha \langle t \rangle \\ & & & & v & \mapsto & t \end{array}$$

provides an endomorphism of $H \rtimes_\alpha \mathbb{Z}$ mapping $h \in H$ to t^λ, a contradiction with condition d..

Finally, assume H is finitely generated. Then, from (4.15) in Lemma 4.5.4, we have that $b(H \rtimes_\alpha \mathbb{Z}) = 1$ if and only if $\operatorname{rk}\ker(\alpha_*^{ab} - \mathrm{id}) = 0$, i.e., if and only if $\ker(\alpha_*^{ab} - \mathrm{id}) = 0$. This proves [b. $\Leftrightarrow$ e.]. $\square$

Definition 4.5.6. We say that an automorphism $\alpha \in \mathrm{Aut}(H)$ is *deranged* if one of (and thus all) the conditions a. – d. in Proposition 4.5.5 hold. Note that when the base group H is finitely generated, condition e. provides another equivalent definition which is, furthermore, clearly algorithmic with respect to the defining automorphism: given $\alpha \in \mathrm{Aut}(H)$ by images of generators, it is easy to algorithmically check whether α is deranged or not.

In particular, every automorphism of a group H with $b(H) = 0$ (i.e., with finite abelianization) is trivially deranged. Note also that derangedness is, in fact, a property of outer automorphisms. The sets of deranged automorphisms and deranged outer automorphisms of a group H will be denoted, respectively, $\mathrm{Aut}_d(H)$ and $\mathrm{Out}_d(H)$.

Consequently, for any finitely generated group H we have

$$[!H\text{-by-}\mathbb{Z}]_{fp} = \{ H \rtimes_\alpha \mathbb{Z} \text{ fp} \mid \alpha \in \mathrm{Aut}_d(H) \},$$

and then,

$$\begin{aligned} [!\text{fg-by-}\mathbb{Z}]_{fp} &= \{ H \rtimes_\alpha \mathbb{Z} \text{ fp} \mid H \text{ fg and } \alpha \in \mathrm{Aut}_d(H) \} \\ &\subseteq [!\text{-by-}\mathbb{Z}]_{fp} = \{ G \text{ fp} \mid b(G) = 1 \}. \end{aligned} \tag{4.18}$$

Note that in (4.18) we wrote inclusion and not an equality because, in principle, it could happen that a finitely presented !-by-$\mathbb{Z}$ group has his unique normal subgroup with quotient isomorphic to $\mathbb{Z}$ being not finitely generated. In the next section we shall construct such a group (see Corollary 4.6.2), showing that this inclusion is strict.

4.6 Undecidability results

In this section, we will be concerned with inclusions between subfamilies of infinite-cyclic extensions. For example, it is immediate that if H is finitely generated or finitely presented, then so is $H \rtimes_\alpha \mathbb{Z}$ for every $\alpha \in \mathrm{Aut}(H)$, i.e.,

$$\begin{aligned} \text{fg-by-}\mathbb{Z} &\subseteq [*\text{-by-}\mathbb{Z}]_{fg}, \\ \text{fp-by-}\mathbb{Z} &\subseteq [*\text{-by-}\mathbb{Z}]_{fp}. \end{aligned}$$

However, it is less obvious that the converse is not true in general: a semidirect product $H \rtimes_\alpha \mathbb{Z}$ can be finitely presented, with H not even being finitely generated. Or, as was hinted few lines above, even worse: there do exist *finitely presented* $\mathbb{Z}$-extensions which are *not* $\mathbb{Z}$-extensions of *any* finitely generated group. In other words, the following inclusion is strict:

$$[\text{fg-by-}\mathbb{Z}]_{fp} \subset [*\text{-by-}\mathbb{Z}]_{fp}.$$

Indeed, this can happen even for unique $\mathbb{Z}$-extensions. This fact follows easily from the next lemma, showing that any free product $K * \mathbb{Z}$ has the form of a certain semidirect product.

Lemma 4.6.1. *Let* $K = \langle X \mid R \rangle$ *be an arbitrary group with generators* $X = \{x_j\}_{j \in J}$, *and consider the free product*

$$\mathop{\ast}_{i \in \mathbb{Z}} K = \left\langle X^{(i)} \, (i \in \mathbb{Z}) \;\middle|\; R^{(i)} \, (i \in \mathbb{Z}) \right\rangle,$$

where $\langle X^{(i)} \mid R^{(i)} \rangle$ $(i \in \mathbb{Z})$ *are disjoint copies of the original presentation for* K*. Then,*

$$\Big(\mathop{\ast}_{i\in\mathbb{Z}} K \Big) \rtimes_\tau \mathbb{Z} \;\simeq\; K * \mathbb{Z}, \tag{4.19}$$

where τ *is the automorphism of* $\ast_{i\in\mathbb{Z}} K$ *defined by*

$$\tau\colon x_j^{(i)} \mapsto x_j^{(i+1)} \;\Big(\forall i \in \mathbb{Z},\, \forall x_j^{(i)} \in X^{(i)}\Big). \tag{4.20}$$

Proof. Naming t the generator of $\mathbb{Z}$, we have

$$\begin{aligned} \Big(\mathop{\ast}_{i\in\mathbb{Z}} K \Big) \rtimes_\tau \mathbb{Z} &= \left\langle t, X^{(i)}\ (i \in \mathbb{Z}) \;\middle|\; \begin{array}{l} R^{(i)}\ (i \in \mathbb{Z}), \\ t\,x_j^{(i)}\,t^{-1} = x_j^{(i+1)}\ \Big(i \in \mathbb{Z},\, x_j^{(i)} \in X^{(i)}\Big) \end{array} \right\rangle \\ &\simeq \left\langle t, X^{(i)}\ (i \in \mathbb{Z}) \;\middle|\; \begin{array}{l} R^{(0)}, \\ t\,x_j^{(i)}\,t^{-1} = x_j^{(i+1)}\ \Big(i \in \mathbb{Z},\, x_j^{(i)} \in X^{(i)}\Big) \end{array} \right\rangle \end{aligned} \tag{4.21}$$

$$\simeq \left\langle t, X^{(0)} \;\middle|\; R^{(0)} \right\rangle = K * \mathbb{Z}. \tag{4.22}$$

To see the last isomorphism, consider the maps from (4.22) to (4.21) given by

$$\begin{aligned} t &\mapsto t, \\ x_j^{(0)} &\mapsto x_j^{(0)}\ (x_j^{(0)} \in X^{(0)}), \end{aligned}$$

and from (4.21) to (4.22) given by

$$\begin{aligned} t &\mapsto t, \\ x_j^{(i)} &\mapsto t^i x_j^{(0)} t^{-i}\ (i \in \mathbb{Z},\ x_j^{(i)} \in X^{(i)}). \end{aligned}$$

It is straightforward to see that they are both well-defined homomorphisms, and one inverse to the other. □

Corollary 4.6.2. *If* K *is a group with finite abelianization (i.e.,* $b(K) = 0$*), then the free product* $K * \mathbb{Z}$ *is a unique* $\mathbb{Z}$*-extension, and the following conditions are equivalent:*

(a) $K * \mathbb{Z}$ *is* fg-by-$\mathbb{Z}$*;*

(b) $K * \mathbb{Z}$ *is* fp-by-$\mathbb{Z}$*;*

(c) $K * \mathbb{Z}$ *is* $\mathsf{abelian}$-by-$\mathbb{Z}$*;*

(d) $K * \mathbb{Z}$ *is* finite-by-$\mathbb{Z}$*;*

(e) $K * \mathbb{Z}$ *is* free-by-$\mathbb{Z}$*;*

(f) $K = 1$.

In particular, $!\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ *is a strict subfamily of* $[!\text{-}\mathsf{by}\text{-}\mathbb{Z}]_{\mathsf{fp}}$ *(and so,* $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ *is a strict subfamily of* $[*\text{-}\mathsf{by}\text{-}\mathbb{Z}]_{\mathsf{fp}}$*).*

Proof. Note that the abelianization of $K * \mathbb{Z}$ is $(K * \mathbb{Z})^{ab} = K^{ab} \oplus \mathbb{Z}$, where $|K^{ab}| < \infty$ by hypothesis; therefore, $b(K * \mathbb{Z}) = 1$. Thus, from Lemma 4.3.4, $K * \mathbb{Z}$ is a unique $\mathbb{Z}$-extension, i.e., it contains a unique normal subgroup with quotient $\mathbb{Z}$. By Lemma 4.6.1, this unique normal subgroup is isomorphic to $\ast_{z \in \mathbb{Z}} K$, which is finitely generated (resp., finitely presented, abelian, finite, free) if and only if K is trivial (the free case being true because $b(K) = 0$).

Taking K to be a nontrivial finitely presented group with finite abelianization, we obtain that $K * \mathbb{Z}$ belongs to $[!\text{-}\mathsf{by}\text{-}\mathbb{Z}]_{\mathsf{fp}}$ but not to $!\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$. □

Next, inspired by a trick initially suggested by Maurice Chiodo, we will prove a stronger result. Not only the family $!\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ is a strict subfamily of $[!\text{-}\mathsf{by}\text{-}\mathbb{Z}]_{\mathsf{fp}}$, but the membership problem between these two families is undecidable: it is impossible to decide algorithmically whether a given finitely presented unique $\mathbb{Z}$-extension is $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ or not, i.e., whether its unique base group is finitely generated or not. To see this, we use a classic undecidability result: there is no algorithm which, on input a finite presentation, decides whether the presented group is trivial or not (see, for example, [Mil92]).

Theorem 4.6.3. *For every group property* $\mathcal{P} \in \{\mathsf{fg}, \mathsf{fp}, \mathsf{abelian}, \mathsf{finite}, \mathsf{free}\}$*, the membership problem for* $\mathcal{P}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ *within* $[!\text{-}\mathsf{by}\text{-}\mathbb{Z}]_{\mathsf{fp}}$ *is undecidable.*

In other words, there exists no algorithm which, on input a finite presentation of a group with Betti number 1*, decides whether it presents a* $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ *(resp.* $\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}$, $\mathsf{abelian}\text{-}\mathsf{by}\text{-}\mathbb{Z}$, $\mathsf{finite}\text{-}\mathsf{by}\text{-}\mathbb{Z}$, $\mathsf{free}\text{-}\mathsf{by}\text{-}\mathbb{Z}$*) group or not.*

Proof. We will proceed by contradiction. Assume the existence of an algorithm, say $\mathfrak{A}$, such that, given as input a finite presentation of a group with Betti number 1, outputs YES if it presents a $\mathcal{P}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ group, and NO otherwise.

Now, consider the following algorithm $\mathfrak{B}$ to check the triviality of an arbitrary finite presentation $K = \langle X \mid R \rangle$ given as input:

1. abelianize K and, using the Classification Theorem for finitely generated abelian groups, check whether K^{ab} is trivial or not; if not, answer NO; otherwise K is a perfect group and so, the new group $K * \mathbb{Z}$ has Betti number 1;

2. apply $\mathfrak{A}$ to the presentation $\langle X, t \mid R \rangle$, to decide whether $K * \mathbb{Z}$ is a $\mathcal{P}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ group or not.

According to Corollary 4.6.2, the output to step (ii) is YES if and only if K is trivial. Hence, the algorithm $\mathfrak{B}$ decides whether the given presentation $\langle X \mid R \rangle$ presents the trivial group or not. This contradicts Adian–Rabin's Theorem on the undecidability of the triviality problem, and therefore there is no such algorithm $\mathfrak{A}$, as we wanted to prove □

Of course, if the membership problem is not decidable within some family $\mathcal{H}$, it is also undecidable within any superfamily of $\mathcal{H}$. So, we immediately get the following consequence.

Corollary 4.6.4. *The membership problems for the families* $\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ *and* $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ *are undecidable.* □

As stated in the introduction, this is exactly the same as saying that the families (of finite presentations) $\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ and $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ are not recursive.

Remark 4.6.5. Note that none of these families is neither Markov nor co-Markov, and thus the two undecidability results in Corollary 4.6.4 are not consequence of the classic result due to Adian–Rabin. Indeed, any finitely presented group is a subgroup of some $\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ (and so, of some $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$) group; therefore the families $\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ and $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ are not Markov. On the other hand, every fp group embeds into some 2-generated simple group (see [Mil92, Corollary 3.10]); since

$$\mathsf{simple} \ \Rightarrow \ \mathsf{perfect} \ \Rightarrow \ \neg(\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}) \ \Rightarrow \ \neg(\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}),$$

the families $\mathsf{fp}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ and $\mathsf{fg}\text{-}\mathsf{by}\text{-}\mathbb{Z}$ are not co-Markov either.

4.7 Implications for the BNS invariant

Since the early 1980s, in a series of papers by R. Bieri, W. Neumann, and R. Strebel (see [BS80; BNS87]), several gradually more general invariants —called Sigma (or BNS) Invariants— have been introduced to deal with finiteness conditions for presentations of groups. Concretely in [BNS87], they present an invariant that characterizes those normal subgroups of a finitely generated group G that are finitely generated and contain the commutator subgroup of G. Over the years, this theory has been reformulated in more geometric terms (for a modern version see the survey [Str12]). Below, we recall this construction and characterization, and discuss some implications of our undecidability results from Section 4.6.

For a finitely generated group G, consider the real vector space $\mathrm{Hom}(G, \mathbb{R})$ of all homomorphisms $\chi \colon G \to \mathbb{R}$ (from G to the additive group of the field of real numbers) which we call *characters* of G. Note that, since $\mathbb{R}$ is abelian and torsion-free, any character χ must factor through π_*^{ab} (abelianizing and then killing the torsion), i.e.,

$$\chi\colon G \xrightarrow{\pi^{ab}} G^{ab} \xrightarrow{\pi_*} G_*^{ab} \longrightarrow \mathbb{R}\,.$$

Thus, $\mathrm{Hom}(G,\mathbb{R}) = \mathrm{Hom}(\mathbb{Z}^r,\mathbb{R}) = \mathbb{R}^r$, where $r = b(G)$. We will consider the set of nontrivial characters modulo the equivalence relation given by positive scaling:

$$\chi_1 \sim \chi_2 \quad\Leftrightarrow\quad \exists \lambda > 0 \text{ s.t. } \chi_2 = \lambda\chi_1\,. \tag{4.23}$$

They form the so-called *character sphere* of G, denoted $S(G) = \mathrm{Hom}(G,\mathbb{R})^*/\sim$ which, equipped with the quotient topology, is homeomorphic to the unit Euclidean sphere of dimension $r-1$ (through the natural identification of each ray emanating from the origin with its unique point of norm 1).

For example, if G is not $*$-by-$\mathbb{Z}$ (i.e., if $b(G) = 0$), then $\mathrm{Hom}(G,\mathbb{R}) = \{0\}$ and the sphere character is empty (so, for this class of groups the BNS theory will be vacuous). More interestingly, if G is $!$-by-$\mathbb{Z}$ (i.e., $b(G) = 1$), then the character sphere of G is a set of just two points, namely $S(G) = \{+1,-1\}$. Similarly, if $b(G) = 2,3,\ldots$, then $S(G)$ is the unit circle in $\mathbb{R}^2$, the unit sphere in $\mathbb{R}^3$, and so on.

For any given (equivalence class of a) nontrivial character χ, consider now the following submonoid of G, called the *positive cone* of χ:

$$G_\chi = \{g \in G \mid \chi(g) \geqslant 0\} = \chi^{-1}([0,+\infty)), \tag{4.24}$$

to be thought of as the full subgraph of the Cayley graph $\mathsf{Cay}(G,X)$ determined by the vertices in G_χ (once a set of generators for G is fixed). The Sigma invariant $\Sigma(G)$ can then be defined as follows (we note that this is not the original definition given in [BNS87], but a more geometrically appealing one, which was not noticed to be equivalent until several years later, see [Mei90, Theorem 3.19]).

Definition 4.7.1. Let $G = \langle X\rangle$ be a finitely generated group, and $\mathsf{Cay}(G,X)$ its Cayley graph. Then the set

$$\Sigma(G) = \{\, [\chi] \in S(G) \mid G_\chi \text{ is connected}\,\} \subseteq S(G) \tag{4.25}$$

does not depend on the choice of the finite generating set X (see [Str12]), and is called the *(first) Sigma* — or *BNS* — *invariant* of G.

Interestingly enough, this notion is quite related with commutativity. The extreme examples are free and free-abelian groups, for which it is easy to see that the BNS invariants are, respectively, the empty set and the full character sphere: $\Sigma(F_r) = \emptyset$, for $r \geqslant 2$; and $\Sigma(\mathbb{Z}^r) = S(\mathbb{Z}^r)$, for $r \geqslant 1$.

The set of characters vanishing on a certain subgroup $H \leqslant G$ determine the subsphere

$$S(G,H) = \{[\chi] \in S(G) \mid \chi(H) = 0\} \subseteq S(G),$$

which happens to contain relevant information about H itself.

Theorem 4.7.2 (Bieri, Neumann, and Strebel, 1987, [BNS87])**.** *Let H be a normal subgroup of a finitely generated group G with G/H abelian. Then, H is finitely generated if and only if $S(G,H) \subseteq \Sigma(G)$. In particular, the commutator subgroup $[G,G]$ is finitely generated if and only if $\Sigma(G) = S(G)$.*

Note that if G is H-by-$\mathbb{Z}$, then $S(G,H) = \{[\pi_H], -[\pi_H]\}$, where $\pi_H \colon G \twoheadrightarrow G/H \simeq \mathbb{Z}$ is the canonical projection modulo H. In this case, Theorem 4.7.2 tells us that

$$H \text{ is fg} \quad \Leftrightarrow \quad [\pi_H], -[\pi_H] \in \Sigma(G). \tag{4.26}$$

It follows a characterization of fg-by-$\mathbb{Z}$ groups which is directly connected with our undecidability result in Corollary 4.6.4.

Proposition 4.7.3. *A finitely generated group G is fg-by-$\mathbb{Z}$ if and only if its BNS invariant contains a pair of antipodal points; i.e.,*

$$G \textit{ is } \text{fg-by-}\mathbb{Z} \quad \Leftrightarrow \quad \exists [\chi] \in S(G) \textit{ s.t. } [\chi], -[\chi] \in \Sigma(G). \tag{4.27}$$

Proof. The implication to the right is clear from (4.26).

The implication to the left follows from the fact that we can always choose such a character χ with cyclic image (i.e., such that $\operatorname{rk}_{\mathbb{Z}} \chi(G) = 1$). To see this, we observe that, given a nontrivial character $\chi \colon G \to \mathbb{R}$, one has $\operatorname{rk}_{\mathbb{Z}} \chi(G) = 1$ if and only if there exists $\lambda > 0$ such that $\lambda\chi$ has integral image, $\lambda\chi \colon G \twoheadrightarrow \mathbb{Z} \subseteq \mathbb{R}$. In other words, rank-one characters correspond, precisely, to those points in the sphere $S(G)$ which are projections of integral (or rational) points from $\mathbb{R}^r \setminus \{0\}$. Thus, rank-one characters form a dense subset of $S(G)$. This, together with the fact that $\Sigma(G)$ is an open subset of $S(G)$ (see [Str12, Theorem A3.3]) allows us to deduce, from the hypothesis, the existence of a pair of antipodal points of rank one. □

As a corollary, and using Theorem 4.6.3, we obtain the main result in this section: the BNS invariant is not uniformly decidable.

Theorem 4.7.4. *There is no algorithm such that, given a finite presentation of a group G and a character $[\chi] \in S(G)$, decides whether $[\chi]$ belongs to $\Sigma(G)$ or not.*

Proof. Given a finite presentation of a !-by-$\mathbb{Z}$ group G (i.e., $b(G) = 1$, and $S(G)$ has two points), we can abelianize and construct the unique two characters $\pm\pi \colon G \twoheadrightarrow \mathbb{Z}$. Assuming an algorithm like in the statement, we could algorithmically decide whether $\pi, -\pi$ both belong to $\Sigma(G)$ or not, i.e., according to Proposition 4.7.3, whether G is fg-by-$\mathbb{Z}$ or not. This contradicts Theorem 4.6.3. □

We note that, in the case of a one-relator group $G = \langle a, b \mid r \rangle$, K. Brown provided an interesting algorithm for deciding whether a given character $\chi \colon G \to \mathbb{R}$ belongs to $\Sigma(G)$ or not, by looking at the sequence of χ-images of the prefixes of the relation r (assumed to be in cyclically reduced form); see [Bro87]. Later, N. Dunfield, J. Button and D. Thurston found applications of this result to 3-manifold theory; see [Dun01; But05; DT06].

4.8 Recursive enumerability of presentations

A *standard* presentation of a given fp-by-$\mathbb{Z}$ group G has been defined as a *finite* presentation of the form

$$\left\langle X, t \;\middle|\; R, txt^{-1} = \alpha(x)\ (x \in X) \right\rangle,$$

where α is an automorphism of $\langle X \mid R \rangle$. It is natural to ask for an algorithm to compute one —or all— standard presentations for such a group G, since this algorithm will provide explicit computable ways to think G as a semidirect product (i.e., an explicit base group H, and an explicit automorphism α, such that $G \simeq H \rtimes_\alpha \mathbb{Z}$).

We have seen that membership for fp-by-$\mathbb{Z}$ is undecidable (Corollary 4.6.4). However, given a finite presentation for a fp-by-$\mathbb{Z}$ group G, we can use Tietze transformations to obtain a recursive enumeration of all the finite presentations for G. In the following proposition we provide a (brute force) filtering process which extracts from it a recursive enumeration of all the standard ones.

Proposition 4.8.1. *Given a finite presentation of a* fp-by-$\mathbb{Z}$ *group* G, *the set of standard presentations for* G *is recursively enumerable.*

Proof. Let P be the finite presentation given (of a fp-by-$\mathbb{Z}$ group G). We will start enumerating all finite presentations of G by successively applying to P chains of elementary Tietze transformations in all possible ways. This process is recursive and eventually visits all finite presentations for G (all standard presentations among them).

Now, it will be enough to construct a recognizing subprocess $\mathfrak{S}$ which, applied to any finite presentation P' for G, if P' is in standard form it halts and returns P', and if not it maybe halts returning "NO, P' is not standard", or works forever. Having $\mathfrak{S}$, we can keep following the enumeration of all finite presentations P' for G via Tietze transformations and, for each one, start and run in parallel the recognizing process $\mathfrak{S}$ for it; we maintain all of them running in parallel (some of them possibly forever), and at the same time we keep opening new ones, simultaneously aware of the possible halts (each one killing one of the parallel processes and possibly outputting a genuine standard presentation for G).

So, we are reduced to design such a recognizing process $\mathfrak{S}$. For a given finite presentation P' of G, let us perform the following steps:

i. Check whether P' matches the scheme

$$\left\langle X, t \;\middle|\; R,\ tx_i t^{-1} = w_i \;\; (x_i \in X) \right\rangle, \tag{4.28}$$

where $X = \{x_1, \ldots, x_n\}$ and $R = \{r_1, \ldots, r_m\}$ are finite, and the w_i's and r_j's are all (reduced) words on X. If P does not match this scheme, then halt and answer "NO, P' is not standard"; otherwise go to the next step.

ii. With P' being of the form (4.28), consider the group $H = \langle X \mid R \rangle = F(X)/\langle\langle R \rangle\rangle$ and let us try to check whether the map $x_i \mapsto w_i$ extends to a well-defined homomorphism $\alpha \colon H \to H$. For this, we must check whether $\alpha(r_j) = 1$ in H or not (but caution! we cannot assume in general a solution to the word problem for H). Enumerate and reduce the elements in $\langle\langle R \rangle\rangle$ and check whether, for every relator $r_j(x_1, \ldots, x_n) \in R$, the word $r_j(w_1, \ldots, w_n)$ appears in the enumeration. If this happens for all $j = 1, \ldots, m$, then go to the next step (with P' being of the form

$$\left\langle X, t \;\middle|\; R,\ tx_i t^{-1} = \alpha(x_i) \;\; (x_i \in X) \right\rangle, \tag{4.29}$$

where $\alpha \in \operatorname{End}(F(X)/\langle\langle R \rangle\rangle)$).

iii. With P' being of the form (4.29), let us try to check now whether α is bijective, looking by brute force for its eventual inverse: enumerate all possible n-tuples $(v_1, \ldots, v_n)$ of reduced words on X and for each one, check simultaneously whether $r_j(v_1, \ldots, v_n) = 1$ in H for all $j = 1, \ldots, m$ (i.e., whether $\beta \colon H \to H$, $x_i \mapsto v_i$ is a well-defined endomorphism of H) and whether $v_i(w_1, \ldots, w_n) = 1$ and $w_i(v_1, \ldots, v_n) = 1$ in H, for all $i = 1, \ldots, n$ (i.e., whether $\alpha\beta = \beta\alpha = \mathrm{id}$ and so $\alpha \in \operatorname{Aut}(H)$). We do this in a similar way as in the previous step: enumerate the normal closure $\langle\langle R \rangle\rangle$ (an infinite process) and wait until all the mentioned words appear in the enumeration. When this happens (if so), halt the process and output P' as a standard presentation for G.

For any given P', step (i) finishes in finite time and either rejects P', or recognizes that P' is of the form (4.28) and sends the control to step (ii). Now step (ii) either works forever, or it halts recognizing P' of the form (4.29) and sending the control to step (iii) (note that, by construction, it is guaranteed that if P' is really in standard form then α is a well-defined endomorphism of H and step (ii) will eventually halt in finite time). Finally, the same happens in step (iii): it either works forever, or it halts recognizing that P' is in standard form (again by construction, it is guaranteed that if P' is really in standard form then α is bijective and step (iii) will eventually catch its inverse and halt in finite time. Process $\mathfrak{S}$ is built and this concludes the proof. □

We remark that we can apply the previous algorithm to an arbitrary finite presentation P of a (arbitrary) group G: if G is a fp-by-$\mathbb{Z}$ group the process will enumerate all its standard presentations, while if G is not fp-by-$\mathbb{Z}$ the process will work forever outputting nothing. So, we can successively apply —in parallel— the previous algorithm to any enumerable family $\mathcal{H}$ of presentations to obtain an enumeration of all standard fp-by-$\mathbb{Z}$ presentations within $\mathcal{H}$. Taking $\mathcal{H} = \mathcal{G}_{\mathsf{fp}}$, we get an enumeration of all standard fp-by-$\mathbb{Z}$ presentations.

Corollary 4.8.2. *The set of standard presentations of* fp-by-$\mathbb{Z}$ *groups is recursively enumerable.* □

Applying all possible Tietze transformations to every standard presentation outputted by this procedure, we obtain an enumeration of all finite presentations of fp-by-$\mathbb{Z}$ groups. This enriches Corollary 4.6.4 in the following way.

Corollary 4.8.3. *The set of finite presentations of* fp-by-$\mathbb{Z}$ *groups is recursively enumerable but not recursive.* □

4.9 Isomorphisms of $\mathbb{Z}$-extensions

Let us consider now problems of the first kind mentioned in Section 4.1: isomorphism problems within families of the form $[\mathcal{P}\text{-by-}\mathbb{Z}]_{\mathsf{fp}}$.

To begin with, we combine Lemma 4.6.1 with the following one, in order to see that a $\mathbb{Z}$-extension can have non-isomorphic base groups. The proof is just a direct writing of the corresponding presentations.

Lemma 4.9.1. *Let* H *be an arbitrary group, and* $\phi \in \mathrm{Aut}(H)$. *Then,*

$$(H \rtimes_{\phi} \mathbb{Z}) \times \mathbb{Z} \simeq (H \times \mathbb{Z}) \rtimes_{\Phi} \mathbb{Z}, \tag{4.30}$$

where $\Phi \in \mathrm{Aut}(H \times \mathbb{Z})$ *is defined by* $(h, t) \mapsto (\phi(h), t)$. □

Corollary 4.9.2. *Isomorphic* $\mathbb{Z}$*-extensions can have non-isomorphic base groups, even of different type. More precisely, there exist a finitely presented group* H*, a non-finitely generated group* H'*, and automorphisms* $\alpha \in \mathrm{Aut}(H)$ *and* $\beta \in \mathrm{Aut}(H')$*, such that* $H \rtimes_{\alpha} \mathbb{Z} \simeq H' \rtimes_{\beta} \mathbb{Z}$. *In particular,* $H \rtimes \mathbb{Z} \simeq H' \rtimes \mathbb{Z} \not\Rightarrow H \simeq H'$.

Proof. Let K be any nontrivial finitely presented group. Consider $H = K * \mathbb{Z}$, which is also finitely presented, and $H' = (\ast_{i \in \mathbb{Z}} K) \times \mathbb{Z}$, which is not finitely generated. Combining (4.19) and (4.30), we get

$$H \times \mathbb{Z} = (K * \mathbb{Z}) \times \mathbb{Z} \simeq \left(\left(\mathop{\ast}_{i \in \mathbb{Z}} K \right) \rtimes_{\tau} \mathbb{Z} \right) \times \mathbb{Z} \simeq \left(\left(\mathop{\ast}_{i \in \mathbb{Z}} K \right) \times \mathbb{Z} \right) \rtimes_{T} \mathbb{Z} = H' \rtimes_{T} \mathbb{Z},$$

where $\tau \in \operatorname{Aut}(\ast_{i \in \mathbb{Z}} K)$ is the automorphism (4.20) defined in Lemma 4.6.1, and $T \in \operatorname{Aut}(H')$ the corresponding one according to Lemma 4.9.1. The result follows taking $\alpha = \mathrm{id}_H$ and $\beta = T$. □

So, there is considerable flexibility in describing cyclic extensions as semidirect products. Even fixing the base group, this flexibility persists within the possible defining automorphisms. For example, one can easily see that $H \rtimes_{\gamma} \mathbb{Z} \simeq H \times \mathbb{Z}$, for every inner automorphism $\gamma \in \operatorname{Inn}(H)$. A bit more generally, the following is a folklore lemma providing sufficient conditions for two automorphisms (of the same base group) to define isomorphic semidirect products by $\mathbb{Z}$.

Lemma 4.9.3. *Let H be an arbitrary group, and let $\alpha, \beta \in \operatorname{Aut}(H)$. If $\beta = \gamma \xi \alpha^{\pm 1} \xi^{-1}$ for some $\gamma \in \operatorname{Inn}(H)$ and some $\xi \in \operatorname{Aut}(H)$, then $H \rtimes_{\alpha} \mathbb{Z} \simeq H \rtimes_{\beta} \mathbb{Z}$.* □

The existence of such $\gamma \in \operatorname{Inn} H$ and $\xi \in \operatorname{Aut}(H)$ is exactly the same as $[\beta]$ being conjugate to $[\alpha]^{\pm 1}$ in $\operatorname{Out}(H)$. This condition turns out to have some protagonism along the rest of the chapter, making convenient to have a general shorthand terminology for it.

Definition 4.9.4. Let G be an arbitrary group. A pair of elements $g, h \in G$ are said to be *semi-conjugate* if g is conjugate to either h or h^{-1}; we denote this situation by $g \sim h^{\pm 1}$.

With this terminology, Lemma 4.9.3 is saying that when the defining automorphisms $\alpha, \beta \in \operatorname{Aut}(H)$ are semi-conjugate in $\operatorname{Out}(H)$, then the corresponding semidirect products $H \rtimes_{\alpha} \mathbb{Z}$ and $H \rtimes_{\beta} \mathbb{Z}$ are isomorphic. Note also the following necessary condition: by Proposition 4.5.5, α is deranged if and only if $b(H \rtimes_{\alpha} \mathbb{Z}) = 1$ so, in order for $H \rtimes_{\alpha} \mathbb{Z}$ and $H \rtimes_{\beta} \mathbb{Z}$ to be isomorphic, a necessary condition is that α and β are either both simultaneously deranged, or both not deranged.

Apart from this, not much is known in general about characterizing or deciding when two $\mathbb{Z}$-extensions of a given group are isomorphic. In [BMV07], Bogopolski–Martino–Ventura proved that, when the base group H is free of rank 2, the converse to Lemma 4.9.3 also holds, providing a quite neat characterization of isomorphism within the family of F_2-by-$\mathbb{Z}$ extensions and (using the decidability of the conjugacy problem in $\operatorname{Out}(F_2)$, see [Bog00]) a positive solution to the isomorphism problem within this family of groups.

Theorem 4.9.5 (Bogopolski–Martino–Ventura, [BMV07]). *Let $\alpha, \beta \in \mathrm{Aut}(F_2)$. Then,*

$$F_2 \rtimes_\alpha \mathbb{Z} \simeq F_2 \rtimes_\beta \mathbb{Z} \;\Leftrightarrow\; [\alpha] \text{ and } [\beta] \text{ are semi-conjugate in } \mathrm{Out}(F_2). \tag{4.31}$$

In particular, the isomorphism problem within the family F_2 - by - $\mathbb{Z}$ is decidable. □

However, in this same paper, a counterexample was given (suggested by Dicks) to see that this equivalence is not true for free groups of higher rank, where the situation is, in general, much more complicated. The example is the following: consider the free group of rank 3, $F_3 = \langle a, b, c \mid \rangle$, and the automorphisms $\alpha\colon F_3 \to F_3$ given by $a \mapsto b \mapsto c \mapsto b^{-1}ab^{-2}c^3$, and $\beta\colon F_3 \to F_3$ by $a \mapsto b \mapsto c \mapsto a^{-1}b^2cb^{-1}$. It happens that $F_3 \rtimes_\alpha \mathbb{Z} \simeq F_3 \rtimes_\beta \mathbb{Z}$ (see [BMV07] for details), while α and β are not semi-conjugate in $\mathrm{Out}(F_3)$ because they abelianize to two 3×3 matrices of determinants, respectively, 1 and -1. As far as we know, the isomorphism problem for F_r - by - $\mathbb{Z}$ groups is open for $r \geqslant 3$.

The goal of the present section is to prove that an equivalence like (4.31) still holds, but under kind of an orthogonal condition: rather than restricting the base group to be F_2, we will leave H arbitrary finitely generated, and impose conditions on the defining automorphism. Note that such an equivalence reduces the isomorphism problem in the family of restricted extensions, to the conjugacy problem in the corresponding family of outer automorphisms of the base group (or even to a weaker problem, if semi-conjugacy is not algorithmically equivalent to conjugacy).

This context strongly suggests defining the semi-conjugacy problem much in the same way that the standard conjugacy problem, and asking for the relationship between them. We state both problems together in order to make the comparison clear.

Definition 4.9.6. Let $\langle X \mid R \rangle$ be a finite presentation for a group G. Then:

- *Conjugacy Problem* for G [$\mathrm{CP}(G)$]: given two words u, v in $X^{\pm}$, decide whether they represent conjugate elements in G ($u \sim_G v$) or not.
- *Semi-conjugacy Problem* for G [$\frac{1}{2}\mathrm{CP}(G)$]: given two words u, v in $X^{\pm}$, decide whether they represent semi-conjugate elements in G ($u \sim_G v^{\pm 1}$) or not.

Question 1. *Is there a (finitely presented) group with decidable semi-conjugacy problem but undecidable conjugacy problem?*

This question looks quite tricky. Of course, if two elements $g, h \in G$ are not semi-conjugate, then they are not conjugate either. But if $g \sim h^{-1}$, it is not clear how this information can help, in general, to decide whether $g \sim h$ or not; in this sense the answer to the question seems reasonable to be negative. But, on the other hand, the two algorithmic problems are so close that it seems hard to construct a counterexample.

In our case, the condition demanded for the defining automorphisms is derangedness (see 4.5.6). The first observation is the following: suppose $H \rtimes_\alpha \mathbb{Z} \simeq K \rtimes_\beta \mathbb{Z}$ for some groups H, K, and some deranged automorphisms $\alpha \in \mathrm{Aut}(H)$ and $\beta \in \mathrm{Aut}(K)$. Then, by construction, H and K are respectively, the unique normal subgroups with quotient isomorphic to $\mathbb{Z}$. Hence $H \simeq K$ and, after expressing β in terms of the generators of H, we can think that both $\alpha, \beta \in \mathrm{Aut}(H)$. The next step is to show that, under the derangedness condition, $H \rtimes_\alpha \mathbb{Z} \simeq H \rtimes_\beta \mathbb{Z}$ implies that $[\alpha], [\beta] \in \mathrm{Out}(H)$ are semi-conjugate. To see this, we need to analyze how homomorphisms between unique $\mathbb{Z}$-extensions look like.

Definition 4.9.7. Let G be a group, and H a subgroup of G. An endomorphism $\Psi \in \mathrm{End}(G)$ is called H-*stable* if $\Psi(H) \leqslant H$ (i.e., if H is invariant under Ψ). The collection of all H-stable endomorphisms of G form a submonoid denoted $\mathrm{End}_H(G) \leqslant \mathrm{End}(G)$. In a similar way, the collection of all H-stable automorphisms of G form a normal subgroup denoted $\mathrm{Aut}_H(G) \trianglelefteqslant \mathrm{Aut}(G)$.

Notation. We will use Greek capital letters $(\Phi, \Psi, \ldots)$ to denote homomorphisms of the extensions $H \rtimes \mathbb{Z}$, and keep using lowercase Greek letters $(\phi, \psi, \alpha, \beta, \ldots)$ to denote homomorphisms of the base group H. Moreover, when a homomorphism $\Phi \in \mathrm{End}(H \rtimes \mathbb{Z})$ is H-stable, we will usually denote with the *same letter in lowercase* the endomorphism induced by Φ in H; i.e., $\phi := \Phi_{|H} \colon H \ni h \mapsto \Phi(h)$.

A general description of the H-stable endomorphisms and automorphisms of infinite-cyclic extensions of H follows.

Proposition 4.9.8. *Let H be a group generated by $X = \{x_i \mid i \in I\}$, and let $\alpha, \beta \in \mathrm{Aut}(H)$. Then, any homomorphism from $H \rtimes_\alpha \mathbb{Z}$ to $H \rtimes_\beta \mathbb{Z}$ mapping H to H is of the form*

$$\begin{aligned} \Phi_{\epsilon,\phi,h_0} \colon H \rtimes_\alpha \mathbb{Z} &\to H \rtimes_\beta \mathbb{Z}, \\ x_i &\mapsto \phi(x_i) \\ t &\mapsto h_0\, t^\epsilon \end{aligned} \tag{4.32}$$

where $\epsilon \in \mathbb{Z}$, $h_0 \in H$, and $\phi \in \mathrm{End}(H)$ are such that $\gamma_{h_0} \beta^\epsilon \phi = \phi\alpha$.

Furthermore, Φ_{ϵ,ϕ,h_0} is an isomorphism if and only if $\epsilon = \pm 1$ and $\phi \in \mathrm{Aut}(H)$. Thus, the set of H-stable automorphisms of $H \rtimes_\alpha \mathbb{Z}$ is

$$\mathrm{Aut}_H(H \rtimes_\alpha \mathbb{Z}) = \left\{ \Phi_{\epsilon,\phi,h_0} \mid \epsilon = \pm 1,\ h_0 \in H,\ \phi \in \mathrm{Aut}(H) \text{ s.t. } \gamma_{h_0} \alpha^\epsilon \phi = \phi\alpha \right\}. \tag{4.33}$$

Proof. Let $\Phi \colon H \rtimes_\alpha \mathbb{Z} \to H \rtimes_\beta \mathbb{Z}$ be a homomorphism leaving H invariant, and let us denote by $\phi \colon H \to H$ its restriction to H. Write $\Phi(t) = h_0\, t^\epsilon$ for some $h_0 \in H$

and $\epsilon \in \mathbb{Z}$. Applying Φ to both sides of the relation $tht^{-1} = \alpha(h)$ in the domain, we get

$$h_0 \cdot \beta^{\epsilon}\phi(h) \cdot h_0^{-1} = h_0 t^{\epsilon} \cdot \phi(h) \cdot t^{-\epsilon} h_0^{-1} = \Phi(tht^{-1}) = \Phi(\alpha(h)) = \phi\alpha(h)\,,$$

for all $h \in H$. Hence, $\gamma_{h_0}\beta^{\epsilon}\phi = \phi\alpha$ and $\Phi = \Phi_{\epsilon,\phi,h_0}$ has the desired form.

Assume now that Φ_{ϵ,ϕ,h_0} is an isomorphism (in particular, $\phi\colon H \to H$ is injective). Then we must have $\epsilon = \pm 1$, otherwise t would not be in the image. On the other hand, since $H \trianglelefteq H \rtimes_{\alpha} \mathbb{Z}$, we have that

$$\phi(H) = \Phi_{\epsilon,\phi,h_0}(H) \trianglelefteq H \rtimes_{\beta} \mathbb{Z} = \Phi_{\epsilon,\phi,h_0}(H \rtimes_{\alpha} \mathbb{Z}) = \langle \phi(H),\, h_0 t^{\epsilon} \rangle\,,$$

and so, any element of $H \rtimes_{\beta} \mathbb{Z}$ can be written in the form $\phi(h)(h_0 t^{\epsilon})^k$, for some $h \in H$ and $k \in \mathbb{Z}$; and it belongs to H if and only if $k = 0$. Thus, $\Phi(H) = H$ and $\phi \in \mathrm{Aut}(H)$. For the converse, it is clear that $\epsilon = \pm 1$ and $\phi \in \mathrm{Aut}(H)$ implies that Φ_{ϵ,ϕ,h_0} is an isomorphism. The final statement follows immediately. □

4.10 Isomorphisms of unique $\mathbb{Z}$-extensions

Note that Proposition 4.5.5 states precisely that $\mathrm{End}_H(H \rtimes_{\alpha} \mathbb{Z}) = \mathrm{End}(H \rtimes_{\alpha} \mathbb{Z})$ if and only if α is deranged. This fact, together with the previous description provides a characterization of isomorphic deranged extensions in terms of semi-congugacy.

Corollary 4.10.1. *Let H and K be two arbitrary groups, and let $\alpha \in \mathrm{Aut}(H)$ and $\beta \in \mathrm{Aut}(K)$ be two deranged automorphisms. Then,*

$$H \rtimes_{\alpha} \mathbb{Z} \simeq K \rtimes_{\beta} \mathbb{Z} \;\; \Leftrightarrow \;\; H \simeq K \text{ and } [\alpha] \sim [\beta']^{\pm 1} \text{ in } \mathrm{Out}(H), \tag{4.34}$$

where $\beta' = \psi^{-1}\beta\psi \in \mathrm{Aut}(H)$, and $\psi\colon H \to K$ is any isomorphism.

Proof. For any isomorphism $\psi\colon H \to K$, it is clear that $K \rtimes_{\beta} \mathbb{Z} = \psi(H) \rtimes_{\beta} \mathbb{Z} \simeq H \rtimes_{\psi^{-1}\beta\psi} \mathbb{Z}$. Hence, the statement is equivalent to saying

$$H \rtimes_{\alpha} \mathbb{Z} \simeq H \rtimes_{\beta} \mathbb{Z} \;\; \Leftrightarrow \;\; [\alpha] \sim [\beta]^{\pm 1} \text{ in } \mathrm{Out}(H),$$

for $\alpha, \beta \in \mathrm{Aut}(H)$. The implication to the left is a general fact (see Lemma 4.9.3), and the implication to the right is a direct consequence of Proposition 4.9.8: since α and β are deranged, any isomorphism from $H \rtimes_{\alpha} \mathbb{Z}$ to $H \rtimes_{\beta} \mathbb{Z}$ must map H to H and so, must be of the form Φ_{ϵ,ϕ,h_0} for some $\epsilon = \pm 1$, $h_0 \in H$, and $\phi \in \mathrm{Aut}(H)$ satisfying $\gamma_{h_0}\beta^{\epsilon}\phi = \phi\alpha$. Hence, $[\alpha] \sim [\beta]^{\pm 1}$ in $\mathrm{Out}(H)$. □

We are now ready to prove the main result in this section: for any family $\mathcal{H}$ of finitely presented groups with decidable isomorphism problem, we characterize when the family $[!\mathcal{H}\text{-by-}\mathbb{Z}]_{\mathsf{fp}}$ has again decidable isomorphism problem, in terms of a certain variation of the conjugacy problem for outer automorphisms of groups in $\mathcal{H}$.

Note that Corollary 4.10.1 clearly insinuates a link between the isomorphism problem for deranged extensions, and the semi-conjugacy problem for deranged outer automorphisms of the base group. However, there is a subtlety at this point: the supposed algorithm solving the isomorphism problem will receive the input (the compared groups) as finite presentations of the $\mathbb{Z}$-extensions. From those, we know how to compute suitable base groups H, K, and automorphisms α, β (see Proposition 4.8.1), but this last ones are given *by images of the generators in the starting presentations*, and not as words in some presentation of the corresponding automorphism groups, which would be the appropriate inputs for the standard conjugacy problem there.

So, in general, one must distinguish between these two close but not necessarily identical situations. As before, we state both problems together to emphasize the difference between them.

Definition 4.10.2. Let $\langle X \mid R \rangle$ be a presentation for a group G, $\langle Y \mid S \rangle$ a presentation for $\mathrm{Aut}(G)$, and assume $|X| < \infty$. Then:

- *(Standard) conjugacy problem* for $\mathrm{Aut}(G)$ [$\mathsf{CP}(\mathrm{Aut}(G))$]: given two automorphisms $\alpha, \beta \in \mathrm{Aut}(G)$ *as words in the presentation of* $\mathrm{Aut}(G)$, decide whether α and β are conjugate to each other in $\mathrm{Aut}(G)$.
- Aut-*conjugacy problem* for G [$\mathsf{CP}_G(\mathrm{Aut}(G))$]: given two automorphisms $\alpha, \beta \in \mathrm{Aut}(G)$ *by images of (the finitely many) generators* X, decide whether α and β are conjugate to each other in $\mathrm{Aut}(G)$.

Similarly, we define the Out-*conjugacy problem* [$\mathsf{CP}_G(\mathrm{Out}(G))$], the Aut-*semi-conjugacy problem* [$\frac{1}{2}\mathsf{CP}_G(\mathrm{Aut}(G))$], and the Out-*semi-conjugacy problem* [$\frac{1}{2}\mathsf{CP}_G(\mathrm{Out}(G))$] for G (in contrast with the standard $\mathsf{CP}(\mathrm{Out}(G))$, $\frac{1}{2}\mathsf{CP}(\mathrm{Aut}(G))$, and $\frac{1}{2}\mathsf{CP}(\mathrm{Out}(G))$).

Note that, in general, these pairs of problems are similar but not identical: from the algorithmic point of view it could be different to have an automorphism of G given as the collection of images of a finite set of generators of G, or as a word (composition of generators for $\mathrm{Aut}(G)$). Consider, for example, the Baumslag–Solitar group $G = \mathrm{BS}(2,4)$, which is finitely generated, but whose automorphism group $\mathrm{Aut}(G)$ is known to be not finitely generated (see [CL83]).

However, knowing in advance a finite set of generators for $\mathrm{Aut}(G)$ (respectively, $\mathrm{Out}(G)$) as images of generators of G, these two kinds of problems turn out to be equivalent.

Proposition 4.10.3. *Let* $\langle X \mid R\rangle$ *be a presentation for a group* G, $X = \{x_1, \ldots, x_n\}$, *and* $\{u_{i,j} \mid i = 1, \ldots, n,\ j = 1, \ldots, N\}$ *a finite set of words in* $X^{\pm}$ *such that* $\{\alpha_j : x_i \mapsto u_{i,j} : j = 1, \ldots, N\}$ *is a well defined finite family of automorphisms generating* $\mathrm{Aut}(G)$. *Then,*

$$\mathsf{CP}_G(\mathrm{Aut}(G)) \text{ is decidable } \Leftrightarrow \mathsf{CP}(\mathrm{Aut}(G)) \text{ is decidable}. \tag{4.35}$$

The same is true replacing conjugacy by semi-conjugacy, or Aut *by* Out.

Proof. Suppose that G has decidable Aut-conjugacy problem. Given two automorphisms $\alpha, \beta \in \mathrm{Aut}(G)$ as words on the α_i's, say $\alpha = a(\alpha_1, \ldots, \alpha_N)$ and $\beta = b(\alpha_1, \ldots, \alpha_N)$, we can compute the corresponding compositions of α_j's and obtain explicit expressions for $\alpha(x_i)$ and $\beta(x_i)$ in terms of X, for $i = 1, \ldots, n$. Now, applying the solution to the Aut-conjugacy problem for G we decide whether α and β are conjugate to each other in $\mathrm{Aut}(G)$.

Conversely, suppose $\mathrm{Aut}(G)$ has decidable conjugacy problem, and we are given two automorphisms $\alpha, \beta \in \mathrm{Aut}(G)$ by the images of the x_i's, say $\alpha(x_i)$ and $\beta(x_i)$, $i = 1, \ldots, n$. We will express α and β as compositions of the α_j's, and then apply the assumed solution to the conjugacy problem for $\mathrm{Aut}(G)$ to decide whether α and β are conjugate to each other, or not. We can do this by a brute force enumeration of all possible formal reduced words w on $\alpha_1, \ldots, \alpha_N$ and, for each one, computing the tuple $(w(x_1), \ldots, w(x_n))$ and trying to check whether it equals $(\alpha(x_1), \ldots, \alpha(x_n))$, or $(\beta(x_1), \ldots, \beta(x_n))$ (following a brute force enumeration of the normal closure $\langle\!\langle R\rangle\!\rangle$, like in the proof of Proposition 4.8.1).

The proofs of the other versions of the statement are completely analogous. For the Out cases we need to add another brute force search layer enumerating all possible conjugators; we leave details to the reader. □

After this proposition we can prove the main result in this section.

Theorem 4.10.4. *Let* $\mathcal{H}$ *be a family of finitely presented groups with decidable isomorphism problem. Then, the isomorphism problem of* $!\mathcal{H}$-by-$\mathbb{Z}$ *is decidable if and only if the* Out_d-*semi-conjugacy problem of* H *is decidable for every* H *in* $\mathcal{H}$*; i.e.,*

$$\mathsf{IP}(!\mathcal{H}\text{-by-}\mathbb{Z}) \text{ decidable } \Leftrightarrow \tfrac{1}{2}\mathsf{CP}_H(\mathrm{Out}_d(H)) \text{ decidable}, \forall H \in \mathcal{H}.$$

Proof. Suppose that every $H \in \mathcal{H}$ has decidable Out-semi-conjugacy problem for deranged inputs. Given finite presentations of two groups G and G' in $!\mathcal{H}$-by-$\mathbb{Z}$, we run Proposition 4.8.1 to compute standard presentations for them, and extract finite presentations for base groups and defining automorphisms (say H and $\alpha \in \mathrm{Aut}(H)$ for G, and K and $\beta \in \mathrm{Aut}(K)$ for G', respectively). We have $G = H \rtimes_\alpha \mathbb{Z}$ and $G' = K \rtimes_\beta \mathbb{Z}$ and, by hypotheses, α and β are deranged. Furthermore, since $b(G) = b(G') = 1$, H and K are the unique normal subgroups of G and G', respectively, with quotient $\mathbb{Z}$; hence, $H, K \in \mathcal{H}$.

Now we apply the isomorphism problem within $\mathcal{H}$ to the obtained presentations for H and K, and decide whether they are isomorphic as groups. If $H \not\simeq K$ then, by Corollary 4.10.1, $G \not\simeq G'$ and we are done. Otherwise, we construct an explicit isomorphism $\psi \colon H \to K$ (by a brute force search procedure like the ones above), we compute $\beta' = \psi^{-1}\beta\psi \in \mathrm{Aut}(H)$, and we apply our solution to the Out-semi-conjugacy problem for $H \in \mathcal{H}$ to the inputs α and β', (which are deranged, by construction). The output on whether $[\alpha]$ and $[\beta']$ are or are not semi-conjugate in $\mathrm{Out}(H)$ is the final answer we are looking for (again by Corollary 4.10.1).

For the converse, assume that the isomorphism problem is decidable in the family $!\mathcal{H}$-by-$\mathbb{Z}$, and fix a finite presentation $\langle X \mid R \rangle$ for a group $H \in \mathcal{H}$. Given two deranged automorphisms $\alpha, \beta \in \mathrm{Out}(H)$ via images of the generators $x_i \in X$, build the corresponding standard presentations for $H \rtimes_\alpha \mathbb{Z}$ and $H \rtimes_\beta \mathbb{Z}$ (which are groups in $!\mathcal{H}$-by-$\mathbb{Z}$, by construction) and apply the assumed solution to the isomorphism problem for this family, to decide whether they are isomorphic or not. By Corollary 4.10.1, the answer is affirmative if and only if $[\alpha]$ and $[\beta]$ are semi-conjugate in $\mathrm{Out}(H)$. □

We apply now Theorem 4.10.4 to special families of groups with decidable isomorphism problem. Some of these corollaries are already known in the literature; other methods provide alternative approaches.

Taking $\mathcal{H}$ to be a single group H, we get the following result.

Corollary 4.10.5. *Let H be a finitely presented group. Then the isomorphism problem is decidable within the family $!H$-by-$\mathbb{Z}$ if and only if H has decidable Out-semi-conjugacy problem for deranged inputs. In particular, if $|\mathrm{Out}(H)| < \infty$, then $!H$-by-$\mathbb{Z}$ has decidable isomorphism problem.* □

Taking $\mathcal{H}$ to be the families of finite, finitely generated abelian, or polycyclic groups, Theorem 4.10.4 provides well-known results, since the obtained extensions turn out to be subfamilies of that of virtually-polycyclic groups for which the isomorphism problem is known to be decidable (see [Seg90]).

For the family of Braid groups $\mathcal{B} = \mathsf{braid} = \{B_n \mid n \geqslant 2\}$, the specially simple structure of its outer automorphism group allows us to state the isomorphism problem within the family $\mathcal{B}$-by-$\mathbb{Z}$.

Corollary 4.10.6. *The isomorphism problem is decidable within the family* $\mathcal{B}$ - by - $\mathbb{Z}$.

Proof. It is well known that, for every $n \geqslant 2$, $\mathrm{Out}(B_n) = \{\overline{\mathrm{id}}, \bar{\iota}\}$, where $\iota \colon B_n \to B_n$ is the automorphism given by $\sigma_i \mapsto \sigma_i^{-1}$ (see [DG81]). Then, (since automorphisms in the same inner class induce the same $\mathbb{Z}$-extension, by Lemma 4.9.3), we have

$$\mathcal{B}\text{-by-}\mathbb{Z} = \{B_n \times \mathbb{Z}, n \geqslant 2\} \cup \{B_n \rtimes_\iota \mathbb{Z}, n \geqslant 2\}. \tag{4.36}$$

We claim that they are all pairwise non-isomorphic. Those in the first term of the union cannot be isomorphic to those in the second one because id is not deranged, while ι is (in other words, $b(B_n \times \mathbb{Z}) = 2$ while $b(B_n \rtimes_\iota \mathbb{Z}) = 1$). Two deranged extensions $B_n \rtimes_\iota \mathbb{Z}$ and $B_m \rtimes_\iota \mathbb{Z}$ can only be isomorphic if their base groups are, and this happens only when $n = m$ (this can be seen, for example, by observing that the center $Z(B_n)$ is generated by the full twist $(\sigma_1\sigma_2 \cdots \sigma_{n-1})^n$ and so, the abelianization of $B_n/Z(B_n)$ is cyclic of order $n(n-1)$). Finally, the same argument shows that $B_n \times \mathbb{Z} \simeq B_m \times \mathbb{Z}$ if and only if $n = m$.

Thus, the isomorphism problem within $\mathcal{B}$ - by - $\mathbb{Z}$ is decidable: given two finite presentations of groups G_1 and G_2 in $\mathcal{B}$ - by - $\mathbb{Z}$, explore the two trees of Tietze transformations until finding standard presentations for them, i.e., until recognizing their number of strands, say n and m. Now $G_1 \simeq G_2$ if and only if $b(G_1) = b(G_2)$ and $n = m$. □

Finally, let us consider the case of finitely generated free groups, $\mathcal{F} = [\mathsf{free}]_{\mathsf{fg}} = \{F_n \mid n \geqslant 0\}$. To start with, the isomorphism problem for $\mathcal{F}$ is decidable like in the case of Braid groups (since $F_n \simeq F_m \Leftrightarrow n = m$). A solution to the conjugacy problem in $\mathrm{Out}(F_n)$ was announced by M. Lustig in the preprints [Lus00; Lus01]. Although this project is not completed (and there is no published version yet), it is believed that $\mathrm{Out}(F_n)$ has decidable conjugacy problem. However, at this moment we can only say to have firm complete solutions for some classes of outer automorphisms:

i. the case of rank 2 is easily decidable because $\mathrm{Out}(F_2) \simeq \mathrm{GL}_2(\mathbb{Z})$;

ii. for finite-order elements of $\mathrm{Out}(F_n)$ an algorithm to solve the conjugacy problem follows from results of S. Krstić (see [Krs89]);

iii. J. Los and, independently, Z. Sela solved the conjugacy problem in $\mathrm{Out}(F_r)$ for irreducible inputs, see [Sel95; Los96; Lus07];

iv. for Dehn twist automorphisms, the conjugacy problem has been solved by Cohen–Lustig, see [CL99];

v. finally, Krstić–Lustig–Vogtmann solved the conjugacy problem in $\mathrm{Out}(F_n)$ for linearly growing automorphisms, i.e., for roots of Dehn twists, see [KLV01].

If the conjugacy problem in $\mathrm{Out}(F_n)$ were decidable in general, we could deduce from Theorem 4.10.4 that the isomorphism problem for the family $![\mathsf{free}]_{\mathsf{fg}}$ - by - $\mathbb{Z}$ is decidable as well. By the moment, we can only restrict our attention to the above mentioned subsets of $\mathrm{Out}(F_n)$, where the conjugacy problem is firmly know to be decidable, and we obtain the isomorphism problem for the corresponding subfamilies (see the proof of Theorem 4.10.4).

Corollary 4.10.7. *If the conjugacy problem for* $\mathrm{Out}(F_n)$ *is decidable, then the isomorphism problem within the family* $![\mathsf{free}]_{\mathsf{fg}}$ - by - $\mathbb{Z}$ *is also decidable.* □

Corollary 4.10.8. *The isomorphism problem within the following families is decidable:*

(i) $!F_2$ - by - $\mathbb{Z}$;

(ii) $\{F_n \rtimes_\alpha \mathbb{Z} \mid \alpha \in \mathrm{Out}(F_n)$ *deranged and finite order*$\}$;

(iii) $\{F_n \rtimes_\alpha \mathbb{Z} \mid \alpha \in \mathrm{Out}(F_n)$ *deranged and irreducible*$\}$;

(iv) $\{F_n \rtimes_\alpha \mathbb{Z} \mid \alpha \in \mathrm{Out}(F_n)$ *deranged and linearly growing*$\}$. □

It is worth mentioning that our approach is somehow opposite to that taken by Dahmani in [**dahmani_suspensions_2013**]. In this interesting preprint the author solves the conjugacy problem for atoroidal automorphisms of F_n. An automorphism $\alpha \in \mathrm{Out}(F_n)$ is *atoroidal* if no proper power of α fixes any nontrivial conjugacy class (note that this notion is similar in spirit to our notion of derangedness, though they do not coincide). Brinkmann proved in [Bri00] that $F_n \rtimes_\alpha \mathbb{Z}$ is hyperbolic if and only if α is atoroidal. And $\alpha, \beta \in \mathrm{Out}(F_n)$ are conjugate to each other if and only if $F_n \rtimes_\alpha \mathbb{Z}$ is isomorphic to $F_n \rtimes_\beta \mathbb{Z}$ with an automorphism mapping F_n to F_n, and t to an element of the form wt^1 (i.e., with an stable and positive automorphism in our language, see the proof of Corollary 4.10.1). Then, Dahmani uses a variation of the celebrated solution to the general isomorphism problem for hyperbolic groups (see [Sel95; DG11]) to determine whether $F_n \rtimes_\alpha \mathbb{Z}$ and $F_n \rtimes_\beta \mathbb{Z}$ are isomorphic through an isomorphism of the above type, and so deciding whether the atoroidal automorphisms α and β are conjugated to each other in $\mathrm{Out}(F_n)$. Our approach has been the opposite: we have used the conjugacy problem in $\mathrm{Out}(F_n)$ (more precisely, those particular cases where it is known to be decidable) to solve the isomorphism problem in the corresponding families of unique $\mathbb{Z}$-extensions.

Free-abelian by free groups

5

5.1 Alphabets, words, and languages

Definition 5.1.1. An *alphabet* is a nonempty (not necessarily finite) set. The elements of an alphabet are called *symbols* or *letters*.

Definition 5.1.2. A *string* (or *word*) of *length* k in A is a (k-)tuple of elements in A. We denote by A^* the set of words in A,

$$A^* = \{\lambda\} \cup \bigcup_{n \geqslant 1} A^n ,$$

where λ denotes de the *empty word* (of *length* 0).

We will denote by $|w|$ the length of a word w.

Notation 5.1.3. We use sans serif typeface to denote both alphabets ($A, B, C, \ldots$), and letters $a, b, c, \ldots \in A$, or words $u, v, w, \ldots \in A^*$.

Notation 5.1.4. We will usually omit parentheses and commas and represent words just as a juxtaposition of letters; i.e.,

$$a_{i_1} a_{i_2} \ldots a_{i_k} \coloneqq (a_{i_1}, a_{i_2}, \ldots, a_{i_k}) .$$

Definition 5.1.5. Given two words $u = a_{i_1} a_{i_2} \ldots a_{i_k}$, $v = b_{j_1} b_{j_2} \ldots b_{j_l}$, we define the *concatenation* of u and v , denoted by uv, to be the word (of length $k + l$) obtained by juxtaposing v after u

$$uv = a_{i_1} a_{i_2} \ldots a_{i_k} b_{j_1} b_{j_2} \ldots b_{j_l} .$$

Remark 5.1.6. Concatenation is associative with neutral element λ. Hence, with the operation given by concatenation, A^* is the *free monoid* over A; and we can recursively define the concatenation of any finite sequence of words in the natural way.

Example 5.1.7. $\{a\}^* \simeq \mathbb{N}$.

Remark 5.1.8. If the alphabet A is finite, then the free monoid A^* has countable cardinal $\aleph_0$; otherwise, A and A^* have the same cardinal. So, a nontrivial free monoid is always infinite.

Notation 5.1.9. Given a word w, we will denote by $w[i]$ the i-th letter in w, and $w[i,j]$ the $[i,j]$-*segment* of w (i.e., the subword of w starting in position i and ending in position j).

We will usually write $w = w(A)$, or $w = w(\overrightarrow{a_i})$ to denote that w is a word in the alphabet $A = \{a_i\}_i$. Also, given a family of words $\{u_i\}_i$ in some other alphabet B, we will denote by $w(\overrightarrow{u_i})$ the word obtained from w by replacing each appearance of the letter a_i in w by the corresponding word $u_i \in B^*$. Formally, $w(\overrightarrow{u_i}) := \Theta(w)$, where $\Theta \colon A^* \to B^*$ is the extension to A^* of the map $a_i \mapsto u_i$.

Definition 5.1.10. A *language* over A is any subset $\mathcal{L} \subseteq A^*$.

Several natural operations can be defined on languages. We introduce below the most common ones.

Definition 5.1.11. Let $\mathcal{L}, \mathcal{L}' \subseteq A^*$. Then,

- the *union of languages* $\mathcal{L} \cup \mathcal{L}'$ is the set:
$$\mathcal{L} \cup \mathcal{L}' = \left\{ w \in A^* \,:\, w \in \mathcal{L} \text{ or } w \in \mathcal{L}' \right\};$$
- the *intersection of languages* $\mathcal{L} \cap \mathcal{L}'$ is the set:
$$\mathcal{L} \cap \mathcal{L}' = \left\{ w \in A^* \,:\, w \in \mathcal{L} \text{ and } w \in \mathcal{L}' \right\};$$
- the *difference of languages* $\mathcal{L} \setminus \mathcal{L}'$ is the set:
$$\mathcal{L} \setminus \mathcal{L}' = \left\{ w \in A^* \,:\, w \in \mathcal{L} \text{ and } w \notin \mathcal{L}' \right\};$$
- the *complementary of language* of $\mathcal{L}$ is the set:
$$\mathcal{L}^c = A^* \setminus \mathcal{L} = \{ w \in A^* \,:\, w \notin \mathcal{L} \};$$
- the *concatenation of two languages* $\mathcal{L}\mathcal{L}'$ is the set:
$$\mathcal{L}\mathcal{L}' = \left\{ ww' \in A^* \,:\, w \in \mathcal{L} \text{ and } w' \in \mathcal{L}' \right\}.$$

Finally, since the concatenation of languages is associative, the languages below are also well defined.

- The n-*th power of a language* $\mathcal{L}$ is the set:
$$\mathcal{L}^n = \mathcal{L} \overset{(n)}{\cdots} \mathcal{L} = \{ w_1 w_2 \cdots w_n \,:\, w_i \in \mathcal{L} \quad (i = 1, \ldots, n) \}, \text{ if } n \geqslant 1;$$
or the *trivial language* $\mathcal{L}^0 = \{\lambda\}$, if $n = 0$;
- the *Kleene star of a language* $\mathcal{L}$ is the set $\mathcal{L}^* = \bigcup_{n \geqslant 0} \mathcal{L}^n$.

Definition 5.1.12. A word $u \in A^*$ is said to be a *prefix* (resp. *suffix*) of a word $w \in A^*$ if there exists a word $v \in A^*$ such that $w = uv$ (resp. $w = vu$). We denote by $w[i{\to}]$ (resp. $w[{\leftarrow}i]$) the i-*th prefix* (resp. the i-*th suffix*) of w; i.e., the subword of w consisting of its first (resp. last) i letters. Finally, we denote by $w[i, j]$ the *segment subword* of w starting at the ith letter and ending at the jth letter.

Definition 5.1.13. The set of *rational languages* on a finite alphabet A, denoted by $\mathrm{Rat}(A^*)$, is the smallest set of A-languages $\mathfrak{R}$ satisfying the following conditions:

(i) $\varnothing \in \mathfrak{R}$, and $\{a\} \in \mathfrak{R}$, for all $a \in A$,

(ii) $\mathfrak{R}$ is closed under finite union, concatenation, and Kleene star; (i.e., if $\mathcal{L}, \mathcal{L}' \in \mathfrak{R}$, then $\mathcal{L} \cup \mathcal{L}', \mathcal{L}\mathcal{L}', \mathcal{L}^* \in \mathfrak{R}$).

Remark 5.1.14. Note that $A^* \in \mathrm{Rat}(A^*)$, and any intersection of sets of languages satisfying (i) and (ii), satisfies again (i) and (ii). Thus, the "smallest set" in Definition 5.1.13 is well defined.

Definition 5.1.15. Given an alphabet A, we denote by A^{-1} the *set of formal inverses* of A. Formally A^{-1} can be defined as any set A' equipotent and disjoint with A, together with a bijection $\iota\colon A \to A'$. Then, for every $a \in A$, we call $a\iota$ the *formal inverse* of a, and we denote by $a^{-1} \coloneqq a\iota$. So, we have $A^{-1} = \{a^{-1} : a \in A\}$, and $A \cap A^{-1} = \varnothing$.

Definition 5.1.16. An alphabet is said to be *involutive* if it is of the form

$$X^{\pm} \coloneqq X \sqcup X^{-1}.$$

Then, we usually extend $\iota\colon x \mapsto x^{-1}$ to an involution $\iota = {}^{-1}$ on $(X^{\pm})^*$ through

$$(x_{i_1} \cdots x_{i_k})^{-1} \coloneqq x_{i_k}^{-1} \cdots x_{i_1}^{-1},$$

and call the resulting pair $((X^{\pm})^*, {}^{-1})$, the *involutive free monoid* on X.

Definition 5.1.17. A word $w \in X^{\pm}$ is said to be *(freely) reduced* if it contains no consecutive mutually inverse letters (i.e., it has no subword of the form xx^{-1}, where $x \in X^{\pm}$).

The following are classic results, whose (long but standard) proofs we omit. See, for example [MKS04; LS01].

Proposition 5.1.18. *The word obtained from w by removing pairs of consecutive inverse letters (i.e., applying the rewriting rules $xx^{-1} \to \lambda$, $\forall x \in X^{\pm}$) is unique, called the* free reduction *of w, and denoted by $\widetilde{w}$.* □

In the same vein, the free reduction of a language $\mathcal{L} \subseteq X^{\pm}$ is the set $\widetilde{\mathcal{L}} = \{\widetilde{w} : w \in \mathcal{L}\}$, and we denote by $\mathcal{R}_X$ the full set of reduced words in $X^{\pm}$.

Proposition 5.1.19. *The free group $\mathbb{F}_X$ (with basis X) is isomorphic to $\mathcal{R}_X$ with the operation of concatenation and reduction (i.e., $u \star v = \widetilde{uv}$).* □

That is, we can think elements in the free group either as equivalence classes or as reduced words. We write $\widetilde{S}$ to refer to the set of reduced words in $X^{\pm}$ representing elements in $S \subseteq \mathbb{F}_X$, i.e., $\widetilde{S} = \{ w \in \mathcal{R}_X : [w] \in S \}$. So, we immediately have a natural isomorphism $H \simeq \widetilde{H}$, for any subgroup $H \leqslant \mathbb{F}_X$.

Remark 5.1.20. If $G = \langle X \mid R \rangle$ is a group generated by X, then we have the following natural sequence of epimorphisms:

$$\begin{array}{ccccc} (X^{\pm})^* & \xrightarrow{\mu_{\mathbb{F}}} & \mathbb{F}_X & \xrightarrow{\phi_G} & G = \mathbb{F}_X / \langle\langle R \rangle\rangle \\ w & \longmapsto & [w]_{\mathbb{F}} & \longmapsto & [w]_G \, . \end{array} \qquad \text{(with } \mu_G : (X^{\pm})^* \twoheadrightarrow G\text{)} \tag{5.1}$$

Fig. 5.1: Natural epimorphisms between free monoids, free groups, and groups

When there is no ambiguity on the group G we are working with, we usually omit the subindexes, and denote with the same letter in italics the group element described by a word w; i.e., we write $w = [w]_G$.

5.2 Digraphs and automata

Definition 5.2.1. A *directed multigraph* (*digraph* for short) is a tuple $\Gamma = (V, E, \iota, \tau)$, where:

1. V is a set called set of *vertices* of Γ,
2. E is a set called set of *arcs* or *directed edges*,
3. $\iota, \tau \colon E \to V$ are (resp. initial and final) *incidence functions*.

Then, for each arc $e \in E$, we say e is *incident* to $\iota(e)$ and $\tau(e)$, which are called *initial vertex* (or *origin*), and *final vertex* (or *target*) of e, respectively.

Note that no cardinal or incidence restrictions have been done, in particular we are admitting the *empty digraph* (where $V = E = \varnothing$); and both the possibility of arcs having the same vertex as origin and target (called *directed loops*), and of different arcs sharing the same origin and target (called *parallel arcs*).

Definition 5.2.2. Given a digraph Γ, we usually denote by $V\Gamma$ (resp., by $E\Gamma$) its set of vertices (resp., set of arcs); and call *order* and *size* of Γ the respective cardinals $\# V\Gamma$, and $\# E\Gamma$.

Definition 5.2.3. A digraph Γ is called *finite* (resp. *countable*) if its global number of vertices and arcs ($\# (V\Gamma \sqcup E\Gamma)$) is finite (resp. countable).

Definition 5.2.4. The *in-neighborhood* $E_p^{\leftarrow}$ (resp. *out-neighborhood* $E_p^{\rightarrow}$) of a vertex p in a digraph is the set of arcs with final (resp initial) vertex p. The (total) *neighborhood* of p is $E_p = E_p^{\leftarrow} \cup E_p^{\rightarrow}$. The *in-degree*, *out-degree*, and *total degree* of a vertex p are the cardinals of $E_p^{\leftarrow}$, $E_p^{\rightarrow}$, and E_p respectively.

Definition 5.2.5. An arc e_2 is said to be *reversed* to an arc e_1 if they swap initial and final vertices; i.e., if $\iota(e_i) = \tau(e_j)$, for distinct $i, j \in \{1, 2\}$.

Definition 5.2.6. We say that two vertices are *adjacent* if there exists and arc indident to both of them. For arcs, we have two notions of incidency; namely, two arcs e, f are:

- *strongly incident* if there exists a vertex p such that $\tau(e) = p = \iota(f)$, or $\tau(f) = p = \iota(e)$.
- *(weakly) incident* if there exists a vertex incident to both e and f.

Accordingly, we have the corresponding notions for walks and connectivity.

Definition 5.2.7. A directed walk (resp., undirected walk) in an digraph Γ is either a single vertex p, or a nonempty finite sequence $e_1 \cdots e_k$ of successively strongly (resp., weakly) incident arcs in Γ. In the first case we say that the walk has *length* 0, and goes from p to p; and in the second case that it has *length* k, and goes from $\iota(e_1)$ to $\tau(e_k)$. We will denote the existence of a directed (resp., undirected) walk from p to q by $p \rightsquigarrow q$ (resp. by $p \leftrightsquigarrow q$).

Definition 5.2.8. A digraph is said to be *(strongly) connected* if for every pair of vertices p, q of Γ there exists a (directed) undirected walk from p to q.

Definition 5.2.9. Being connected is an equivalence relation between the vertices of a digraph Γ, whose equivalence classes are called *connected components* of Γ.

Definition 5.2.10. A *labelled digraph* is a digraph together with a map $\ell \colon E \to A$ from the set of arcs to some set A. Then, we say that A is the set of *labels* of Γ, and that Γ is an A- *digraph*. Similarly, an arc e labelled by a is called an a-*arc*, and two parallel arcs with the same label a are called a-parallel.

Remark 5.2.11. The set A is usually thought to be a set of symbols (alphabet) with no further structure. However, we will admit more general labellings along the chapter (see Definition 5.7.3).

Remark 5.2.12. Distinguishing two subsets of (initial and terminal) vertices in a (finite) A-labelled digraph constitutes the standard notion of (finite) automata. We, however, will only use automata having a unique terminal vertex equal to the initial one, which is not a superfluous restriction (see Remark 5.2.21).

Definition 5.2.13. A *(pointed)* A-*automaton* $\Gamma = (V, E, \iota, \tau, \ell, \odot)$ is a connected digraph $\Gamma = (V, E, \iota, \tau)$ with a labelling $\ell \colon E \to A$, and a distinguished vertex $\odot$ (called *base vertex* or *basepoint* of Γ).

In this context, the vertices of Γ are usually called the *states* of the automaton, the basepoint $\odot$ is the *initial* and (unique) *final* state, and the A-labelled arcs define *transitions* between states (an arc $p \xrightarrow{a} q$ defines a transition from state p to q with label a). An automaton is called *finite* if it has a finite number of states, and transitions (i.e., if $\# V < \infty$ and $\# E < \infty$).

Given a subset $S \subseteq A$, an A-automaton is said to be *S-complete* (resp., *S-regular*) if for every vertex p, and for every letter $a \in S$, there exists at least (resp., exactly) one arc departing p with label a. An A-complete A-automaton is simply called *complete*.

Remark 5.2.14. Hereafter, if not stated otherwise, we will assume all automata to be pointed, and we will simply refer them as automata.

Also, if not stated otherwise, we will assume that all the walks appearing in this chapter are directed, and thus its labelling determines (spells) a word in A^*.

Definition 5.2.15. The *label of a walk* γ in an A-automaton Γ is defined to be

$$\ell(\gamma) := \begin{cases} \lambda, & \text{if } \gamma = p, & \text{(an empty walk in } \Gamma), \\ \ell(e_1) \cdots \ell(e_k), & \text{if } \gamma = e_1, \ldots, e_k & \text{(a nonempty walk in } \Gamma). \end{cases}$$

Then, we say that γ *reads* or *spells* the word $w = \ell(\gamma) \in A^*$. We denote the existence of a walk from p to q reading the word $w \in A^*$ by $p \overset{w}{\rightsquigarrow} q$.

Also, if G is a group generated by X, we will denote by $\ell_G(\gamma)$ the element in G described by the word $w \in (X^{\pm})^*$ read by γ; formally (see diagram (5.1)):

$$\ell_G(\gamma) := (\ell(\gamma))\mu_G \,.$$

Definition 5.2.16. We say that a vertex p in Γ is *accessible* (resp. *coaccessible*) if there exists a directed walk in Γ from the basepoint to p (resp. from p to the basepoint); i.e.,

$$p \text{ is accessible} \iff \odot \rightsquigarrow p\,,$$
$$p \text{ is coaccessible} \iff p \rightsquigarrow \odot\,.$$

Walks starting and ending in the basepoint are called $\odot$-closed walks ($\odot$-walks, for short).

Remark 5.2.17. A vertex is both accessible and coaccessible if and only if it belongs to some $\odot$-walk.

Definition 5.2.18. Let Γ be a A-automaton. A word $w \in A^*$ is said to be *accepted* (or *recognized*) by Γ if there exists a $\odot$-walk in Γ reading w; i.e.,

$$w \in A^* \text{ is accepted by } \Gamma \iff \odot \overset{w}{\rightsquigarrow} \odot \text{ in } \Gamma\,.$$

Definition 5.2.19. The *language recognized* (or *accepted*) by an A-automaton Γ is the set $\mathcal{L}(\Gamma)$ consisting of all the words in A accepted by Γ; i.e.,

$$\begin{aligned}\mathcal{L}(\Gamma) &= \left\{ w \in A^* : \circledcirc \overset{w}{\rightsquigarrow} \circledcirc \text{ in } \Gamma \right\} \\ &= \left\{ \ell(\gamma) : \gamma \text{ is a } \circledcirc\text{-walk in } \Gamma \right\} \subseteq A^* .\end{aligned}$$

Remark 5.2.20. Note that allowing parallel arcs with the same label (i.e., redundant transitions) do not contribute to the language recognized by the automaton. We admit them because they appear in a natural way from the techniques that follow (e.g. closed foldings (5.7)).

Remark 5.2.21. Our definition of (pointed) automata in Definition 5.2.13 supposes an actual restriction in terms of the possible recognized languages, since, any language recognized by our automata must contain, for example, the empty string (which is not the case for a general automaton).

Remark 5.2.22. For any A-automaton Γ, the restriction of the labelling to $\circledcirc$-walk s is an epimorphism (of monoids):

$$\begin{array}{rcl}\ell_{\circledcirc} \colon \{\circledcirc\text{-walks in } \Gamma\} & \twoheadrightarrow & \mathcal{L}(\Gamma) \\ \gamma & \mapsto & \ell(\gamma)\end{array} \tag{5.2}$$

which is not injective in general (see Corollary 5.4.5). In particular $\mathcal{L}(\Gamma) = \operatorname{im}(\ell_{\circledcirc})$ is closed by concatenation, and hence not every language $\mathcal{L} \subseteq A^*$ can be recognized by a A-automaton. That is, the map

$$\begin{array}{rcl}\{\text{A-automata}\} & \to & \{\text{subsets of } A^*\} \\ \Gamma & \mapsto & \mathcal{L}(\Gamma) .\end{array}$$

is not onto.

Below, we formalize transformations preserving the language recognized by an automaton.

Definition 5.2.23. Let Γ, Γ' be A-automata. A *homomorphism of automata* between Γ and Γ' is a map $\phi : V(\Gamma) \to V(\Gamma')$ between their respective vertex sets which preserves basepoints, and labelled adjacency. More formally, ϕ is such that:

1. If p_0 is the basepoint of Γ, then $p_0\phi$ is the basepoint of Γ'.
2. If $p \xrightarrow{a} q$ is an arc of Γ, then $p\phi \xrightarrow{a} q\phi$ is an arc of Γ'.

Remark 5.2.24. Note that the second condition allows us to extend the homomorphism ϕ to arcs in a natural (but not necessarily unique) way: for every arc $p \xrightarrow{a} q$ of Γ, define its image by ϕ to be (one of the arcs) of the form $p\phi \xrightarrow{a} q\phi$.

This ambiguity disappears if we remove redundancies (parallel arcs with the same label) from the automata; then the previous agreement provide a unique map $\phi\colon E(\Gamma) \to E(\Gamma')$ between arcs. This fact justifies the usual abuse of notation consisting in denoting homomorphisms of automata by $\phi\colon \Gamma \to \Gamma'$.

The distinctive property of homomorphism of automata summarized in the following result is easily proved by induction.

Lemma 5.2.25. *If $\phi\colon \Gamma \to \Gamma$ is a homomorphism of automata, then for any vertices p, q in Γ, and every word $w \in A^*$:*

$$p \xrightarrow{w} q \implies p\phi \xrightarrow{w} q\phi\,.$$

Remark 5.2.26. So, if there exists an homomorphism of automata between Γ and Γ', then any word accepted by Γ is also accepted by Γ', i.e.,

$$\exists \text{ homomorphism } \phi\colon \Gamma \to \Gamma' \implies \mathcal{L}(\Gamma) \subseteq \mathcal{L}(\Gamma')\,. \tag{5.3}$$

The converse implication, however, is easily shown to be false. In order it to be true, we need an important extra condition for the image of ϕ (see Proposition 5.4.12).

5.3 Involutive automata

In order to deal with subgroups of groups given by presentations $\langle X \mid R\rangle$, it is natural to consider automata with an involutive alphabet (corresponding to the involutive closure $X^{\pm}$ of the generating set $X = \{x_j\}_j$ in the presentation), and consistent with involution. We make this requirements precise in the following definition.

Definition 5.3.1. An *involutive* X*-automaton* is an automaton having an (involutive) alphabet $X^{\pm}$, and an involution ${}^{-1}\colon E \to E$, $e \mapsto e^{-1}$ on its arcs such that:

1. No arc e can be the inverse of itself (i.e., $e^{-1} \neq e$, for every $e \in E$).
2. Inverse arcs are reversed (i.e., $\iota e^{-1} = \tau e$, for every $e \in E$).
3. Arc involution is compatible with label inversion (i.e., $\ell(e^{-1}) = \ell(e)^{-1}$, for every $e \in E$).

Definition 5.3.2. In an involutive automaton, we define the *reduced label*, denoted by $\widetilde{\ell}$, of a walk γ in the natural way, i.e., $\widetilde{\ell}(\gamma) \coloneqq \widetilde{\ell(\gamma)}$.

Thus, in an involutive automaton, for every labelled arc $e \equiv p \xrightarrow{a} q$, there exists a "unique" reversed arc $e^{-1} \equiv p \xrightarrow{a^{-1}} q$ (called the inverse of e). That is, arcs appear by pairs:

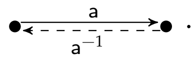

(5.4)

Fig. 5.2: Arcs in an involutive automaton

We say that an arc is *positive* (resp. *negative*) if it is labelled with a letter in X (resp. X^{-1}). We denote by E^+ and E^- the sets of positive and negative arcs, respectively.

Remark 5.3.3. An involutive automaton is strongly connected by construction; in particular, every vertex is both accessible and coaccessible.

Definition 5.3.4. The *(undirected) underlying graph* of an involutive automaton Γ is the undirected graph, denoted by $\underline{\Gamma}$, obtained by omitting the labelling, and identifying all pairs of respectively inverse arcs in Γ.

Remark 5.3.5. Every undirected graph can be defined in this way. Note that then an edge in $\underline{\Gamma}$ is an unordered pair $\{e, e^{-1}\}$, where $e \in E^+\Gamma$, and $e^- \in E^-\Gamma$. Hence, the total degree of any vertex in an involutive automaton Γ doubles the degree of its underlying graph (which in turn coincides with both the in-degree and the out-degree of Γ). Thus,

$$\# E\underline{\Gamma} = \# E^+\Gamma = \# E^-\Gamma = \# E\Gamma/2\,. \tag{5.5}$$

Remark 5.3.6. Note that for an involutive X-automata it is enough to represent positively-labelled arcs (i.e., X-labelled arcs) in order to fully describe it; with the tacit assumption that for every shown positive arc there is also a hidden — dashed in (5.4) — inverse arc.

Convention 5.3.7. It will be convenient to think involutive X-automata just as connected X^+-automata whose arcs can be crossed in both directions (with the assumption that an x-arc crossed backwards is read x^{-1}).

This convention, together with the relations in (5.5) makes it natural to abuse language and refer to positive arcs as edges, to the cardinal of $E^+\Gamma$ as the size (number of edges) of an involutive automaton Γ, and of the cardinal $E^+(p)$ as the degree of any vertex p in Γ.

In the same vein — interpreting undirected graphs as (inverse-arc) identified digraphs — if Δ is subgraph of the underlying graph of Γ, we denote by $\boldsymbol{\Delta}$ the involutive subautomaton of Γ having underlying graph Δ. In this way, we can easily refer to paths, trees, etc. (as involutive subautomata) within an involutive automaton.

Definition 5.3.8. The *(circuit) rank* of an undirected graph Γ, denoted $\operatorname{rk}\Gamma$ is the minimum number of edges that must be removed from Γ to make it acyclic. For finite nonempty connected graphs, Euler's formula states that

$$\operatorname{rk}\Gamma = \#\mathsf{E}\Gamma - \#\mathsf{V}\Gamma + 1\,. \tag{5.6}$$

(That is, the rank of a connected graph is the cardinal of the set of edges outside any spanning tree.) The *rank of an involutive automaton* $\mathbf{\Gamma}$ is the rank of its underlying graph; i.e.,

$$\operatorname{rk}\mathbf{\Gamma} = \operatorname{rk}\underline{\mathbf{\Gamma}} = \#(\mathsf{E}\mathbf{\Gamma} \setminus \mathsf{E}\mathbf{T})\,,$$

where $\mathbf{T}$ is any spanning tree of $\mathbf{\Gamma}$.

Definition 5.3.9. A labelled arc e_2 is said to be *reversed* to a labelled arc e_1 in an involutive automaton $\mathbf{\Gamma}$ if they are reversed, and the labels are inverse of each other.

The existence and computability of spanning trees for *arbitrary* graphs is a subtle matter. We summarize without proof the results on this topic which are relevant to us (the result for finite graphs is folklore; for details on spanning trees for infinite graphs see [Sti93, section 2.1.5], and [Sou08]).

Theorem 5.3.10. *Let Γ be a connected undirected graph. Depending on the order of Γ we distinguish the following cases:*

- *If Γ is finite, then a spanning tree always exists and can be computed in linear time.*
- *If Γ is infinite, then the existence of spanning trees for Γ is equivalent to the axiom of choice* (AC)*.*
- *If Γ is recursively enumerable, then (assuming* AC*) there exists a recursively enumerable spanning tree for Γ.*

So, throughout the chapter, AC will be assumed wherever the existence of general spanning trees is needed.

Definition 5.3.11. A walk γ in an involutive automaton is said to be *reduced* if no successive edges in γ are mutually reversed (i.e., a walk is reduced if it has no vertex backtracking). Otherwise we say that the walk γ is *degenerated*.

Note that we can remove any vertex backtracking from any walk γ to obtain a reduced walk, denoted $\widetilde{\gamma}$, reading the same word once reduced, i.e., such that:

$$\widetilde{\ell}(\gamma) = \widetilde{\ell}(\widetilde{\gamma})\,.$$

We say that $\widetilde{\gamma}$ is the *walk-reduction* of γ.

Remark 5.3.12. Since we are admitting parallel arcs with the same label, the inverse arc e^{-1} can be not the only arc reversed to e, and thus walk-reduction is not necessarily unique in a general involutive automaton.

Remark 5.3.13. Note that if γ is a walk in an involutive automaton, then:

$$\ell(\gamma) \text{ is reduced } \implies \gamma \text{ is reduced}\,,$$

but the converse is not necessarily true, since subwalks of the form

$$\bullet \overset{a}{\dashrightarrow} \bullet \overset{a^{-1}}{\dashrightarrow} \bullet \tag{5.7}$$

(reading aa^{-1}) can appear in a reduced walk. Thus, in an involutive automaton, we always have $|\ell(\widetilde{\gamma})| \leqslant |\widetilde{\ell}(\gamma)|$, and, in general $\widetilde{\ell}(\gamma) \neq \ell(\widetilde{\gamma})$ (see Lemma 5.4.9).

Involutive automata can be used to represent subgroups. Namely, if X is a set of generators for a group G as in (5.1), and Γ is an involutive X-automaton, then the labels of the set of walks starting and ending at the basepoint (i.e., the $\odot$-closed walks) of Γ describe a subgroup of G. More precisely, the following result is immediate from the algebra of walks in the automaton.

Proposition 5.3.14. *Let Γ be an involutive X-automaton, then the language recognized by Γ contains the empty word, and is closed by concatenation and taking inverses (i.e., is an involutive submonoid of $(X^{\pm})^*$).* □

The following immediate lemma about the behaviour of walks within involutive trees is essential to translate the previous proposition (about monoids) into Proposition 5.3.16 (about groups).

Lemma 5.3.15. *Let T be an involutive X-automaton, whose underlying graph is a tree. Then,*

(i) *for any two vertices p, q in T, there exist a unique reduced walk from p to q along T. We denote it by $\gamma_{\mathsf{T}}[p,q]$.*

$$\gamma_{\mathsf{T}}[p,q] \;\equiv\; p \overset{\mathsf{T}}{\rightsquigarrow} q$$

Fig. 5.3: The (unique) reduced walk from p to q along the tree T

(ii) *any $\odot$-walk γ within T can be obtained from the empty $\odot$-walk by recursively inserting a finite number of subwalks of the form $\eta_i\eta_i^{-1}$ (where the η_i's are walks in T as well). Therefore, the group element represented by γ is always trivial (i.e., $\ell_G(\gamma) = 1$, for every $G = \langle X \rangle$).* □

Proposition 5.3.16. *Let Γ be an involutive X-automaton, let $G = \langle X \mid R \rangle$ be a group, and let $\mu_G \colon (X^{\pm})^* \to G$ be the natural projection (5.1). Then, the set of words recognized by Γ describe (through μ_G) a subgroup of G denoted by $\langle \Gamma \rangle_G$; namely:*

$$\langle \Gamma \rangle_G \coloneqq (\mathcal{L}_\Gamma)\mu_G = \{\, \ell_G(\gamma) : \gamma \text{ is a } \odot\text{-walk in } \Gamma \,\} \tag{5.8}$$

is a subgroup of G with generating set

$$\{\, \ell_G(\odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\ e\ } \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot) : e \in E^+(\Gamma \setminus \mathbf{T}) \,\}, \tag{5.9}$$

for any spanning tree $\mathbf{T}$ of Γ (where $\xrightarrow{\ e\ }$ denotes the arc named e, and $\overset{\mathbf{T}}{\rightsquigarrow}$ denotes the unique reduced walk through $\mathbf{T}$ between vertices).

In particular, if Γ is finite, then $\langle \Gamma \rangle_G$ is finitely generated.

Proof. Since every walk $\odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\ e\ } \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot$ is a $\odot$-walk in Γ, it is clear that the subgroup generated by (5.9) is included in $\langle \Gamma \rangle_G$.

To see the opposite inclusion, we need to prove that every element in $\langle \Gamma \rangle_G$ can be written as a product of elements of (5.9) and its inverses.

Consider a nontrivial element $h \in \langle \Gamma \rangle_G$. Then, since the automaton is involutive , there exists a $\odot$-walk

$$\gamma \equiv \quad \odot \xrightarrow{e_1^{\epsilon_1}} p_1 \xrightarrow{e_2^{\epsilon_2}} p_2 \cdots p_{k-1} \xrightarrow{e_k^{\epsilon_k}} \odot$$

such that $e_i \in E^+$ and $\epsilon_i = \pm 1$ for all $i = 1, \ldots, k$, and $\ell(\gamma) =_G h$.

Now, let γ' be the walk obtained from γ by going to the basepoint through $\mathbf{T}$ and back, after visiting every intermediate vertex in γ:

$$\odot \xrightarrow{e_1^{\epsilon_1}} p_1 \overset{\mathbf{T}}{\rightsquigarrow} \odot \overset{\mathbf{T}}{\rightsquigarrow} p_1 \xrightarrow{e_2^{\epsilon_2}} p_2 \overset{\mathbf{T}}{\rightsquigarrow} \odot \overset{\mathbf{T}}{\rightsquigarrow} p_2 \overset{\mathbf{T}}{\rightsquigarrow} \odot \cdots$$
$$\cdots \odot \overset{\mathbf{T}}{\rightsquigarrow} p_{k-2} \xrightarrow{e_{k-1}^{\epsilon_{k-1}}} p_{k-1} \overset{\mathbf{T}}{\rightsquigarrow} \odot \overset{\mathbf{T}}{\rightsquigarrow} p_{k-1} \xrightarrow{e_k^{\epsilon_k}} \odot .$$

Then, defining $p_0 = p_k = \odot$, we have

$$\ell(P') = \left(\prod\nolimits_{i=1}^{k-1} \ell(e_i)^{\epsilon_i} \cdot \ell(\bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet) \right) \cdot \ell(e_k)^{\epsilon_k} . \tag{5.10}$$

Now (from Lemma 5.3.15),

1. since for every vertex p, $\ell_G(p \overset{\mathbf{T}}{\rightsquigarrow} \odot \overset{\mathbf{T}}{\rightsquigarrow} p) = 1$, we have

$$\ell_G(P') = \ell_G(e_1)^{\epsilon_1} \cdot \ell_G(e_2)^{\epsilon_1} \cdots \ell_G(e_k)^{\epsilon_1} = \ell_G(P) = h, \tag{5.11}$$

and so γ' also describes the element h in G.

2. since for every arc e_i in $\mathbf{T}$,

$$\ell_G\left(\circledcirc \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{e_i} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \circledcirc\right) = 1\,,$$

we can remove all the arcs in $\mathbf{T}$ from the product in (5.10) (and so, use only arcs in $\Gamma \setminus \mathbf{T}$).

Therefore, every element $h \in \langle \Gamma \rangle_G$ can be generated by the family (5.9), as we wanted to prove. This concludes the proof. □

Definition 5.3.17. The subgroup in (5.8) is called *subgroup recognized* by Γ in G, and denoted by $\langle \Gamma \rangle_G$ (or simply $\langle \Gamma \rangle$ if the background group G is clear).

So, every involutive X-automaton describes some subgroup of $G = \langle X \rangle$. On the other hand, given a word $w = w(X)$ in the generators, we can always consider the involutive $\circledcirc$-walk $\mathsf{Fl}(w)$ spelling w (or w^{-1} if read in the opposite direction). This is the petal associated to w.

Recall that, for involutive automata, we only represent arcs with positive labels, and assume that for every arc labelled $x \in X$, there is always a unique reversed arc labelled x^{-1}.

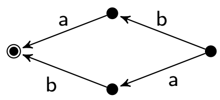

Fig. 5.4: The petal automaton associated to the word $a^{-1}b^{-1}ab$

Now it is straightforward to build an involutive X-automaton recognizing any given subgroup of $G = \langle X \rangle$.

Definition 5.3.18. Let $S = \{w_i(X)\}_i$ be a set of words in the generators of a group G. The *flower automaton* of S, denoted by $\mathsf{Fl}(S)$, is the involutive X-automaton obtained from the petals reading the words $w_i \in S$ by identifying all their basepoints.

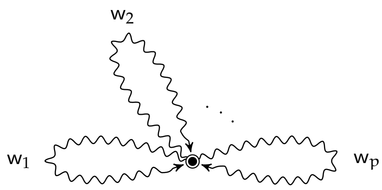

Fig. 5.5: The flower automaton of $\{w_1, w_2, \ldots, w_p\} \subseteq (X^{\pm})^*$

It is clear that the subgroup of G recognized by the flower automaton $\mathsf{Fl}(S)$ is precisely the subgroup generated by S. Thus, any subgroup $H \leqslant G$ is represented by some involutive X-automaton and vice versa. Moreover, the subgroup is finitely (resp., countably) generated if and only if it can be represented by some finite (resp., countable) involutive X-automaton.

So, given a group $G = \langle X \mid R \rangle$, the map

$$\begin{array}{rcl} \{\text{involutive X-automata}\} & \to & \{\text{subgroups of } G\} \\ \Gamma & \mapsto & \langle \Gamma \rangle_G\,, \end{array} \tag{5.12}$$

sending every involutive automaton to the subgroup it recognizes in G, is well-defined and onto. Of course it is not injective since you can, for example, consider flower automata associated to different generating sets for the same subgroup; or, more importantly, the petal automata associated to any relation in the group describes the same (trivial) subgroup of G.

Remark 5.3.19. Of course, if G is Noetherian (resp., countable), then it is enough to consider finite (resp., countable) automata in Equation (5.12). This facts have important implications related to computability of subgroups through automata.

Definition 5.3.20. Let G be a group generated by X, and let Γ, Γ' be two involutive X-automata. Then, we say that Γ, Γ' are *equivalent automata* with respect to G, denoted by $\Gamma \equiv_G \Gamma'$ if they recognize the same subgroup of G; i.e.,

$$\Gamma \equiv_G \Gamma' \iff \langle \Gamma \rangle_G = \langle \Gamma' \rangle_G$$

As usual, if there is no ambiguity about the group G, we will omit the references to it, and simply talk of equivalent automata, denoted by $\Gamma \equiv \Gamma'$.

5.4 Stallings automata

A natural goal in this scenario is trying to normalize the map (5.12) (i.e., to distinguish, for every subgroup $H \leqslant G$, a canonical representative among the automata recognizing H. This would allow us to restrict (5.12) to a bijection, which we would like to be computable in the finite case.

This is essentially what Stallings did (from a more topological wiewpoint) for free groups in his celebrated paper [Sta83]. Although, it is often claimed — and not very surprising — that the arguments can be translated to combinatorial language in full generality, we know no detailed survey of such treatment. Below, we recreate this nice theory using algorithmic language (see [KM02; BS10] for similar approaches mostly focused in the finitely generated case).

Let $\mathbb{F}_X = \langle X \mid - \rangle$ denote the free group with basis X.

Remark 5.4.1. Recall that, according to Proposition 5.1.19, we can identify elements in the free group (thought as classes of words) with freely reduced words; i.e., for each word $w \in (X^{\pm})^*$, we can assume $w\mu = \widetilde{w}$. Similarly, for every X-walk γ, we have that $\ell_{\mathbb{F}_X}(\gamma) = \widetilde{\ell}(\gamma)$.

From the previous discussion, we have the following natural sequence of onto maps:

$$\begin{array}{ccccc} \{\text{X-}\mathbf{Involutive}\} & \xrightarrow{\mathcal{L}} & \{\text{involutive submonoids of } (X^{\pm})^*\} & \xrightarrow{\mu} & \{\text{subgroups of } \mathbb{F}_X\} \\ \Gamma & \longmapsto & \mathcal{L}_\Gamma & \longmapsto & \langle \Gamma \rangle , \end{array} \tag{5.13}$$

where $\mathcal{L}$ and μ, and so its composition $\Gamma \mapsto \langle \Gamma \rangle$, are not injective in general. We will see that we can overcome this last ambiguity by imposing two conditions in the automata: namely *determinism*, and *trimness*.

Definition 5.4.2. An A-automaton is called *deterministic* if it has no different arcs with the same label exiting the same vertex (i.e., for every vertex there is at most one arc exiting it with the same label).

Namely, an automaton is deterministic if and only if *none* of the following situations occur:

Fig. 5.6: Nondeterministic situations

Remark 5.4.3. In a deterministic A-automaton Γ:

1. There are no parallel arcs in Γ with the same label. Note that for such (non-redundant) automata, one can define the set of transitions as a subset of $V(\Gamma) \times A \times V(\Gamma)$, recovering a standard (a bit more tight) notion of automata.

2. For each $p \in V(\Gamma)$, and each $a \in A$ there exists at most one vertex q such that $p \xrightarrow{a} q$ (in particular, the out-degree of any vertex is at most $\# A$). Thus, in these automata we can define transitions through a (partial) function

$$\begin{array}{ccc} V(\Gamma) \times A & \to & V(\Gamma) \\ (p, a) & \mapsto & q \end{array} \tag{5.14}$$

which is called *transition function* of the deterministic automaton, and can be extended to a partial function $V \times A^* \to V$ in the natural way.

3. The degree of any vertex in Γ is at most #A. A deterministic A-automata is complete if and only if it is A-regular.

An important consequence of the definition is easily proved by induction.

Lemma 5.4.4. *In a deterministic automaton Γ, if two walks starting at the same vertex read the same word, then they are identical (and, in particular the final vertex is the same).* □

So, for any fixed vertex p in a deterministic automaton Γ, the following map is injective:

$$\begin{array}{rcl} \{\text{walks in } \Gamma \text{ starting at } p\} & \to & A^* \\ \gamma & \mapsto & \ell(\gamma) \end{array}$$

Corollary 5.4.5. *Let Γ be a deterministic A-automaton. Then, the map:*

$$\begin{array}{rcl} \ell_{\odot} \colon \{\odot\text{-walks in } \Gamma\} & \to & \mathcal{L}(\Gamma) \\ \gamma & \mapsto & \ell(\gamma) \end{array} \tag{5.15}$$

is an isomorphism of monoids. □

For an involutive automaton, being deterministic implies two crucial properties: one regarding the relation between automata and the languages they recognize (Corollary 5.4.14), and the other regarding the behaviour of reduced words in the automata.

Example 5.4.6 (Schreier coset digraph). If H is a subgroup of a group X with generating set X, the *(right) Schreier coset digraph* of H relative to X, denoted by $\mathrm{Sch}(H, X)$, is the directed graph with vertices the right cosets $Hg \in H\backslash G$, and for every vertex Hg, and every element $x \in X^{\pm}$, an arc labelled x from Hg to Hgx.

Note that setting the coset H as basepoint confers to $\mathrm{Sch}(H, X)$ structure of pointed automata. Then the following result is immediate.

Lemma 5.4.7. *If X is a subset of a group G, then the Schreier coset digraph $\mathrm{Sch}(H, X)$ is a (complete) deterministic involutive X-automaton recognizing H.* □

The following results state key facts about deterministic involutive automata.

Remark 5.4.8. In a deterministic involutive automata, the inverse arc e^{-1} is the unique arc reversed to a given arc e; i.e.,

$$e' \text{ is reversed to } e \iff e' = e^{-1} (\text{the inverse of } e)\,.$$

Lemma 5.4.9. *Let γ be a walk in a deterministic involutive X-automaton. Then, the following statements are equivalent.*

(a) $\ell(\gamma) = aa^{-1}$*, for some* $a \in X^{\pm}$.

(b) $\gamma = ee^{-1}$*, for some edge* e *in* Γ *(labelled* a*).*

So, in a deterministic involutive automaton,

$$\gamma \text{ is reduced} \iff \ell(\gamma) \text{ is reduced},$$

and $\ell(\widetilde{\gamma}) = \widetilde{\ell}(\gamma)$. ☐

Lemma 5.4.10. *Let* Γ *be a deterministic involutive* X*-automaton. Then, for any vertices* p, q *in* Γ*, and any words* w, u, v *in* $X^{\pm}$*,*

$$p \xrightarrow{uww^{-1}v} q \implies p \xrightarrow{uv} q\,.$$

So, if Γ *is involutive and deterministic, then* $\widetilde{\mathcal{L}}_{\Gamma} \subseteq \mathcal{L}_{\Gamma}$.

Proof. Let $\gamma \equiv p \xrightarrow{uww^{-1}v} q$, and suppose that $w = w'a$, with $a \in X^{\pm}$. From Lemma 5.4.9, γ must have two successively inverse edges (reading aa^{-1}) that we can remove to obtain a new (shorter) walk

$$\gamma' \equiv p \xrightarrow{uw'(w')^{-1}v} q\,.$$

The result follows from repeating this procedure until $w' = \lambda$. ☐

Corollary 5.4.11. *If* Γ *is a deterministic involutive* X*-automaton, the isomorphism* $\ell_{\circledast}$ *in Corollary 5.4.5 restricts to the following isomorphism of (free) groups:*

$$\begin{aligned} \ell_{\circledast} \colon \{\text{reduced } \circledast\text{-walks in } \Gamma\} &\to \widetilde{\mathcal{L}}(\Gamma) \simeq \langle \Gamma \rangle \\ \gamma &\mapsto \ell(\gamma) = \widetilde{\ell}(\gamma)\,. \end{aligned} \tag{5.16}$$

The existence of a homomorphism $\Gamma \to \Gamma'$, trivially implies that $\mathcal{L}(\Gamma) \subseteq \mathcal{L}(\Gamma')$. Now we prove that, for involutive automata, the converse is also true if the codomain Γ' is deterministic.

Proposition 5.4.12. *Let* Γ, Γ' *be involutive* X*-automata, with* Γ' *deterministic. Then:*

$$\mathcal{L}(\Gamma) \subseteq \mathcal{L}(\Gamma') \iff \exists \text{ homomorphism } \phi \colon \Gamma \to \Gamma'\,.$$

and, if so, the homomorphism is unique.

Proof. Unicity is consequence of the determinism of Γ': indeed, suppose that we have homomorphisms $\phi, \phi' \colon \Gamma \to \Gamma'$, and a vertex p in Γ such that $p\phi = p' \neq p'' = p\phi'$. Then the images $\gamma\phi, \gamma\phi'$ of any walk $\gamma \equiv \circledast \rightsquigarrow p$ in Γ, would be walks in Γ' with the same starting points and labels, but different end point, in contradiction with Γ' being deterministic (Lemma 5.4.4).

[$\Leftarrow$] This implication is immediate from the definition of homomorphism of automata in 5.2.23.

[$\Rightarrow$] Suppose $\mathcal{L}(\Gamma) \subseteq \mathcal{L}(\Gamma')$ with Γ' deterministic. Then, given a vertex p in Γ, consider a $\odot$-walk in Γ through p

$$\odot \overset{\mathsf{u}}{\rightsquigarrow} \mathsf{p} \overset{\mathsf{v}}{\rightsquigarrow} \odot \,.$$

accepting some word uv. Since $\mathcal{L}(\Gamma) \subseteq \mathcal{L}(\Gamma')$, we can read the same word in Γ', i.e.,

$$\odot' \overset{\mathsf{u}}{\rightsquigarrow} \mathsf{p}' \overset{\mathsf{v}}{\rightsquigarrow} \odot' \,,$$

where p' is some vertex in Γ'. We claim that such a vertex p' is uniquely determined, and the corresponding map

$$\begin{aligned} \phi \colon \mathsf{V}(\Gamma) &\to \mathsf{V}(\Gamma') \\ \mathsf{p} &\mapsto \mathsf{p}' \end{aligned} \tag{5.17}$$

(sending basepoint to basepoint) obtained in this way defines a homomorphism between Γ and Γ'.

For the first claim, suppose that $\odot \overset{\mathsf{u}'}{\rightsquigarrow} \mathsf{p} \overset{\mathsf{v}'}{\rightsquigarrow} \odot$ is another $\odot$-walk through p in Γ. Then,

$$\odot \overset{\mathsf{u}}{\rightsquigarrow} \mathsf{p} \overset{\mathsf{v}'}{\rightsquigarrow} \odot$$

is also a $\odot$-walk through p in Γ. Therefore, both uv, and uv' must be read by $\odot'$-walks in Γ':

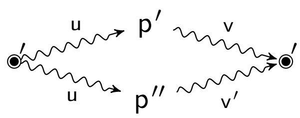

Then, since Γ' is deterministic, from Lemma 5.4.4 we have $\mathsf{p}' = \mathsf{p}''$, and the map ϕ in (5.18) is well defined.

Finally, to see that ϕ is a homomorphism of automata, consider a labelled arc $\mathsf{e} \equiv \mathsf{p} \xrightarrow{\mathsf{a}} \mathsf{q}$ in Γ. Then (since p, q are both accessible and coaccessible), there exists a $\odot$-walk through e in Γ:

$$\odot \overset{\mathsf{u}}{\rightsquigarrow} \mathsf{p} \xrightarrow{\mathsf{a}} \mathsf{q} \overset{\mathsf{v}}{\rightsquigarrow} \odot \,.$$

Thus, the word uav belongs to $\mathcal{L}_{\Gamma}$, and so to $\mathcal{L}_{\Gamma'}$; which according to definition (5.18) means that in Γ' we have the walk

$$\odot \overset{\mathsf{u}}{\rightsquigarrow} \mathsf{p}\phi \xrightarrow{\mathsf{a}} \mathsf{q}\phi \overset{\mathsf{v}}{\rightsquigarrow} \odot \,.$$

In particular, we have an arc $p\phi \xrightarrow{a} q\phi$ in Γ'. Thus $\phi\colon \Gamma \to \Gamma'$ is a homomorphism of automata, as we wanted to prove. $\square$

The following two corollaries are immediate, and important for us.

Corollary 5.4.13. *If Γ is a deterministic involutive automaton, then the only homomorphism $\Gamma \to \Gamma$ is the identity.* $\square$

Corollary 5.4.14. *If two deterministic involutive* X*-automata Γ, Γ' recognize the same language, then they are isomorphic; i.e., for deterministic involutive* X*-automata Γ, Γ',*

$$\mathcal{L}_{\Gamma} = \mathcal{L}_{\Gamma'} \;\Leftrightarrow\; \Gamma \simeq \Gamma'. \qquad \square$$

We have already seen (Corollary 5.4.11) that, in a deterministic involutive automaton Γ, reduced $\odot$-walks biunivocally describe the recognized subgroup $\langle \Gamma \rangle$. So, it is clear that the vertices in Γ not laying in any reduced $\odot$-walk play no role in this description. Note, however, that the accepted language $\mathcal{L}_{\Gamma}$ can change if one removes those vertices.

Definition 5.4.15. A vertex (resp. an arc) in an involutive automaton is called *useful* if it belongs to some reduced $\odot$-walk, and *superfluous* otherwise.

Definition 5.4.16. An involutive automaton Γ is called *trim* (or *core*) if it has no superfluous vertices. (i.e., every vertex in Γ is useful).

Any transformation consisting in removing superfluous vertices from an involutive automaton is also called a *trim*. A *total trim*, denoted τ, is the transformation consisting in removing from an inverse automata all its superfluous vertices. The *core* of an involutive automaton Γ, denoted by $\mathrm{core}(\Gamma)$, is the automaton obtained after applying a total trim to Γ, i.e., $\mathrm{core}(\Gamma) \coloneqq (\Gamma)\tau$.

Remark 5.4.17. The core of an involutive automaton Γ is precisely the (trim) automaton obtained by removing from Γ all the (finite or infinite) hanging trees not containing the basepoint. The finite case, however, admits a local description: a finite involutive automaton is trim if and only if it has no non-base vertices of degree one.

Remark 5.4.18. The subgroup recognized by an automaton Γ coincides with the subgroup recognized by its core; i.e., $\langle \Gamma \rangle = \langle \mathrm{core}(\Gamma) \rangle$.

Lemma 5.4.19. *Let Γ be a involutive automaton. Then, Γ is trim (i.e., every vertex in Γ is useful) if and only if every arc in Γ is useful.*

Proof. One of the implications is obvious. For the other one suppose, that every vertex in Γ is useful, and consider an arc $e \equiv p \xrightarrow{a} q$ in Γ. From the usefulness of

p and q we have reduced $\circledast$-closed walks throug p and q in Γ. If e belongs to any of those walks, we are done. Otherwise, we have the following situation:

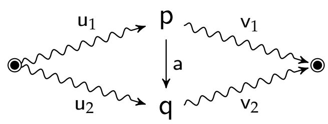

Now it is clear that the walk $\circledast \xrightarrow{u_1} p \xrightarrow{a} q \xrightarrow{v_2} \circledast$ is reduced. The claimed result follows. □

Definition 5.4.20. An *inverse* X*-automaton* is an involutive X-automaton which is deterministic and trim. We denote by X-**Inverse** the full family of X-inverse automata.

Remark 5.4.21. In an inverse automata the basepoint still can have degree one, when this happens we call *basepoint thread* the path from the basepoint to the first vertex with degree strictly grater than 2.

Remark 5.4.22. The degree (resp., total degree) of a vertex in an inverse X-automaton is at most $\#X$ (resp., $2 \cdot \#X$).

Example 5.4.23. If S is a subset of a group $G = \langle X \rangle$, then the core of the Schreier coset digraph of the subgroup $H = \langle S \rangle$ relative to X is a (not necessarily complete) inverse X-automaton recognizing H. We will see in Theorem 5.4.32 that, for the free group, this is (modulo isomorphism) the unique inverse automata recognizing a subgroup $H \leqslant \mathbb{F}_X$.

For an involutive automaton, determinism (through Lemma 5.4.9) and trimness easily produce the following variant of Proposition 5.4.12.

Proposition 5.4.24. *Let* Γ, Γ' *be inverse* X*-automata. Then:*

$$\widetilde{\mathcal{L}}(\Gamma) \subseteq \widetilde{\mathcal{L}}(\Gamma') \;\Leftrightarrow\; \exists \text{ homomorphism } \phi\colon \Gamma \to \Gamma' .$$

and, if so, the homomorphism is unique.

Proof. Unicity is consequence of the determinism of Γ' (see the proof of Proposition 5.4.12).

[$\Leftarrow$] This implication is immediate from the definitions of homomorphism of automata in 5.2.23, and reduction of a language.

[$\Rightarrow$] Suppose $\widetilde{\mathcal{L}}(\Gamma) \subseteq \widetilde{\mathcal{L}}(\Gamma')$. Then, given a vertex p in Γ (which is trim) we can consider a reduced $\circledast$-walk in Γ through p

$$\circledast \xrightarrow{u} p \xrightarrow{v} \circledast \,.$$

accepting some word uv, that will be reduced (from Lemma 5.4.9). Since $\widetilde{\mathcal{L}}(\Gamma) \subseteq \widetilde{\mathcal{L}}(\Gamma')$, we can read the same reduced word in Γ', and from lemmas 5.4.9 and 5.4.10, there exists a reduced walk

$$\odot' \overset{u}{\rightsquigarrow} p' \overset{v}{\rightsquigarrow} \odot',$$

where p' is some vertex in Γ'. We claim that such a vertex p' is uniquely determined, and the corresponding map

$$\begin{aligned} \phi\colon V(\Gamma) &\to V(\Gamma') \\ p &\mapsto p' \end{aligned} \tag{5.18}$$

(sending basepoint to basepoint) obtained in this way defines a homomorphism between Γ and Γ'.

For the first claim, we note that the argument used in Proposition 5.4.12 works for reduced walks. Finally, to see that ϕ is a homomorphism of automata, consider a labelled arc $e \equiv p \xrightarrow{a} q$ in Γ. Then (from Lemma 5.4.19), there exist a reduced $\odot$-walk in Γ of the form:

$$\odot \overset{u}{\rightsquigarrow} p \xrightarrow{a} q \overset{v}{\rightsquigarrow} \odot$$

for some words u, v. Now, using $\widetilde{\mathcal{L}}(\Gamma) \subseteq \widetilde{\mathcal{L}}(\Gamma')$ together with Lemma 5.4.9 (in both sides), we have that there exists a reduced $\odot$-walk in Γ' of the form:

$$\odot \overset{u}{\rightsquigarrow} p\phi \xrightarrow{a} q\phi \overset{v}{\rightsquigarrow} \odot\,.$$

In particular, we have an arc $p\phi \xrightarrow{a} q\phi$ in Γ'. Thus $\phi\colon \Gamma \to \Gamma'$ is a homomorphism of automata, as we wanted to prove. □

Finally, combining Proposition 5.4.12 with Proposition 5.4.24, we obtain one of the main goals of Stallings theory: the family X-**Inverse** of inverse X-automata univocally represent the subgroups of the free group with basis X.

Corollary 5.4.25. *Let Γ, Γ' be inverse X-automata. Then, the following statements are equivalent:*

1. *Γ and Γ' are isomorphic (i.e., $\Gamma \simeq \Gamma'$).*
2. *Γ and Γ' recognize the same language (i.e., $\mathcal{L}_\Gamma = \mathcal{L}_{\Gamma'}$).*
3. *Γ and Γ' recognize the same reduced language (i.e., $\widetilde{\mathcal{L}}_\Gamma = \widetilde{\mathcal{L}}_{\Gamma'}$).*
4. *Γ and Γ' recognize the same subgroup of $\mathbb{F}_X$ (i.e., $\langle \Gamma \rangle = \langle \Gamma' \rangle$).* □

Below, we see that for every involutive X-automaton, there exists a (unique) inverse automaton equivalent to it. Note that a constructive proof of this fact (see Theorem 5.4.43) can not be given in the general setting.

Definition 5.4.26. Let Γ be an involutive X-automaton. We say that two vertices p, q in Γ are equivalent, denoted by $\mathsf{p} \equiv \mathsf{q}$, if there exist a walkbetween both vertices recognizing the trivial element in $\mathbb{F}_X$. That is,

$$\mathsf{p} \equiv \mathsf{q} \overset{\text{(def)}}{\Longleftrightarrow} \exists w \in \mathbb{F}_X : \mathsf{p} \xrightarrow{1_{\mathbb{F}_X}} \mathsf{q}\,.$$

Lemma 5.4.27. *Equivalence of vertices is an equivalence relation compatible with the automaton structure; i.e., the map*

$$\begin{array}{ccc} \mathsf{V}\Gamma & \twoheadrightarrow & \mathsf{V}\Gamma/\equiv \\ \mathsf{p} & \mapsto & [\mathsf{p}] \end{array}$$

induces (identifying parallel arcs with the same label) a well-defined epimorphism of automata $\Gamma \twoheadrightarrow \Gamma/\equiv$. □

Proposition 5.4.28. *Let* Γ *be an* X*-automaton. Then, the core of the quotient automaton* $\Gamma/\equiv$ *is an inverse* X*-automaton recognizing the same subgroup as* Γ*. It is called the* Stallings reduction *of* Γ*, and denoted* $\mathsf{St}(\Gamma)$.

Proof. It is enough to note that any path in the quotient graph $\Gamma/\equiv$ can be lifted to a path in Γ whose label represents the same element in $\mathbb{F}_X$. This is so because given any path of length 2 in $\Gamma/\equiv$,

$$[\mathsf{p}] \xrightarrow{x_1} [\mathsf{q}] \xrightarrow{x_2} [\mathsf{r}]\,, \tag{5.19}$$

there must exist 1-paths $\mathsf{p}_1 \xrightarrow{x_1} \mathsf{q}_1$ and $\mathsf{q}_2 \xrightarrow{x_2} \mathsf{r}_1$ in Γ, where p_1 and r_1 are preimages of $[\mathsf{p}]$ and $[\mathsf{r}]$ respectively, and $\mathsf{q}_1, \mathsf{q}_2$ are (maybe different) preimages of $[\mathsf{q}]$. And since $\mathsf{q}_1, \mathsf{q}_2$ are equivalent, from Definition 5.4.26 there must also exist walks in Γ from the basepoint $\odot$ to q_1 and q_2 reading the same element, say $w \in \mathbb{F}_X$. But then,

$$\mathsf{p}_1 \xrightarrow{x_1} \mathsf{q}_1 \xrightarrow{w^{-1}} \odot \xrightarrow{w} \mathsf{q}_2 \xrightarrow{x_2} \mathsf{r}_1$$

is a walk in Γ recognizing x_1x_2 (the same element recognized by (5.19)). This completes the proof. □

Now, Corollary 5.4.25 allows us to use the notion of Stallings automorphism modulo equivalent automata.

Definition 5.4.29. Let H be a subgroup of a free group $\mathbb{F}_{\mathsf{X}}$. Then, the *Stallings automaton* of H with respect to X, denoted by $\mathsf{St}(\mathsf{H},\mathsf{X})$ is the Stallings reduction of any X-automata recognizing H; i.e.,

$$\mathsf{St}(\mathsf{H},\mathsf{X}) \coloneqq \operatorname{core}(\mathsf{St}(\Gamma)),$$

where Γ is an X-automata recognizing H.

If there is no ambiguity about the underlying basis (e.g. when it is clear that we refer to the one in the given presentation), we will sometimes omit any reference to it and simply talk about the Stallings automaton of the subgroup H, denoted by $\mathsf{St}(\mathsf{H})$.

Proposition 5.4.30. *Let H be a subgroup of a free group $\mathbb{F}_{\mathsf{X}}$. Then, the Stallings automaton of H with respect to X is isomorphic to the core of the Schreier automaton of H; i.e.,*

$$\mathsf{St}(\mathsf{H},\mathsf{X}) \simeq \operatorname{core}(\mathsf{Sch}(\mathsf{H},\mathsf{X})).$$

Proof. By construction, $\operatorname{core}(\mathsf{Sch}(\mathsf{H},\mathsf{X}))$ is trim, deterministic, and recognizes the subgroup H. Then, by the uniqueness in Corollary 5.4.25, we are done. □

Since the unique (from Corollary 5.4.25) inverse X-automaton recognizing a subgroup $\mathsf{H} \leqslant \mathbb{F}_{\mathsf{X}}$ always exists (Definition 5.4.29), we can finally state the characterization theorem for subgroups of arbitrary free groups (note that we are not assuming the subgroups, or even the ambient free group to be finitely generated).

Notation 5.4.31. For every arc e outside $\mathbf{T}$, we will denote by $\gamma_{\mathsf{T}}[\mathsf{e}]$ the (unique) reduced $\odot$-walk trough $\mathbf{T}$, across e, and back to $\odot$ through $\mathbf{T}$; i.e.,

$$\gamma_{\mathsf{T}}[\mathsf{e}] \equiv \odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathsf{e}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot \quad \text{(reduced)};$$

and we will denote by $\mathsf{w}_{\mathsf{T}}[\mathsf{e}]$ its label, i.e.,

$$\mathsf{w}_{\mathsf{T}}[\mathsf{e}] = \ell(\gamma_{\mathsf{T}}[\mathsf{e}]) = \widetilde{\ell}(\odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathsf{e}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot).$$

(Note that, since $\gamma_{\mathsf{T}}[\mathsf{e}]$ is reduced, and Γ is inverse, the word $\mathsf{w}_{\mathsf{T}}[\mathsf{e}]$ is also reduced.)

Theorem 5.4.32 (Stallings, 1983, [Sta83])**.** *Let $\mathbb{F}_{\mathsf{X}}$ be a free group with basis X. Then, the map*

$$\begin{array}{rcl} \{\textit{Subgroups of } \mathbb{F}_{\mathsf{X}}\} & \to & \{\textit{Inverse } \mathsf{X}\textit{-automata}\} \\ \mathsf{H} & \mapsto & \mathsf{St}(\mathsf{H},\mathsf{X}) \end{array}, \tag{5.20}$$

(sending every subgroup $\mathsf{H} \leqslant \mathbb{F}_{\mathsf{X}}$ to its Stallings automaton) is a bijection with inverse $\Gamma \mapsto \langle \Gamma \rangle$.

Moreover, for any spanning tree $\mathbf{T}$ *of an inverse* X*-automaton* Γ*, the set*

$$\mathrm{B}_{\Gamma,\mathbf{T}} \;=\; \left\{ \mathrm{w}_{\mathbf{T}}[\mathrm{e}] \,:\, \mathrm{e} \in \mathrm{E}^{+}(\Gamma \smallsetminus \mathbf{T}) \right\}, \tag{5.21}$$

is a free basis for the subgroup H *recognized by* Γ*. It is called the canonical basis for* H *with respect to* $(\Gamma, \mathbf{T})$*.*

Proof. The bijectivity of (5.20) is clear from Corollary 5.4.25.

On the other hand, from Proposition 5.3.16, we have that the family $\mathrm{B}_{\Gamma,\mathbf{T}}$ in (5.21) is a generating set for H (recall that $\ell_{\mathbb{F}_X} = \widetilde{\ell}$, and Notation 5.4.31). Now we prove that, if Γ is an inverse automaton, then $\mathrm{B}_{\Gamma,\mathbf{T}}$ is indeed a free family (and thus a basis of H).

It is enough to see that any nontrivial word w *reduced in* $\mathrm{B}_{\Gamma,\mathbf{T}}$ (i.e., such that

$$\mathrm{w} \;=\; \mathrm{w}_{\mathbf{T}}[\mathrm{e}_{i_1}]\, \mathrm{w}_{\mathbf{T}}[\mathrm{e}_{i_2}] \,\cdots\, \mathrm{w}_{\mathbf{T}}[\mathrm{e}_{i_k}]\,, \tag{5.22}$$

where $\mathrm{e}_{i_j} \in \mathrm{E}(\Gamma \smallsetminus \mathbf{T})$, and no two successive e_{i_j}'s are inverse of each other) represents a nontrivial element in $\mathbb{F}_X$ (i.e., that $\widetilde{\mathrm{w}} \neq 1$). Since $\odot$-walks within trees read trivial group elements, we have:

$$\begin{aligned}
\widetilde{\mathrm{w}} \;&=\; \widetilde{\ell}\,(\gamma_{\mathbf{T}}[\mathrm{e}_{i_1}]\, \gamma_{\mathbf{T}}[\mathrm{e}_{i_2}] \,\cdots\, \gamma_{\mathbf{T}}[\mathrm{e}_{i_k}]) \\
&=\; \widetilde{\ell}\,(\odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathrm{e}_{i_1}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathrm{e}_{i_2}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot \cdots \odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathrm{e}_{i_k}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot) \\
&=\; \widetilde{\ell}\,(\odot \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathrm{e}_{i_1}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathrm{e}_{i_2}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \bullet \cdots \bullet \overset{\mathbf{T}}{\rightsquigarrow} \bullet \xrightarrow{\mathrm{e}_{i_k}} \bullet \overset{\mathbf{T}}{\rightsquigarrow} \odot)\,.
\end{aligned} \tag{5.23}$$

Note that the walk in (5.23) is nontrivial and already reduced (since the arcs e_{i_j} are outside $\mathbf{T}$). But, since Γ is deterministic, this means (Corollary 5.4.11) that $\widetilde{\mathrm{w}} \neq 1$, as we wanted to prove. □

Corollary 5.4.33. *For any subgroup* H *of a free group* $\mathbb{F}_X$*,* $\operatorname{rk} \mathrm{H} = \operatorname{rk} \mathrm{St}(\mathrm{H}, X)$*. In particular:*

$$\mathrm{H} \textit{ is finitely generated } \Leftrightarrow \mathrm{St}(\mathrm{H}, X) \textit{ is finite}\,. \qquad \square$$

5.4.1 Finitely generated subgroups of free groups

For finitely generated subgroups of free groups, the previous characterization can be made algorithmic. In this case, it is enough to sequentially remove from any (finite) starting automata the obvious redundances coming from non-determinism, and non-trimness. To this end, let us introduce two kinds of transformations on inverse automata that obviously keep the recognized language unchanged: *foldings* and *trims*.

We call foldings the natural identifications of arcs preformed in order to fix local non-determinism, namely the situations shown in Figure 5.6.

Definition 5.4.34. Let $\Gamma = (\mathrm{V}, \mathrm{E}, \iota, \tau, \ell, \circledast)$ be an involutive X-automaton, and let e, f be two positive arcs in Γ with the same origin and label (say $\mathrm{e} \equiv \mathrm{p} \xrightarrow{a} \mathrm{q}$, and $\mathrm{f} \equiv \mathrm{p} \xrightarrow{a} \mathrm{r}$, where q and r are not necesarily different).

Then, the involutive X-automaton Γ' obtained by identifying e and e' (and its respective inverses) in Γ is called an (elementary) *Stallings reduction* of e and f in Γ.

Formally, $\Gamma' = (\mathrm{V}', \mathrm{E}', \iota', \tau', \ell', \circledast)$ is given by:

1. vertex set $\mathrm{V}' = (\mathrm{V} \setminus \{\mathrm{q}, \mathrm{r}\}) \sqcup \{\mathrm{q}'\}$,
2. arc set $\mathrm{E}' = \mathrm{E} \setminus \mathrm{E}_{\{\mathrm{q},\mathrm{r}\}} \sqcup \left\{ \mathrm{f}', \forall \mathrm{f} \in \mathrm{E}_{\{\mathrm{q},\mathrm{r}\}} \right\}$, where
   2.1. for all $\mathrm{f} \in \mathrm{E}_{\{\mathrm{q},\mathrm{r}\}}$, if the origin $\iota\mathrm{f}$ (resp. end $\tau\mathrm{f}$) of f is q or r, then the origin (resp. end) of f' is q',
   2.2. all other incidence relation for an arc $\mathrm{e}' \in \mathrm{E}'$ is inherited from the corresponding arc $\mathrm{e} \in \mathrm{E}$.
3. labelling
   3.1. for all $\mathrm{f} \in \mathrm{E}_{\{\mathrm{q},\mathrm{r}\}}$, $\ell'(\mathrm{f}') = \ell(\mathrm{f})$,
   3.2. $\ell'(\mathrm{e}) = \ell(\mathrm{e})$ for every other arc $\mathrm{e} \in \mathrm{E}'$.

Then, the map

$$\begin{aligned} \varphi_{\{\mathrm{e},\mathrm{f}\}} \colon \Gamma &\to \Gamma' \\ \mathrm{q}, \mathrm{r} &\mapsto \mathrm{q}' \\ \mathrm{s} &\mapsto \mathrm{s}, \text{ for all } \mathrm{s} \notin \{\mathrm{q}, \mathrm{r}\}, \end{aligned}$$

is an epimorphism of automata called (elementary) *Stallings folding* of the arcs e and f. The action of a folding transformation on Γ to obtain Γ' is denoted by

$$\Gamma \xrightarrow{\varphi} \Gamma'.$$

Note that an inverse automaton is deterministic if and only if it has no available foldings (the term *folded automaton* is also used).

The following results are clear by construction.

Proposition 5.4.35. *Foldings do not change the language (and hence the subgroup) recognized by an involutive automaton.* □

Remark 5.4.36. A folding decreases the number of arcs by 1 (2 if we count inverse arcs). So, if the starting automaton is finite, we will reach a folded automaton after a finite number of foldings.

Definition 5.4.37. Let Γ be a finite inverse automaton. Then, we define the *loss* of a folding $\varphi\colon \Gamma \to \Gamma'$ to be the difference of the ranks of Γ and Γ'; i.e.,

$$\operatorname{loss}\varphi = \operatorname{rk}\Gamma - \operatorname{rk}\Gamma' .$$

The loss of a finite sequence of elementary foldings is the sum of the losses of each folding in the sequence.

The result below is immediate from the definition of automaton rank.

Proposition 5.4.38. *The loss of a folding on a finite inverse automaton is* 0 *if the identified arcs are nonparallel, and* 1 *otherwise.* □

Definition 5.4.39. We can divide foldings in two disjoint types:

(a) *Open foldings* (with loss 0): those which are identifications of two nonparallel adjacent arcs with the same labels (and its corresponding inverses).

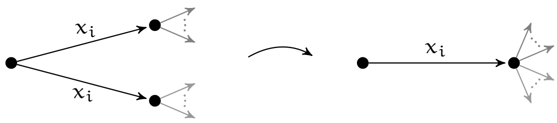

Fig. 5.7: Open folding

(b) *Closed foldings* (with loss 1): i.e., those which identify two parallel arcs with the same label (and its corresponding inverses).

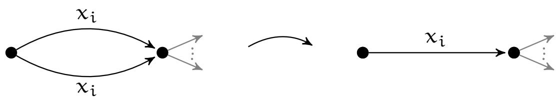

Fig. 5.8: Closed folding

Remark 5.4.40. Note that the natural act of neglecting an explicit cancellation determined by two non mutually inverse arcs within an induced walk:

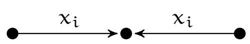

in an inverse X-automata, can be achieved by a folding followed by a trim. So we can, without loss of generality, avoid such situations in the considered automata.

The computability of finite inverse automata is now straightforward.

Proposition 5.4.41. *Any finite* X*-automaton* Γ *recognizing* H *can be converted into an inverse automata — recognizing the same subgroup* H *— after a finite number of foldings, followed by a total trim.*

Proof. This is clear, because each folding decreases the number of arcs by one. Since the starting automaton is finite, after a finite number of foldings we will reach a deterministic automaton, that will be inverse after performing the final complete trim. □

$$\Gamma \xrightarrow{\varphi^{(1)}} \Gamma^{(1)} \xrightarrow{\varphi^{(2)}} \cdots \xrightarrow{\varphi^{(p)}} \Gamma^{(p)} \xrightarrow{\tau} \Gamma^{(p+1)} \tag{5.24}$$

Fig. 5.9: Folding process followed by a total trim

Corollary 5.4.42. *The Stallings automaton of a finitely generated subgroup* H *of a free group is computable given any finite generating set for* H. □

Proof. Given S a finite generating set for H, it is enough to compute the flower automaton $\mathsf{Fl}(\mathsf{S})$ and then apply Proposition 5.4.41 on it. Recall that $\mathsf{Fl}(\mathsf{S})$ (an so, every automaton in the folding sequence) is trim, therefore, the final T transformation in (5.24) is not necessary in this case. □

The computation of the Stallings automaton of a finitely generated subgroup of the free group is well known to be fast (in [Tou06], Touikan shows it to be $\mathcal{O}(\mathsf{N}\log^* \mathsf{N})$, where N is the total number of letters in the input generating set).

On the other side, given a finite inverse X-automaton Γ, there exist well known algorithms to compute a spanning tree, and hence — using (5.21) — a basis for the recognized subgroup $\langle \Gamma \rangle$. The version of Theorem 5.4.32 for finitely generated subgroups follows.

Theorem 5.4.43 (Stallings, 1983, [Sta83])**.** *Let* $\mathbb{F}_\mathsf{X}$ *be a finitely generated free group with basis* X*. Then, the map*

$$\begin{array}{rcl} \{\textit{f.g. subgroups of } \mathbb{F}_\mathsf{X}\} & \to & \{\textit{Finite inverse } \mathsf{X}\textit{-automata}\} \\ \mathsf{H} & \mapsto & \mathsf{St}(\mathsf{H},\mathsf{X}) \end{array} \tag{5.25}$$

sending a finitely generated subgroup $\mathsf{H} \leqslant \mathbb{F}_\mathsf{X}$ *to its (finite) Stallings automaton, is a computable bijection.* □

Corollary 5.4.44. *For any finitely generated subgroup* H *of a free group* $\mathbb{F}$,

$$\operatorname{rk}\mathsf{H} = \#\mathsf{E}\Gamma - \#\mathsf{V}\Gamma + 1,$$

where Γ *is the Stallings automaton of* H. □

5.5 Some applications of Stallings theory

The results on the last section provide a natural geometric reinterpretation and proof of many classic results about subgroups of free groups. The first natural application is the solvability of the membership problem for free groups. We state this problem (a.k.a. generalized word problem) below, for an arbitrary finitely presented group G.

(Subgroup) membership problem, $\mathsf{MP}(G)$**.** *Given a finite set of words* $w, w_1, \ldots, w_n$*, in the generators of* G*, decide whether* w *represents an element in* $\langle w_1, \ldots, w_n \rangle_G$*; and, in affirmative case, compute an expression of* w *as a word in the* w_i*'s.*

Theorem 5.5.1. *The subgroup membership problem is decidable for free groups.*

Proof. Note that, for this problem, one can always assume the ambient group to be finitely generated: indeed, for any finite input $S = \{w, w_1, \ldots, w_n\} \subseteq \mathbb{F}_X$, it is enough to consider the (finitely generated) free ambient with basis the support set consisting of all the generators of X appearing in S.

Let $w, w_1, \ldots, w_n$ be reduced words in some (we can assume finite) free basis X. In order to decide whether w belongs to the subgroup generated by $S = \{w_1, \ldots, w_n\}$, it is enough to compute the Stallings automaton $\Gamma = \mathsf{St}(\langle S \rangle)$, and check whether the word w is recognized by Γ. Since Γ recognizes exactly the elements in the subgroup $\langle S \rangle$, and this test can be performed algorithmically, the solvability of the membership (decision) problem follows.

For the search problem, let $w, w_1, \ldots, w_n$ be freely reduced words in X, such that w belongs to the subgroup generated by $S = \{w_1, \ldots, w_n\}$. Once built the Stallings automaton Γ of $\langle S \rangle$, the spelling of the word w will draw a $\odot$-walk γ in Γ. Then, unfold back Γ to the flower automaton $\mathsf{Fl}(S)$ keeping track of the walk γ. The walk obtained from γ in $\mathsf{Fl}(S)$ will describe w as a product of the starting generators $w_1, \ldots, w_n$. □

Remark 5.5.2. Recall that, although the membership search problem is always solvable by brute force (see Figure 5.49), the Stallings construction allows a much more efficient search of the expression of a word $w \in H$ in terms of a finite family of generators S: namely, lifting back the closed walk in $\mathsf{St}(H)$ reading w through the tower of foldings to the Flower automaton $\mathsf{Fl}(S)$.

Also, recall that Theorem 5.4.32 and Theorem 5.4.43 contain the classic results below.

Proposition 5.5.3. *Two free groups are isomorphic if and only if their respective bases have the same cardinality, i.e.,*

$$\mathbb{F}_X \simeq \mathbb{F}_Y \;\Leftrightarrow\; \#X = \#Y\,. \qquad \square$$

Theorem 5.5.4 (Nielsen–Schreier Theorem)**.** *Every subgroup of a free group is itself free.* $\square$

Theorem 5.5.5. *Bases for finitely generated subgroups of free groups are computable (from any finite family of generators).* $\square$

The close relationship between the Stallings automaton and the Schreier digraph of a subgroup of a free group, provides useful information. For example, if the ambient group is finitely generated, the following implications are straightforward:

$$\begin{aligned} [\,\mathbb{F} : H\,] < \infty &\iff \mathsf{Sch}\,(H) \text{ is finite} \\ &\implies \mathsf{St}\,(H) \text{ is finite} \implies H \text{ is f.g.} \end{aligned} \tag{5.26}$$

(Note that for the last implication, a finitely generated ambient is necessary.)

Indeed, a more precise statement can be made.

Proposition 5.5.6. *Let* $\mathbb{F}_n$ *a free group of finite rank* n*, and* $H \leqslant \mathbb{F}_n$*. Then,*

$$[\,\mathbb{F}_n : H\,] < \infty \iff \mathsf{St}\,(H) \textit{ is complete,} \tag{5.27}$$

and, if so,

1. *The index of* H *in* $\mathbb{F}_n$ *is the number of vertices of* $\mathsf{St}\,(H)$*.*
2. *Schreier index formula:*

$$\operatorname{rk} H - 1 = [\mathbb{F}_n : H] \cdot (n-1)\,. \tag{5.28}$$

Proof. The equivalence (5.27) and the first claim, are clear from the identity $\mathsf{St}\,(H, X) = \mathsf{core}(\mathsf{Sch}\,(H, X))$, assuming X to be finite. For the second claim, recall that $\operatorname{rk} H = \operatorname{rk} \mathsf{St}\,(H, X)$, and use Euler's formula (5.6). $\square$

Therefore, we can algorithmically decide whether a subgroup of $\mathbb{F}_n$ *given by a finite family of generators* has finite index, just building its Stallings automaton, and then checking whether every vertex has degree $2n$. Note also that, in affirmative case (if every vertex is complete), we can immediately obtain a family of coset representatives just reading words from the basepoint to the vertices.

Theorem 5.5.7. *The finite index problem is solvable for free groups of finite rank.* $\square$

The relation between the rank of a subgroup and its Stallings automaton also provides easily the following classic result.

Theorem 5.5.8. *Finitely generated free groups are Hopfian (and not co-Hopfian).*

Proof. The non-coHopfianicity of free groups is evident (to prove it in our setting, it is enough to to build a Stallings automaton of rank n not being equal to the bouquet).

For the Hopfianicity, let S be a generating set for a subgroup $H \leqslant \mathbb{F}_n$. We want to prove that, if $\#S = \operatorname{rk} H$, then S is indeed a basis for H. It is enough to observe that then, both the flower automaton $\mathsf{Fl}(S)$, and the Stallings automaton $\mathsf{St}(H)$ have the same rank, and thus no closed foldings are possible in the folding process. The result follows. □

However, this does not mean that, in a free group, any minimal generating set in $\mathbb{F}_n$ has exactly n elements. Indeed, using the Stallings automata, it is not difficult to build examples satisfying a slightly stronger claim.

Lemma 5.5.9. *There exist minimal generating sets for $\mathbb{F}_n$ of any (finite) size $r \geqslant n$.* □

Conjugation and normality also have natural translations in the geometric setting. For example, the following result is straightforward from the graphic interpretation.

Proposition 5.5.10. *Two subgroups $H, K \leqslant \mathbb{F}_X$ are conjugate, say $H = K^w$, for certain $w \in \mathbb{F}$, if and only if their Schreier coset digraphs $\mathsf{Sch}(H)$ and $\mathsf{Sch}(K)$ coincide except maybe for the location of the basepoint. In this case, the respective basepoints are connected through a walk reading the conjugating word w.* □

The corollaries below follow easily from Proposition 5.5.10.

Corollary 5.5.11. *A subgroup of a free group is normal if and only if its Schreier digraph is vertex-transitive.* □

Remark 5.5.12. If H is a nontrivial normal subgroup of the free group, then its Schreier digraph is core (and hence coincides with its Stallings automaton)

Corollary 5.5.13. *A nontrivial subgroup $H \leqslant \mathbb{F}_X$ is normal if and only if its Stallings automaton $\mathsf{St}(H, X)$:*

(i) *is complete (i.e., is $X^{\pm}$-regular), and*

(ii) *has a vertex-transitive subjacent digraph (i.e., for every vertex p in $\mathsf{St}(H)$, there exist an automorphism* of digraphs *of* $\mathsf{St}(H)$, *sending the basepoint to p).* □

Remark 5.5.14. The Stallings automaton of a nontrivial normal subgroup of $\mathbb{F}_n$ is $2n$-regular.

Corollary 5.5.15. *A nontrivial normal subgroup of $\mathbb{F}_n$ is finitely generated if and only if it has finite index.*

Proof. The converse implication always holds in a finitely generated ambient, see (5.26). The direct implication follows from Remark 5.5.12: if H is nontrivial and normal in $\mathbb{F}_n$, then its (finite) Stallings automaton coincides with its Schreier digraph. Since the vertices of the Schreier digraph are precisely the right cosets of H in $\mathbb{F}_n$, the result follows. □

Corollary 5.5.16. *There exists an algorithm which, given $w_1, \ldots, w_k \in \mathbb{F}_n$, decides whether the subgroup $H = \langle w_1, \ldots, w_k \rangle$ is normal in $\mathbb{F}_n$; and, if not, finds a word $w \in \mathbb{F}_n$ such that $H^w \neq H$.* □

5.5.1 Intersection problems for free groups

Problems concerning intersections within free groups deserve its own section, not only as a nice application of Stallings machinery, but also because their importance in modern geometric group theory, and the protagonism they have throughout the whole dissertation (see Figure 5.49).

For free groups, the first important result on this topic is the following theorem of Howson.

Theorem 5.5.17 (Howson, 1954, [How54]). *The intersection of two finitely generated subgroups of a free group is again finitely generated.* □

The property that appears in the theorem above has aroused so much interest that finally took the name of the author who first considered it.

Definition 5.5.18. A group G is said to satisfy *Howson's property* (or to be a *Howson group*, for short), if the intersection of any two of its finitely generated subgroups is again finitely generated.

Remark 5.5.19. Note that for this problem, the rank of the ambient group can be assumed to be finite. Namely, we only need to consider generators appearing in any pair of finite generating sets for the subgroups.

After Theorem 5.5.17, the natural question about the rank of the intersection subgroup became popular, and ended up being one of the most famous open problems in geometric group theory during the last decades. The first (partial) answer to this question was given by Howson himself, providing also examples of subgroups $H, K \leqslant \mathbb{F}_n$ where $\operatorname{rk}(H \cap K) - 1 = (\operatorname{rk} H - 1)(\operatorname{rk} K - 1)$.

Shortly after, Hanna Neumann improved Howson's bound.

Proposition 5.5.20 (Neumann, 1956, [Neu56]). *If* H, K *are non-trivial finitely generated subgroups of a free group, then:*

$$\operatorname{rk}(H \cap K) - 1 \leqslant 2\,(\operatorname{rk} H - 1)\,(\operatorname{rk} K - 1)\,. \tag{5.29}$$

In the same paper Neumann conjectured that the bound still holds if we remove the factor 2 from the right term in (5.29). This last claim has become known as the "Hanna Neumann conjecture" (see Theorem 5.5.21).

Many partial proofs, related claims, and generalizations of this conjecture have been done since its statement [Bur71; Neu90; Tar92; Dic94; Tar96; Arz00; DF01; Iva01; MW02; Kha02; Iva08; DI08; Wis05], but no full proof was given since Igor Mineyev and Joel Friedman independently proved it quite recently (see also the remarkable two-page version of the proof using Bass-Serre theory by Warren Dicks, available at [Dic12], and the related paper [Nos16]).

Theorem 5.5.21 (Mineyev, 2012, [Min12]; Friedman, 2015, [Fri15]). *If* H, K *are non-trivial finitely generated subgroups of a free group, then:*

$$\operatorname{rk}(H \cap K) - 1 \leqslant (\operatorname{rk} H - 1)\,(\operatorname{rk} K - 1)\,. \tag{5.30}$$

(Note that, according to Howson's examples, this bound is tight.) □

Below, we see how to use Stallings automata to obtain Howson's result, and Hanna Neumann bound Equation (5.29), and indeed compute a basis for the intersection, given generators for the intersecting subgroups.

Definition 5.5.22. Let $\Gamma = (V, E, \iota, \tau, \ell, \circledast)$ and $\Gamma' = (V', E', \iota', \tau', \ell', \circledast')$ be X-automata. Then, the *(tensor) product* or *pull-back* of Γ and Γ', denoted by $\Gamma \times \Gamma'$, is the pointed X-digraph:

1. with vertex set the cartesian product $V \times V'$,
2. such that an arc $(p, p') \longrightarrow (q, q')$ exists and is labelled by $x \in X$, if and only if there exist both an arc $p \xrightarrow{x} q$ in Γ, and an arc $p' \xrightarrow{x} q'$ in Γ', and
3. having the vertex $(\circledast, \circledast')$ as basepoint.

The following lemmas contain claims which are clear from the definition.

Lemma 5.5.23. *If* Γ, Γ' *are deterministic automata, then the product* $\Gamma \times \Gamma'$ *is also deterministic. Moreover, the degree of any vertex* (p, p') *in* $\Gamma \times \Gamma'$ *is at most equal to the minimum of the degrees of* p *and* p' *in* Γ *and* Γ' *respectively.* □

Remark 5.5.24. The product of two inverse (i.e., involutive, deterministic, and trim) automata is again involutive and deterministic, but not necessarily connected nor trim.

Lemma 5.5.25. *The language recognized by (the connected component containing the basepoint in) the product $\Gamma \times \Gamma'$ is exactly the intersection of the languages recognized by Γ and Γ'.*

In particular, if two languages $\mathcal{L}, \mathcal{L}'$ are regular, then their intersection is also regular.

Finally, removing the contributions from the possible superfluous vertices in the product we reach the automata recognizing subgroup intersections.

Definition 5.5.26. Let Γ, Δ be inverse X-automata. Then, the core of the product $\Gamma \times \Delta$ is called the *junction automaton* of Γ and Δ, and is denoted by $\Gamma \wedge \Delta$; i.e.,

$$\Gamma \wedge \Delta \coloneqq \mathrm{core}(\Gamma \times \Delta)\,.$$

Proposition 5.5.27. *Let H, K be finitely generated subgroups of a free group. Then, the Stallings automaton of the intersection $H \cap K$, is the junction of the Stallings automata of H and K; that is:*

$$\mathsf{St}\,(H \cap K) \,=\, \mathsf{St}\,(H) \wedge \mathsf{St}\,(K)\ .$$

Proof. It is enough to realize that the junction automaton $\mathsf{St}\,(H) \wedge \mathsf{St}\,(K)$:

1. recognizes the subgroup $H \cap K$. This is clear from Lemma 5.5.25, since the only missing words can be obtained by inserting word-cancellations.
2. is deterministic (Lemma 5.5.23), and
3. is trim (obvious because it is a core automata). □

Remark 5.5.28. Since the product of two finite automata is finite, Proposition 5.5.27 automatically proves Howson's property for free groups, and it is not difficult to also obtain Hanna Neuman's bound in (5.29) from the product construction. Namely, if H, H' are finitely generated subgroups of $\mathbb{F}_n$, let us denote by Γ and Γ' the Stallings automaton of H and H' respectively. Then, removing from Γ and Γ'

basepoint threads if necessary (note that this does not affect the rank, and so our argument):

$$
\begin{aligned}
2(\operatorname{rk} H-1)(\operatorname{rk} H'-1) &= 2(\#E\Gamma-\#V\Gamma)(\#E\Gamma'-\#V\Gamma') \\
&= 2\left(\frac{1}{2}\sum\nolimits_{p\in V\Gamma}(\deg p-2)\right)\left(\frac{1}{2}\sum\nolimits_{p'\in V\Gamma'}(\deg p'-2)\right) \\
&= \frac{1}{2}\sum\nolimits_{p\in V\Gamma}\sum\nolimits_{p'\in V\Gamma'}(\deg p-2)(\deg p'-2) \\
&= \frac{1}{2}\sum\nolimits_{(p,p')\in V\Gamma\times V\Gamma'}(\deg p-2)(\deg p'-2) \\
&\geqslant \frac{1}{2}\sum\nolimits_{(p,p')\in V(\Gamma\times\Gamma')}(\deg(p,p')-2)^2 \\
&\geqslant \frac{1}{2}\sum\nolimits_{(p,p')\in V(\Gamma\times\Gamma')}\left|\deg(p,p')-2\right| \\
&\geqslant \frac{1}{2}\sum\nolimits_{(p,p')\in V(\mathrm{St}(H\cap H'))}(\deg(p,p')-2) \\
&= \#E(\mathrm{St}\,(H\cap H'))-\#V(\mathrm{St}\,(H\cap H')) \\
&= \operatorname{rk}(H\cap H')-1,
\end{aligned}
$$

where we have used the rank formula (5.4.44), the handshaking lemma, and the inequalities $\deg(p,p') \leqslant \max(\deg p, \deg p')$, and $\#V(\mathrm{St}\,(H\cap H')) \leqslant \#V(\Gamma\times\Gamma')$.

Finally, both the decision intersection problem (which obviously always answers YES, since free groups are Howson), and the computability of a basis for the intersection of two finitely generated subgroups, are clear consequences of Proposition 5.5.27.

Theorem 5.5.29. *The (full) subgroup intersection problem is solvable for free groups. That is, there exists an algorithm which, given elements $u_1,\dots,u_k,v_1,\dots,v_l \in \mathbb{F}_X$, computes a basis for the intersection $\langle u_1,\dots,u_k\rangle \cap \langle v_1,\dots,v_l\rangle$.* □

Example 5.5.30. Let $H = \langle x^3, yx, y^3xy^{-2}\rangle$, and $H' = \langle x^2, yxy^{-1}\rangle \leqslant F_2 = \langle x,y \mid -\rangle$. In order to compute (a basis for) the intersection $H\cap H'$ we first compute the Stallings automata $\mathrm{St}\,(H)$ and $\mathrm{St}\,(H')$, and then build the product:

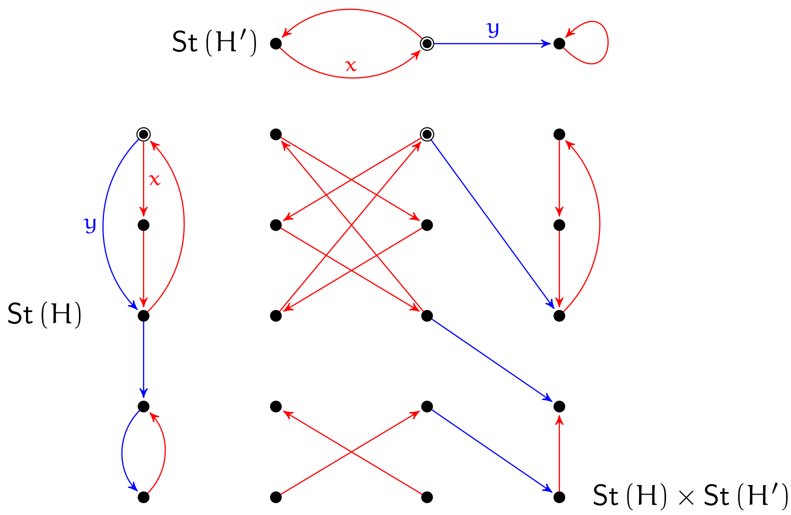

Fig. 5.10: Pull-back of St (H) and St (H′)

Note that the product in Figure 5.10 is neither connected nor trim; thus, in order to obtain the Stallings automata of the intersection $H \cap H'$, we need to consider the connected component in $\mathsf{St}(\Gamma) \times \mathsf{St}(\Delta)$ containing the basepoint, and then apply a total trim. After rearranging the resulting automata we get:

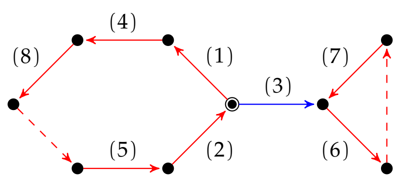

Fig. 5.11: Stallings automaton of $H \cap H'$

where the dashed arcs indicate the arcs outside the spanning tree built according to the order $x \prec x^{-1} \prec y \prec y^{-1}$ in the generating set $\{x, y\}^{\pm}$ (the labels indicate the order of appearance of each arc in this construction). Note that Figure 5.10 also shows that, for the purpose of intersection, the element y^3xy^{-2} in H (producing only a hanging thread, and a non-basic disconnected component in the product) is completely irrelevant.

Now, we can use Theorem 5.4.43 to compute the basis $\{x^6, yx^3y^{-1}\}$ for the intersection $H \cap H'$, which therefore has rank 2.

Finally we state without proof two more easy consequences of the pullback construction for free groups.

Theorem 5.5.31. *The (full) coset intersection problem is solvable for free groups.* □

Corollary 5.5.32. *A subgroup* H *of a free group is malnormal if and only if every connected component of the product of* $\mathrm{St}(H)$ *with itself is a tree, except the diagonal one (which is clearly isomorphic to* $\mathrm{St}(H)$*).* □

5.6 Free-abelian by free groups

We aim to extend the Stallings machinery for free groups to a broadened scenario. Namely, let

- $\mathbb{Z}^m = \langle T \mid [T, T] \rangle$ be a free-abelian group of finite rank m, with free-abelian basis $T = \{t_1, \ldots, t_m\}$,
- $\mathbb{F}_X = \langle X \mid - \rangle$ be a free group (of arbitrary rank) with basis $X = \{x_j\}_{j \in J}$, and
- $\mathbf{A}_j \in \mathrm{GL}_m(\mathbb{Z})$, for each $j \in J$, be automorphisms of $\mathbb{Z}^m$;

and consider the semidirect product $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$, with action the homomorphism given by

$$\begin{array}{rcl} \mathbf{A}_\bullet \colon \mathbb{F}_X & \to & \mathrm{GL}_m(\mathbb{Z}) \\ x_j & \mapsto & \mathbf{A}_j \,. \end{array} \tag{5.31}$$

Notation 5.6.1. Of course, once fixed a free-abelian basis, we can interpret the automorphisms in $\mathrm{GL}_m(\mathbb{Z})$ as invertible integer matrices, that we will assume acting on the right. That is, we will write $\mathbf{bA}$ the image of $\mathbf{b} \in \mathbb{Z}^m$ under the matrix $\mathbf{A}$, and $\mathbf{AB}$ the composition $\mathbb{Z}^m \xrightarrow{\mathbf{A}} \mathbb{Z}^m \xrightarrow{\mathbf{B}} \mathbb{Z}^m$.

If $w = w(\overrightarrow{x_j}) = x_{j_1}^{\epsilon_1} x_{j_2}^{\epsilon_2} \cdots x_{j_p}^{\epsilon_p}$ is a word in X, then we will write

$$\mathbf{A}_w := w(\overrightarrow{\mathbf{A}_j}) = \mathbf{A}_{j_1}^{\epsilon_1} \mathbf{A}_{j_2}^{\epsilon_2} \cdots \mathbf{A}_{j_p}^{\epsilon_p}$$

(the product of matrices obtained replacing each appearance of a letter x_j in w, by the corresponding matrix $\mathbf{A}_j$). Note that $\mathbf{A}_w = \mathbf{A}_{\mathit{w}}$, for each word $w = w(X)$ representing an element $\mathit{w} \in \mathbb{F}_X$; in particular, $\mathbf{A}_\lambda = \mathbf{A}_1 = \mathbf{I_m}$. In the same vein, we denote by $\mathbf{A}_S$ the (set of matrices) image of any subset $S \subseteq \mathbb{F}_X$ under the action $\mathbf{A}_\bullet$; i.e., $\mathbf{A}_S = \{\mathbf{A}_{\mathit{w}} : \mathit{w} \in S\}$. Note that then, $\mathbf{A}_{\langle S \rangle} = \langle \mathbf{A}_S \rangle \leqslant \mathrm{GL}_m(\mathbb{Z})$.

With the previous notation, our target group can be written:

$$\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m = \left\langle T, X \;\middle|\; \begin{array}{ll} t_i t_k = t_k t_i & \forall i, k \in [1, m] \\ x_j^{-1} t_i x_j = t_i \mathbf{A}_j & \forall i \in [1, m], \forall j \in J \end{array} \right\rangle . \tag{5.32}$$

Remark 5.6.2. The case where $\mathbf{A}_\bullet \colon x_j \mapsto \mathbf{I_m}$ (for all j) corresponds exactly to the direct product $\mathbb{F}_X \times \mathbb{Z}^m$. Thus, all the results in this chapter apply to the groups considered in Part I and [DV13], and some of them constitute generalizations of the results found there.

Given a word $w = w(X, T)$, we can always use the relations in (5.32) to move t's to one side (say to the right), and obtain a normal form for the element represented by w.

Definition 5.6.3. A *normal form* for the element in $\mathbb{G}_{\mathbf{A}}$ represented by the word w is defined to be

$$w\, t^{\mathbf{a}} \;=\; w \cdot t_1^{a_1} t_2^{a_2} \cdots t_m^{a_m}\,, \tag{5.33}$$

where w is the element in $\mathbb{F}_X$ represented by the free part of w, (i.e., $w = w\pi_{\mathbb{F}} \in \mathbb{F}_X$), $\mathbf{a} = (a_1, a_2, \ldots, a_m) \in \mathbb{Z}^m$; and we have abbreviated $t^{\mathbf{a}} := t_1^{a_1} t_2^{a_2} \cdots t_m^{a_m}$. In particular, if $\{\mathbf{e}_i\}_i$ is the canonical basis in $\mathbb{Z}^m$, then $t^{\mathbf{e}_i} = t_i$, for all $i \in [1, m]$. We say that w and $t^{\mathbf{a}}$ are respectively the *free part*, and *free-abelian* part of the word $w\, t^{\mathbf{a}}$.

Remark 5.6.4. The abbreviated notation $w\, t^{\mathbf{a}}$ allows us to think elements of $\mathbb{F}_X \ltimes \mathbb{Z}^m$ as pairs $(w, \mathbf{a}) \in \mathbb{F}_X \ltimes \mathbb{Z}^m$ but with multiplicative behaviour, which is convenient when working in a noncommutative group.

With this notation, the semidirect conjugation relation in $\mathbb{G}_{\mathbf{A}}$ becomes

$$w^{-1}\, t^{\mathbf{a}}\, w \;=\; t^{\mathbf{a}\mathbf{A}_w}\ \text{,where } \mathbf{a} \in \mathbb{Z}^m\,,\ \text{and } w \in \mathbb{F}_X\,.$$

This means that a subword $t^{\mathbf{a}}$ can jump to the right (resp. left) of a free subword $w \in \mathbb{F}_X$ at the price of applying the matrix $\mathbf{A}_w$ (resp. $\mathbf{A}_w^{-1}$) to the vector $\mathbf{a}$; i.e.,

$$t^{\mathbf{a}}\, w \;=\; w\, t^{\mathbf{a}\mathbf{A}_w} \quad \text{and} \quad w\, t^{\mathbf{a}} \;=\; t^{\mathbf{a}\mathbf{A}_w^{-1}}\, w \;=\; t^{\mathbf{a}\mathbf{A}_{-w}}\, w\,. \tag{5.34}$$

Remark 5.6.5. Obviously, the normal form of a given word is always computable. So (for computable purposes) we can always assume, without loss of generality, the elements in $\mathbb{G}_{\mathbf{A}}$ given in normal form.

Iterating the rules in (5.34) we obtain the expressions below, that will be useful later.

Lemma 5.6.6. *Let* $w(\overrightarrow{u_i\, t^{\mathbf{a}_i}}) \;=\; u_{i_1}\, t^{\mathbf{a}_{i_1}}\, u_{i_2}\, t^{\mathbf{a}_{i_2}} \cdots u_{i_p}\, t^{\mathbf{a}_{i_p}}$ *be a word of length* $p \geqslant 1$ *in* $\{u_i\, t^{\mathbf{a}_i}\}_i \subseteq \mathbb{Z}^m \times \mathbb{F}_n$. *Then*

$$\begin{aligned} w(\overrightarrow{u_i\, t^{\mathbf{a}_i}}) \;&=\; u_{i_1}\, t^{\mathbf{a}_{i_1}}\, u_{i_2}\, t^{\mathbf{a}_{i_2}} \cdots u_{i_p}\, t^{\mathbf{a}_{i_p}} \\ &=_{\mathbb{G}_{\mathbf{A}}}\; u_{i_1} u_{i_2} \cdots u_{i_p}\, t^{\mathbf{a}_{i_1} \mathbf{A}_{u_{i_2}} \cdots \mathbf{A}_{u_{i_p}}}\, t^{\mathbf{a}_{i_2} \mathbf{A}_{u_{i_3}} \cdots \mathbf{A}_{u_{i_p}}} \ldots t^{\mathbf{a}_{i_{p-1}} \mathbf{A}_{u_{i_p}}}\, t^{\mathbf{a}_{i_p}} \\ &=_{\mathbb{G}_{\mathbf{A}}}\; w(\overrightarrow{u_i})\; t^{\sum_{j=1}^{p} \left(\mathbf{a}_{i_j} \prod_{k=j+1}^{p+1} \mathbf{A}_{u_{i_k}}\right)}\,, && (5.35) \\ &=_{\mathbb{G}_{\mathbf{A}}}\; w(\overrightarrow{u_i})\; t^{\sum_{j=1}^{p} \mathbf{a}_{i_j} \mathbf{A}_{w[j+1,p](\overrightarrow{u_i})}}\,, && (5.36) \end{aligned}$$

where $w[p+1, p]$ *is the empty word, and so* $\mathbf{A}_{w[p+1,p](\overrightarrow{u_i})} = \mathbf{I}_m$. □

Below, we summarize the expressions for the simplest situations.

Corollary 5.6.7. *Let* $u, u_1, \ldots, u_p \in \mathbb{F}_n$*,* $\mathbf{a}, \mathbf{a_1}, \ldots, \mathbf{a_p} \in \mathbb{Z}^m$*, where* $p \geqslant 1$*. Then,*

(i) $\prod_{i=1}^{p} u_i\, t^{\mathbf{a_i}} = \left(\prod_{i=1}^{p} u_i\right) t^{\sum_{i=1}^{p}\left(\mathbf{a_i} \prod_{j=i+1}^{p+1} \mathbf{A_j}\right)}$*, where* $\mathbf{A_{p+1}} = \mathbf{I_m}$.

(ii) $(u\, t^{\mathbf{a}})^p = u^p\, t^{\mathbf{a}\left(\sum_{i=1}^{p} \mathbf{A}_u^{p-i}\right)} = u^p\, t^{\mathbf{a}\left(\sum_{i=0}^{p-1} \mathbf{A}_u^{i}\right)}$.

(iii) $(u\, t^{\mathbf{a}})^{-1} = u^{-1}\, t^{-\mathbf{a}\mathbf{A}_u^{-1}}$.

(Note that $(u\, t^{\mathbf{a}})^{\epsilon} = u^{\epsilon}\, t^{(\epsilon\mathbf{a})\mathbf{A}_u^{(\epsilon-1)/2}}$*, for* $\epsilon = \pm 1$*.)* □

Recall that semidirect products correspond to split short exact sequences (see Section 3.2); in our case we can take:

$$1 \longrightarrow \mathbb{Z}^m \longrightarrow \mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes \mathbb{Z}^m \xrightarrow{\pi_{\mathbb{F}}} \mathbb{F}_X \longrightarrow 1\,, \tag{5.37}$$

with the natural "identity section" $w \mapsto w\, t^{\mathbf{0}}$ of $\pi_{\mathbb{F}}$. Note that we have omitted the action name (as we will often do hereinafter) in order to lighten notation.

Definition 5.6.8. According to this description, the groups $\mathbb{F}_X \ltimes \mathbb{Z}^m$ are called *free-abelian by free* (FABF), and $\mathbb{Z}^m$ is said to be the *full basegroup* of $\mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$ with quotient $\mathbb{F}_X$.

Remark 5.6.9. Note that one does not necessarily need the full basegroup $\mathbb{Z}^m$ to generate the whole group $\mathbb{G}_{\mathbf{A}}$ through the action $\mathbf{A}_\bullet$. Namely, any (abelian) contribution that can be obtained by conjugating other abelian elements is not necessary in order to generate the whole group (as far as we dispose of the conjugated element).

Definition 5.6.10. A *basegroup* for $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$ (with respect to $\mathbb{Z}^m$) is any subgroup $L \leqslant \mathbb{Z}^m$ that normally generates the full basegroup; i.e., any group that satisfies one (and thus all) of the following equivalent conditions:

(a) $\langle L \cup X \rangle_{\mathbb{G}_{\mathbf{A}}} = \mathbb{G}_{\mathbf{A}}$.

(b) $\langle\!\langle L \rangle\!\rangle_{\mathbb{G}_{\mathbf{A}}} = \mathbb{Z}^m$.

(c) $\sum_{i=1}^{n} L\mathbf{A_i} = \mathbb{Z}^m$.

It is clear that the full basegroup is unique, and the largest possible basegroup for $\mathbb{G}_{\mathbf{A}}$.

5.6.1 Subgroups of free-abelian by free groups

Our objective is to associate to every subgroup $H \leqslant \mathbb{F}_X \ltimes \mathbb{Z}^m$ a "unique" (in a sense that will be detailed) automaton recognizing exactly the elements in H. Let us first describe these subgroups from an algebraic point of view.

Proposition 5.6.11. *Let $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$ as in (5.32). Then, any subgroup $H \leqslant \mathbb{G}_{\mathbf{A}}$ admits the decomposition*

$$H \simeq H\pi_{\mathbb{F}} \ltimes (H \cap \mathbb{Z}^m)\,, \tag{5.38}$$

where $H \cap \mathbb{Z}^m$ is the full basegroup of H, and the action is

$$\begin{array}{rcl} H\pi_{\mathbb{F}} & \to & GL(H \cap \mathbb{Z}^m) \\ u & \mapsto & \mathbf{A}_{u|H\cap\mathbb{Z}^m} \colon \left\{ \begin{array}{rcl} H\cap\mathbb{Z}^m & \to & H\cap\mathbb{Z}^m \\ t^{\mathbf{b}} & \mapsto & t^{\mathbf{b}\mathbf{A}_u} \end{array} \right. . \end{array} \tag{5.39}$$

The expression (5.38) is called the split decomposition *of H as FABF group.*

In particular, every subgroup of a free-abelian by free group is again free-abelian by free.

Proof. Let $H = \langle \{w_k\, t^{\mathbf{a}_k}\}_k \rangle$ be a subgroup of $\mathbb{G}_{\mathbf{A}}$, and consider the restriction of the short exact sequence (5.37) to H:

$$1 \longrightarrow H \cap \mathbb{Z}^m \longrightarrow H \xrightarrow{\pi_{\mathbb{F}|H}} H\pi_{\mathbb{F}} \longrightarrow 1\,. \tag{5.40}$$

Now it is clear that:

- $H\pi_{\mathbb{F}}$ is a subgroup of $\mathbb{F}_X$ with basis, say $\{v_l\}_l$, where each v_l admits an expression $v_l = v_l(\overrightarrow{w_k})$ as a word in the generators $\{w_k\}_k$ of $H\pi_{\mathbb{F}}$.

Moreover, since for every $t^{\mathbf{b}} \in H \cap \mathbb{Z}^m$, and every $ut^{\mathbf{a}} \in H$,

$$(ut^{\mathbf{a}})^{-1} t^{\mathbf{b}} (ut^{\mathbf{a}}) = t^{-\mathbf{a}} u^{-1} t^{\mathbf{b}} u t^{\mathbf{a}} = t^{-\mathbf{a}} u^{-1} u t^{\mathbf{b}\mathbf{A}_u} t^{\mathbf{a}} = t^{-\mathbf{a}+\mathbf{b}\mathbf{A}_u+\mathbf{a}} = t^{\mathbf{b}\mathbf{A}_u} \in H \cap \mathbb{Z}^m\,,$$

we also have that:

- $\ker \pi_{\mathbb{F}|H} = H \cap \mathbb{Z}^m$ is (a free-abelian subgroup of $\mathbb{Z}^m$ of rank at most m) normal in H. Note that this also proves that $\forall u \in H\pi_{\mathbb{F}}$, $H \cap \mathbb{Z}^m$ is invariant under $\mathbf{A}_u$, and thus the homomorphism (5.39) is well defined.
- The map $v_l \mapsto v_l(\overrightarrow{w_k t^{\mathbf{a}_k}})$ defines a section $\sigma \colon H\pi_{\mathbb{F}} \to H$ of $\pi_{\mathbb{F}|H}$ which acts by conjugation on $H \cap \mathbb{Z}^m$ according (5.39).

This completes the proof. ☐

Corollary 5.6.12. *Let $\mathbb{G}_{\mathbf{A}}$ be a free-abelian by free group, and H a subgroup of $\mathbb{G}_{\mathbf{A}}$. Then,*

$$H \text{ is finitely generated} \;\Leftrightarrow\; H\pi_{\mathbb{F}} \text{ is finitely generated}\,,$$

where $\pi_{\mathbb{F}}\colon \mathrm{t}^{\mathbf{a}}\, u \mapsto u$ *is the natural projection in* (5.37).

In particular, if the action is trivial, a particularly simple description for the rank of a subgroup follows immediately from (5.38).

Corollary 5.6.13. *Let* H, H' *be subgroups of a free-abelian times free group* $\mathbb{F}_X \times \mathbb{Z}^m$. *Then,*

$$\mathrm{rk}(H) \;=\; \mathrm{rk}(H \cap \mathbb{Z}^m) + \mathrm{rk}(H\pi_{\mathbb{F}})\,, \qquad \square$$

Below we generalize to FABF groups a concept already used in Part I.

Definition 5.6.14. Given a subgroup $H \leqslant \mathbb{F}_X \ltimes \mathbb{Z}^m$, and an element $w \in \mathbb{F}_n$, we define the *abelian completion of* w *in* H as

$$\mathcal{C}_H(w) \;=\; \{\mathbf{a} \in \mathbb{Z}^m \;:\; w\,\mathrm{t}^{\mathbf{a}} \in H\} \;\subseteq\; \mathbb{Z}^m.$$

In the same vein, the *full completion* of an element $w \in \mathbb{F}_n$ in H is defined to be the set $w\,\mathrm{t}^{\mathcal{C}_H(w)} \coloneqq \{w\,\mathrm{t}^{\mathbf{a}} \;:\; \mathbf{a} \in \mathcal{C}_H(w)\}$ if $w \in H\pi_{\mathbb{F}}$, and the empty set otherwise.

Corollary 5.6.15. *According to Lemma 5.6.6, if* $\{u_i\, \mathrm{t}^{\mathbf{a}_i}\}_i$ *is a generating set for a subgroup* $H \leqslant \mathbb{F}_X \ltimes \mathbb{Z}^m$, *and* $w \in \mathbb{F}_X$, *then:*

(a) *if* $w \notin H\pi_{\mathbb{F}}$, *then* $\mathcal{C}_H(w) = \varnothing$, *and*

(b) *if* $w \in H\pi_{\mathbb{F}}$, *then* $\mathcal{C}_H(w) = \sum_{j=1}^{p} \left(\mathbf{a}_{i_j} \prod_{k=j+1}^{p+1} \mathbf{A}_{u_{i_k}}\right) + L\,,$
where $w = \mathrm{w}(\overrightarrow{u_i}) = u_{i_1} u_{i_2} \cdots u_{i_p}$ *is an expression of* w *as a word in the free projection of the generating set for* H.

So, every completion $\mathcal{C}_H(w) \subseteq \mathbb{Z}^m$ *is either empty, or a coset of* $L = H \cap \mathbb{Z}^m$. $\square$

5.7 Enriched automata

The following discussion suggests natural restrictions on the family of involutive automata (see Section 5.3), in order to describe subgroups of free-abelian by free groups.

Remark 5.7.1. Given S a family of words (we can assume in normal form) in $T \sqcup X$ generating a subgroup $H \leqslant \mathbb{F}_X \ltimes \mathbb{Z}^m$, recall that we denote by S_T the set of words $\{\mathrm{w}_j(T)\}_j \subseteq S$ using only letters in T. Then, it is obvious that S_T generates a (f.g. free-abelian) subgroup $\langle\{\mathbf{w}_j\}_j\rangle \leqslant H \cap \mathbb{Z}^m$, where $\mathbf{w}_j = \mathrm{w}_j^{\mathrm{ab}}$. So, in the flower automaton $\mathsf{Fl}(S)$, we could abbreviate — say as a label attached to the basepoint — all the abelian petals by the subgroup $L = \langle S_T \rangle$ they generate, which is obviously computable if S is finite.

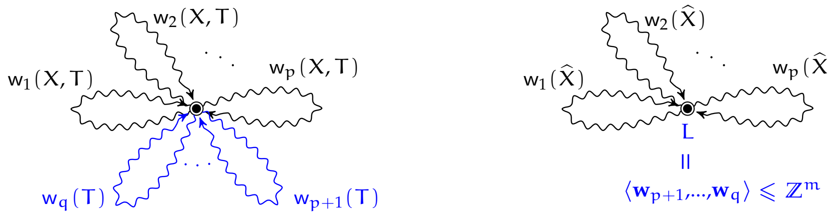

Fig. 5.12: The standard (left), and abbreviated (right) flower automaton $\mathsf{Fl}(S)$ of the set $S = \{w_1(X,T), w_2(X,T), \ldots, w_p(X,T), w_{p+1}(T), \ldots, w_q(T)\}$ of words in the generators of the FABF group (5.32)

Remark 5.7.2. If $w = x_{j_1}^{\epsilon_1} x_{j_2}^{\epsilon_2} \cdots x_{j_p}^{\epsilon_p}$ is a word in $X^{\pm}$, and $\mathbf{b} \in \mathbb{Z}^m$. Then, the element in $\mathbb{F}_X \ltimes \mathbb{Z}^m$ described by the word $wt^{\mathbf{b}}$ can be written in the form

$$(t^{-\mathbf{a}_1} x_{j_1}^{\epsilon_1} t^{\mathbf{b}_1})\,(t^{-\mathbf{a}_2} x_{j_2}^{\epsilon_2} t^{\mathbf{b}_2})\,\cdots\,(t^{-\mathbf{a}_p} x_{j_p}^{\epsilon_p} t^{\mathbf{b}_p}) \tag{5.41}$$

(for some $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_p, \mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p \in \mathbb{Z}^m$) if and only if

$$-\mathbf{a}_1 \mathbf{A}_{j_1}^{\epsilon_1} \mathbf{A}_{j_2}^{\epsilon_2} \cdots \mathbf{A}_{j_p}^{\epsilon_p} + (\mathbf{b}_1 - \mathbf{a}_2)\mathbf{A}_{j_2}^{\epsilon_2} \cdots \mathbf{A}_{j_p}^{\epsilon_p} + \cdots + (\mathbf{b}_{p-1} - \mathbf{a}_p)\mathbf{A}_{j_p}^{\epsilon_p} + \mathbf{b}_p \;=\; \mathbf{b}\,, \tag{5.42}$$

$$\sum\nolimits_{k=1}^{p+1} (\mathbf{b}_{k-1} - \mathbf{a}_k) \prod\nolimits_{i=k}^{p} \mathbf{A}_{j_i}^{\epsilon_i} \;=\; \mathbf{b}\,, \tag{5.43}$$

where we are assuming $\mathbf{a}_0 = \mathbf{b}_{p+1} = \mathbf{0}$. Hence, any noncentral element in $\mathbb{F}_X \ltimes \mathbb{Z}^m$ can be written as a product of *enriched noncentral generators* $\widehat{x}_i = t^{-\mathbf{a}_i} x_i t^{\mathbf{b}_i}$, where $\mathbf{a}_i, \mathbf{b}_i \in \mathbb{Z}^m$, in infinitely many ways.

To provide flexibility in the forthcoming transformations (see Section 5.8), it will we convenient to enable our automata to recognize any fragmentation of the form (5.41) of words describing elements in the target subgroup.

The conjunction of the last two remarks provides a compact definition for "flower automaton" within free-abelian by free groups, and suggests a convenient (restricted) family of automata for describing their subgroups.

The following definitions are meant to provide a formal description of these automata.

Definition 5.7.3. A $\mathbb{Z}^m$*-enriched* X*-automaton* (EA) $\widehat{\Gamma}_L$ is a pointed $(\mathbb{Z}^m \times X \times \mathbb{Z}^m)$-automaton, with a subgroup of $\mathbb{Z}^m$ attached to its basepoint.

In more detail, a $\mathbb{Z}^m$-enriched X-automaton is a tern $\widehat{\Gamma}_L = (\Gamma, \widehat{\ell}, L)$, consisting of:

1. A pointed digraph Γ (the *underlying digraph* of $\widehat{\Gamma}_L$).

2. An arc-labelling $\widehat{\ell}\colon E\Gamma \to \mathbb{Z}^m \times X \times \mathbb{Z}^m$ (the *enriched labelling* of $\widehat{\Gamma}_L$).

3. A subgroup $L \leqslant \mathbb{Z}^m$ (the *basepoint subgroup* of $\widehat{\Gamma}_L$).

Remark 5.7.4. So, $\mathbb{Z}^m$-enriched X-automata are nothing more than convenient abbreviations for standard $\{T, X\}$-automata in the context of free-abelian by free groups. Namely, the basepoint subgroup is meant to abbreviate a set of free-abelian petals attached to the basepoint (see Remark 5.7.1); and the enriched arcs, graphically represented:

$\widehat{e} \equiv$ 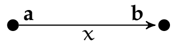

Fig. 5.13: An enriched arc

That is, the *middle label* is a "free letter" $x \in X$, and the two extreme free-abelian labels $\mathbf{a}, \mathbf{b} \in \mathbb{Z}^m$ (named *initial and final labels* of the $\widehat{e}$, respectively) are meant to represent the oriented $T \sqcup X$-walk reading $t^{-\mathbf{a}} \, x \, t^{\mathbf{b}}$.

That is, the abelian labels $\mathbf{a}, \mathbf{b}$ can be interpreted as abbreviations for the directed T-walks reading respectively $t^{-\mathbf{a}}$ and $t^{\mathbf{b}}$ before and after an arc labelled by x.

It will be usually clear from context which label interpretation are we using for an enriched arc $\widehat{e}$ as in (5.13); but to avoid any confusion, we will distinguish between:

$$\widehat{e} = (\mathbf{a}, x, \mathbf{b}) \qquad \text{and} \qquad \ell(\,\widehat{e}\,) = t^{-\mathbf{a}} \, x \, t^{\mathbf{b}} \,,$$

where we have slightly abused notation (in the first one).

Notation 5.7.5. A $\mathbb{Z}^m$-enriched X-automaton can be always thought as an standard X-automaton enriched with free-abelian labels (elements at the extremes of every edge, and a subgroup labelling the basepoint). In this context, we generically accentuate with a hat "$\widehat{\mathcal{O}}$" to mean "object $\mathcal{O}$" with its arcs enriched with abelian labels. We use this notation in order to refer to any $\mathbb{Z}^m$-enriched object through its X-skeleton (see Definition 5.7.7).

So, an enriched automata $\widehat{\Gamma}_L$ is a standard X-automata Γ (its X-skeleton), with and abelian arc-labelling (denoted with a hat, $\widehat{\Gamma}$), plus a free-abelian subgroup attached to its basepoint (denoted subscript, $\widehat{\Gamma}_L$). In the same vein, a walk will be denoted by $\widehat{\gamma}$ or γ depending on whether it is supposed to be read in $\widehat{\Gamma}$ (i.e., including abelian labels) or just in Γ. For example, we may denote by $\widehat{e}$ an arc with enriched label $(\mathbf{a}, x, \mathbf{b})$, and by e the same arc with label x.

Example 5.7.6. The enriched arc

$\widehat{e} \equiv$ 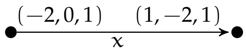

corresponds to the path

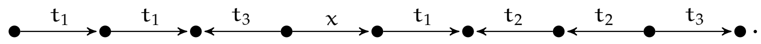

So, we have:

$$\widehat{e} = ((-2,0,1), x, (1,-2,1)), \qquad \ell(\widehat{e}) = t^{(-2,0,1)}\, x\, t^{(1,-2,1)},$$
$$e = (\mathbf{0}, x, \mathbf{0}), \qquad \ell(e) = x.$$

The previous conventions immediately provide a natural translation for null abelian labels within enriched automata: we can neglect any abelian label equal to $\mathbf{0}$ in $\widehat{\Gamma}_L$. For example, we write $\widehat{\Gamma}_{\mathbf{0}} = \widehat{\Gamma}$, if the basepoint subgroup is trivial.

In the same vein, we identify any enriched automata $\widehat{\Gamma}_L$ having all its abelian labels equal to $\mathbf{0}$, with the X-automaton Γ obtained by removing them. Hence, we can think standard automata as particular (null abelian) cases of enriched automata.

Definition 5.7.7. The X-*skeleton* of an enriched automaton $\widehat{\Gamma}_L$, denoted by $\mathrm{sk}(\widehat{\Gamma}_L)$, is the X-automaton Γ obtained after removing all the abelian labels from $\widehat{\Gamma}_L$.

Notation 5.7.8. Whenever possible, we will denote enriched automata by $\widehat{\Gamma}_L$, which allows us to refer to its skeleton simply by Γ.

The meaning of a walk and its label within an enriched automaton is a natural adaptation of Definitions 5.2.7 and 5.2.15 taking into account the meaning of the basepoint label (see Remark 5.7.1).

Definition 5.7.9. Let $\widehat{\Gamma}_L$ be an enriched automaton. A *walk* $\widehat{\gamma}$ in $\widehat{\Gamma}_L$ is either a single vertex in $\widehat{\Gamma}_L$, or a (strongly adjacent) sequence of enriched arcs $\widehat{e}_1 \cdots \widehat{e}_k$ in $\widehat{\Gamma}_L$ such that the corresponding sequence of arcs $e_1 \cdots e_k$ is a walk in the skeleton Γ.

Then, the *label of an enriched walk* in $\widehat{\Gamma}_L$ is recursively defined as follows:

- The label of a non-basepoint vertex is the empty word λ (i.e., $\ell(\bullet) = \lambda$).
- The label of the basepoint vertex is the set t^L (i.e., $\ell(\odot) = t^L$).
- The label of a nontrivial walk $\widehat{\gamma} \equiv \widehat{e}_1 \cdots \widehat{e}_k$ is

$$\ell(\widehat{\gamma}) = \ell(\widehat{e}_1 \cdots \widehat{e}_{k-1})\, \ell(\iota e_k)\, t^{-\mathbf{a}_k}\, \ell(e_k)\, t^{\mathbf{b}_k} \ell(\tau e_k),$$

  where ιe_k and τe_k are respectively the initial and terminal vertices of e_k.

Note that Definition 5.7.9 extends unambiguously the notion of label in an standard automaton; i.e., if γ is a walk in the skeleton Γ of $\widehat{\Gamma}_L$, then $\ell(\gamma) = \ell(e_1) \cdots \ell(e_k)$ both as automaton label, and as enriched automaton label of Γ.

As usual, if $\mu_G \colon ((X \sqcup T)^{\pm})^* \to \mathbb{G}_{\mathbf{A}}$ is the natural projection, then we denote by $\ell_{\mathbb{G}_{\mathbf{A}}}(\widehat{\gamma})$ the element in $\mathbb{G}_{\mathbf{A}}$ represented by $\ell(\widehat{\gamma})$ (i.e., $\ell_{\mathbb{G}_{\mathbf{A}}}(\widehat{\gamma}) \coloneqq (\ell(\widehat{\gamma}))\mu_G$).

Corollary 5.7.10. *Let $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$ be as in (5.32), and consider a basepoint-free enriched walk $\widehat{\gamma} = \widehat{e}_1 \widehat{e}_2 \cdots \widehat{e}_p$, with $\widehat{e}_k = (\mathbf{a}_k, x_{j_k}, \mathbf{b}_k)$. Then,*

$$\ell_{\mathbb{G}_{\mathbf{A}}}(\widehat{\gamma}) = \ell(\gamma)\, t^{\sum_{k=1}^{p+1} (\mathbf{b}_{k-1} - \mathbf{a}_k) \prod_{i=k}^{p} \mathbf{A}_{j_i}}, \tag{5.44}$$

where γ is the X-skeleton of $\widehat{\gamma}$, and we assume $\mathbf{a}_0 = \mathbf{b}_{p+1} = \mathbf{0}$.

In particular, if $\widehat{e} = (\mathbf{a}, x, \mathbf{b})$, then $\ell_{\mathbb{G}_{\mathbf{A}}}(\widehat{e}) = x\, t^{\mathbf{b} - \mathbf{a}\mathbf{A}_x}$. □

Finally, our interpretation of enriched automata as a particular type of standard automata (see Remark 5.7.4) also provides a natural notion for involutive automata in this context.

Definition 5.7.11. An enriched automata is said to be *involutive* if for every arc labelled $(\mathbf{a}, x, \mathbf{b}) \in \mathbb{Z}^m \times X^{\pm} \times \mathbb{Z}^m$, there is always a unique reversed arc labelled $(-\mathbf{b}, x^{-1}, -\mathbf{a})$.

That is, in an involutive enriched automata, for every arc positively labelled $(\mathbf{a}, x_j, \mathbf{b})$, we are always assuming the existence of a (unique) reversed arc, dashed in the following picture, but usually omitted in the graphic representation.

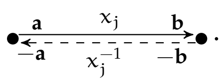

Fig. 5.14: Arcs in an involutive enriched automaton

Remark 5.7.12. Put another way, we have the following rule of thumb. Any visible (positive) arc 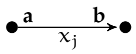 in an involutive enriched automaton is meant to be read:

- $t^{-\mathbf{a}}\, x_j\, t^{\mathbf{b}}$ when crossed forward (from left to right), and
- $t^{-\mathbf{b}}\, x_j^{-1}\, t^{\mathbf{a}}$ when crossed backwards (from right to left).

In particular, the skeleton of an involutive $(\mathbb{Z}^m, X)$-EA is an involutive X-automaton.

The notions of determinism and trimness also admit a definition on enriched automata, conveniently adapted in terms of its interpretation as standard automata. However, we will see in Definition 5.9.2 that a much more concise characterization can be given.

As usual, we only need to depict positive arcs to describe involutive enriched automata.

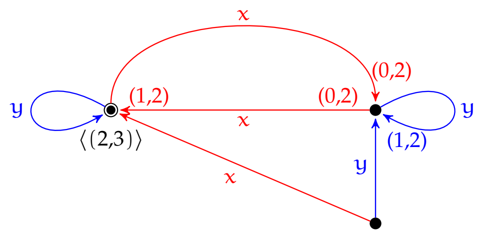

Fig. 5.15: A $\mathbb{Z}^2$-enriched $\{x, y\}$-automaton

We have reached the *restricted* family of automata that we will use to represent subgroups of free-abelian by free groups $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes \mathbb{Z}^m$; namely, that of *involutive enriched automata*. Hereinafter, if not stated otherwise, we will assume all the enriched automata appearing in the chapter to be involutive.

The description of the subgroup of $\mathbb{G}_{\mathbf{A}}$ recognized by an involutive enriched automaton follows from the general Definition 5.3.17 and the previous conventions.

Proposition 5.7.13. *Let $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$, and let $\widehat{\Gamma}_L$ be a $\mathbb{Z}^m$-enriched X-automaton. Then, the subgroup of $\mathbb{G}_{\mathbf{A}}$ recognized by $\widehat{\Gamma}_L$ is:*

$$\langle \widehat{\Gamma}_L \rangle \simeq \langle \Gamma \rangle \ltimes \overline{L^\Gamma}\,, \tag{5.45}$$

where $\overline{L^\Gamma}$ is the normal closure of L in $\langle \widehat{\Gamma}_L \rangle$ (see Definition 5.8.1), and the action is given by $\mathbf{A}_{\langle \Gamma \rangle | \overline{L^\Gamma}}$.

In particular, the projection to the free part of the subgroup recognized by an enriched automaton $\widehat{\Gamma}_L$ is exactly the subgroup recognized by its skeleton Γ; i.e., if H is the subgroup recognized by $\widehat{\Gamma}_L$, then $\langle \Gamma \rangle = H\pi_{\mathbb{F}}$. □

Remark 5.7.14. Note that, since flower automata are clearly of this kind, involutive enriched automata are enough to describe all the subgroups of $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$. Concretely, the map

$$\begin{array}{rcl} \{\text{involutive } \mathbb{Z}^m\text{-enriched } X\text{-automata}\} & \to & \{\text{subgroups of } \mathbb{G}_{\mathbf{A}}\} \\ \widehat{\Gamma}_L & \mapsto & \langle \widehat{\Gamma}_L \rangle_{\mathbb{G}_{\mathbf{A}}} \end{array} \tag{5.46}$$

is well defined and onto. However, this is clearly very far from being injective (which is our goal). Indeed, the sources of looseness within the preimage in (5.46) of a given subgroup are well determined. Namely, looseness can appear:

(a) in the basepoint subgroup L, either through the action $\mathbf{A}_\bullet$ (see Remark 5.8.3), or through parallel X-arcs (see Remark 5.8.7).

(b) by rearranging the abelian arc-labels according multiplication rules in $\mathbb{G}_{\mathbf{A}}$.

(c) in the skeleton Γ, either by nondeterminism or nontrimness (as it happened in the free case, see Section 5.4).

We will see in the next section that these cases contain all possible ambiguity in an enriched automata recognizing a subgroup of $\mathbb{G}_{\mathbf{A}}$.

5.8 Transformations on enriched automata

In this section we introduce different kinds of transformations on enriched automata in order to fix all the kinds of looseness detailed in Remark 5.7.14, and hence compute a unique representative among all the enriched automata recognizing a *finitely generated* subgroup of $\mathbb{G}_{\mathbf{A}}$.

Let $\widehat{\Gamma}_{\mathrm{L}}$ be an involutive $\mathbb{Z}^m$-enriched X-automata recognizing a subgroup H of $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$. First of all, recall that $\odot$-walks in $\widehat{\Gamma}$ acting on L by conjugation can produce abelian elements in $H \cap \mathbb{Z}^m$ which are not included in L; i.e., the subgroup $L \leqslant H \cap \mathbb{Z}^m$ is not necessarily $\mathbf{A}_{H\pi_{\mathbb{F}}}$-invariant. In order to deal with this kind of looseness, we introduce the notion below.

Definition 5.8.1. Let $\widehat{\Gamma}_{\mathrm{L}}$ be an $\mathbb{Z}^m$-enriched X-automaton recognizing H in $\mathbb{G}_{\mathbf{A}}$. Then, the *closure of* L *in* $\widehat{\Gamma}_{\mathrm{L}}$ (with respect to $\mathbf{A}_\bullet$) is the subgroup

$$\overline{\mathrm{L}}^{\Gamma} := \langle\!\langle \mathrm{L} \rangle\!\rangle_{\langle \widehat{\Gamma}_{\mathrm{L}} \rangle} = \mathrm{L}^{\langle \widehat{\Gamma}_{\mathrm{L}} \rangle} = \mathrm{L}^{\langle \Gamma \rangle} = \mathrm{L}^{H\pi_{\mathbb{F}}} = (\mathrm{L})\mathbf{A}_{H\pi_{\mathbb{F}}} \leqslant H \cap \mathbb{Z}^m . \tag{5.47}$$

If the acting skeleton Γ is clear, we will usually omit any reference to it, and denote the closure of the basepoint L simply by $\overline{\mathrm{L}}$.

The basepoint group L is said to be *closed* (in $\widehat{\Gamma}_{\mathrm{L}}$ with respect to $\mathbf{A}_\bullet$) if it equals its own closure; i.e., if $\mathrm{L} = \overline{\mathrm{L}}^{\Gamma}$. An enriched automaton having a closed basepoint subgroup is also called *closed*.

Again, since the ambient action $\mathbf{A}_\bullet$ is given from the start, we will usually omit any reference to it, and talk simply about closed basepoint group, tacitly assuming the given action.

Remark 5.8.2. The closed subgroup $\overline{\mathrm{L}}^{\Gamma}$ is $\mathbf{A}_{H\pi_{\mathbb{F}}}$-invariant by construction.

Remark 5.8.3 (Basepoint subgroup looseness of type I). Let $\widehat{\Gamma}_{\mathrm{L}}$ be an enriched automaton. Then, $\mathrm{L} \subseteq \overline{\mathrm{L}}^{\Gamma} \subseteq \langle \widehat{\Gamma}_{\mathrm{L}} \rangle$, and therefore

$$\langle \widehat{\Gamma}_{\overline{\mathrm{L}}^{\Gamma}} \rangle = \langle \widehat{\Gamma}_{\mathrm{L}} \rangle .$$

That is, $\widehat{\Gamma}_{\mathrm{L}}$ and $\widehat{\Gamma}_{\overline{\mathrm{L}}^{\Gamma}}$ are enriched automata recognizing the same subgroup of $\mathbb{G}_{\mathbf{A}}$. (Note however, that $\overline{\mathrm{L}}^{\Gamma}$ is still not necessarily equal to the full base subgroup $H \cap \mathbb{Z}^m$, since other abelian contributions can arise in H through composition within $\widehat{\Gamma}_{\mathrm{L}}$, see Remark 5.8.7.)

In order to fix the first type of basepoint subgroup ambiguity, we only need to consider the transformation defined below.

Definition 5.8.4. We call *basegroup closure* (for a given action $\mathbf{A}_\bullet$) the transformation consisting in replacing the subgroup basepoint in an enriched automaton by its conjugate closure.

$$(\text{A0}) \quad \widehat{\Gamma}_{L} \rightharpoonup \widehat{\Gamma}_{\overline{L^{\Gamma}}}$$

Fig. 5.16: (Abelian) basegroup closure transformation

Below, we prove that this transformation can be done algorithmically when the enriched automata is finite.

Lemma 5.8.5. *If the automata $\widehat{\Gamma}_{L}$ is finite, then, a basis for the conjugate closure $\overline{L^{\Gamma}}$ is computable.*

Proof. Firstly, recall that we can always use standard linear algebra (over $\mathbb{Z}$) to compute a basis for a subgroup of $\mathbb{Z}^m$, given any *finite* family of generators.

Now, in order to compute a basis for $\overline{L^{\Gamma}}$, start computing a basis B for L, and a basis W for $H\pi_{\mathbb{F}} = \langle \Gamma \rangle$. Then, repeat the following procedure: Check whether $L = \langle B \rangle$ is invariant under the action by conjugation of W; i.e., check whether

$$(B)\mathbf{A}_{W^{\pm 1}} \subseteq \langle B \rangle \,. \tag{5.48}$$

(Note that, since both B and W are finite, this is again linear algebra.) Now, if inclusion (5.48) holds, B is already a basis for $\overline{L^{\Gamma}}$, and we are done. Otherwise, compute a new basis B' for the subgroup generated by $B \cup \mathbf{A}_{W^{\pm 1}}$, update the basis B to B', and restart the process checking (5.48) again for the new B. Repeat until (5.48) is hold, then return the corresponding basis, and stop.

Algorithm 5.8.1: Algorithm to compute the starting base subgroup $\overline{L^{\Gamma}}$

1 START
2 compute a basis B for L;
3 compute a basis W for $\langle \Gamma \rangle$;
4 WHILE $(B)\mathbf{A}_{W^{\pm 1}} \subsetneq \langle B \rangle$ DO
5   compute a basis B' for $\langle\, B \cup (B)\mathbf{A}_{W^{\pm 1}} \,\rangle$;
6   update B to B';
7 RETURN B;
8 STOP

We claim that the previous algorithm returns a basis for the closed subgroup $\overline{L^{\Gamma}}$ (on input a finite enriched automaton for $\widehat{\Gamma}_{L}$).

- It is clear that only elements in $\overline{L^{\Gamma}}$ can appear in the output.
- It is also clear that every element in the basegroup $\overline{L^{\Gamma}}$ will belong to the subgroup $\langle\, B \cup (B)\mathbf{A}_{W^{\pm 1}} \,\rangle$, computed in line 5, at some stage of the procedure.

- So, it only remains to prove that Algorithm 5.8.1 finishes on every (finite) input $\widehat{\Gamma}_L$. This is clear because every enlargement of the candidate subgroup (performed in line 5) supposes either:

  E1 increasing the rank of the base subgroup, or

  E2 decreasing the index of the base subgroup (if the rank keeps constant).

  But, in a free-abelian group $\mathbb{Z}^m$ the rank of a subgroup is bounded above by the ambient rank m; and the index of a subgroup is bounded below by 1. So, none of the steps E1 or E2 (and so line 5 in Algorithm 5.8.1) can be performed infinitely many times, and thus the algorithm must end in finite time.

This concludes the proof. □

It is clear that the basepoint group of a full EA $\widehat{\Gamma}_L$ is the largest possible basepoint group among all the possible basepoint groups for $\widehat{\Gamma}$ recognizing the same subgroup.

Remark 5.8.6. A folded EA is closed iff it is full. (there are no possible contributions from parallel X-arcs).

Remark 5.8.7 (Basepoint subgroup looseness of type II). The second kind of basepoint subgroup looseness is created by parallel enriched x-arcs with the same free label. Indeed, consider the following situation:

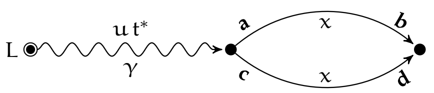

Fig. 5.17: Basepoint subgroup looseness of type II

Note that looping around the parallel x-arc loop in Figure 5.17 (starting at the left vertex), we can read the family of words

$$\left[\left(t^{-\mathbf{a}}\, x\, t^{\mathbf{b}-\mathbf{d}}\, x^{-1} t^{\mathbf{c}}\right)^{\pm 1}\right]^{*},$$

which, naming $\mathbf{e} \coloneqq \mathbf{c}-\mathbf{a}+(\mathbf{b}-\mathbf{d})\mathbf{A}_x^{-1}$, corresponds (in $\mathbb{G}_{\mathbf{A}}$) to the abelian subgroup $\langle\, t^{\mathbf{e}}\,\rangle$; which, in turn, once moved to the basepoint (through γ), provides the abelian contribution

$$\langle t^{\mathbf{e}\mathbf{A}_u^{-1}}\rangle \leqslant \mathbb{Z}^m,$$

Now, if $\mathbf{e}\mathbf{A}_u^{-1} \notin L$, it is clear that $\widehat{\Gamma}_L$ and $\widehat{\Gamma}_{L+\langle \mathbf{e}\mathbf{A}_u^{-1}\rangle}$ are enriched automata with distinct subgroup basepoint recognizing the same subgroup.

We will fix this kind of looseness at the same time that we fix nondeterminism in the skeleton automata Γ, see transformations of type (FII).

In order to specify and fix the looseness that can appear in the arc-labelling of an enriched automaton $\widehat{\Gamma}_L$, we introduce the (abelian) transformations below, which clearly do not affect the recognized subgroup.

Definition 5.8.8. Let us consider the following elementary abelian transformations on enriched automata:

(AI) *Basepoint transformations*: consisting in replacing any abelian label $\mathbf{a}$ in the immediate abelian neighborhood of the basepoint $\odot$, by the label $\mathbf{a}+\mathbf{l}$, for any $\mathbf{l} \in L$ (the basepoint subgroup in $\widehat{\Gamma}_L$).

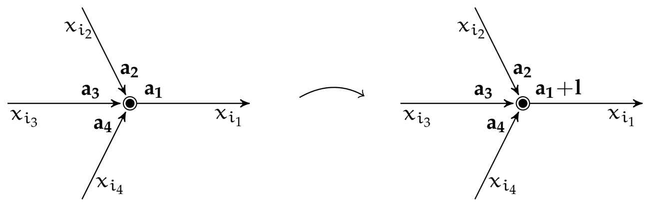

Fig. 5.18: Basepoint transformation

(AII) *Arc transformations*: consisting in respectively adding to the initial and final (abelian) labels of an edge (with free label, say x_i), any element $\mathbf{c} \in \mathbb{Z}^m$, and its corresponding image $\mathbf{c}\mathbf{A}_i$.

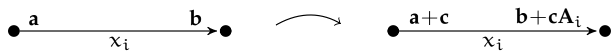

Fig. 5.19: Arc transformation

(AIII) *Non-basepoint transformations*: consisting in adding any element $\mathbf{c} \in \mathbb{Z}^m$ to every abelian label in the abelian neighborhood of a non-base vertex. □

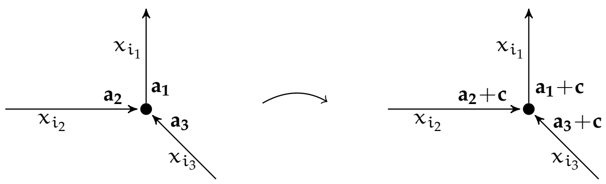

Fig. 5.20: Non-basepoint transformation

The following result is straightforward to check, and we leave the details to the reader.

Lemma 5.8.9. *Elementary abelian transformations do not affect the recognized subgroup in $\mathbb{G}_{\mathbf{A}}$. That is, if $\widehat{\Gamma}_L$ is an enriched $\mathbb{Z}$-automaton, and $\widehat{\Gamma}_L \curvearrowright \widehat{\Gamma}'_{L'}$ is any of the elementary abelian transformations introduced in Definition 5.8.8; then, the subgroups recognized by $\widehat{\Gamma}_L$ and $\widehat{\Gamma}'_{L'}$ in $\mathbb{G}_{\mathbf{A}}$ do coincide; i.e., $\langle \widehat{\Gamma}_L \rangle = \langle \widehat{\Gamma}'_{L'} \rangle \leqslant \mathbb{G}_{\mathbf{A}}$.* □

Remark 5.8.10. We can combine transformations of types AII and AIII to move and spread abelian labels throughout the automata without changing the recognized subgroup. A particularly important application od this fact is used in Lemma 5.8.17

Remark 5.8.11. An AIII-like transformation performed on the basepoint would correspond to a (in general nontrivial) conjugation by an element of the form $t^{\mathbf{c}}$, of the recognized subgroup.

Example 5.8.12. In the particular case of direct products $\mathbb{F}_n \times \mathbb{Z}^m$ (i.e., when $\mathbf{A}_i = \mathrm{id}$, for all i), the previous transformations take a particularly simple form. Namely, abelian arcs work modulo L, and we can freely add any abelian element, either to both extremes of any arc, or to the immediate abelian neighborhood of any vertex (Note that in this case abelian and free labels commute, so there is no point in excluding basevertex neighborhood transformations in AIII since the corresponding conjugations are trivial).

Lemma 5.8.13. *The abelian labels in any bridge-arc in an enriched automaton do not affect the subgroup it recognizes in $\mathbb{F}_n \times \mathbb{Z}^m$. In particular, we can remove all abelian labels from any bridge-arc in this situation.*

Proof 1. Since any ◉-walk would cross any bridge-arc the same number of times in both directions, the neat abelian contribution of any bridge-arc will be zero regardless the actual value of the abelian labels in it. □

Proof 2. Alternatively, use AII to agglutinate all the abelian contribution of the bridge-arc, say $\mathbf{c} \in \mathbb{Z}^m$ in the arc extreme closer to the basepoint; then (using transformations of type AIII) subtract $\mathbf{c}$ of every vertex neighborhood in the component containing the basepoint; and finally (using transformations of type AII) add $\mathbf{c}$ to both the head and tail of every arc, again in the basepoint connected component. □

Of course, if an enriched automaton $\widehat{\boldsymbol{\Delta}}_M$ is the result of applying a finite sequence of elementary abelian transformations to a given enriched automaton $\widehat{\boldsymbol{\Gamma}}_L$ then they recognize the same subgroup of $\mathbb{G}_{\mathbf{A}}$.

Clearly, the converse is not true in general. Think, for example, in what happens with the corresponding skeletons — which are enriched automata as well — in the free group case. So, we need to adapt the classic Stallings foldings to free-abelian by free groups $\mathbb{G}_{\mathbf{A}}$.

The transformations introduced below are aimed to this end: open foldings work essentially in the same way as its free counterparts. However, the extension of closed foldings to $\mathbb{G}_{\mathbf{A}}$ is not that neat. Recall that, in our new ambient, parallel arcs with the same free label can provide new abelian contributions, which we also need to take into account (see Remark 5.8.7).

Definition 5.8.14. Let us consider the following elementary folding transformations on enriched automata:

(FI) *Open (enriched) foldings*: consisting in identifying a pair of nonparallel enriched arcs with exactly the same (free and abelian) labelling is called an *enriched open folding*.

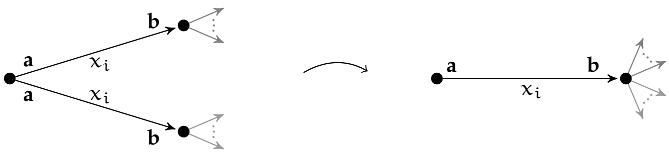

Fig. 5.21: Open enriched folding

(FII) *Closed (enriched) foldings*: consisting in identifying a pair of parallel enriched arcs with the same free label[1] at the price of conveniently updating the basepoint (according with Remark 5.8.7) are called *closed enriched foldings*.

([1]Note however that, without loss of generality, we can restrict these transformations only to parallel x-arcs having null initial labels, previously using abelian arc transformations to cancel the initial labels out.)

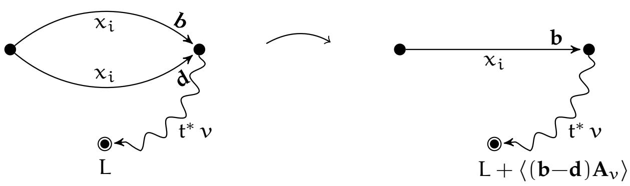

Fig. 5.22: Closed enriched folding

Remark 5.8.15. Note that this last kind of transformation depends on some choices. Namely,

(a) the chosen abelian label in the head of the transformed automaton could have been, for example, $\mathbf{d}$ instead of $\mathbf{b}$, with the corresponding change in the resulting basegroup, and

(b) the chosen walk back to basepoint could have been some other, with the corresponding change in the action, and hence in the basegroup.

However this choices are only relevant locally: since the final basepoint of te folding process will be closed (see Theorem 5.9.12), no matter which choices we make in (a) and (b), the final basepoint in the folded automata will be the same.

Again, the following lemma is natural and straightforward to check.

Lemma 5.8.16. *Enriched foldings do not affect the recognized subgroup in $\mathbb{G}_{\mathbf{A}}$. That is, if $\widehat{\Gamma}_L$ is an enriched $\mathbb{Z}$-automaton, and $\widehat{\Gamma}_L \curvearrowright \widehat{\Gamma}'_{L'}$ is any of the transformations introduced in Definition 5.8.14; then, the subgroups recognized by $\widehat{\Gamma}_L$ and $\widehat{\Gamma}'_{L'}$ in $\mathbb{G}_{\mathbf{A}}$ do coincide; i.e., $\langle \widehat{\Gamma}_L \rangle = \langle \widehat{\Gamma}'_{L'} \rangle \leqslant \mathbb{G}_{\mathbf{A}}$.* □

So, foldings in enriched automata can be performed essentially in the same way as standard foldings in the free group case, but only after moving adequately the abelian labels, and taking into account the abelian contributions that can arise from closed foldings.

Since abelian transformations do not alter the recognized subgroup, it makes sense to look for possible foldings modulo abelian transformations. In this vein, we abuse language and say that two arcs in an enriched automaton $\widehat{\Gamma}_L$ admit an open (resp., closed) folding if they admit it after a suitable abelian transformation on $\widehat{\Gamma}_L$.

The following key lemma shows that the aforementioned parallelism between enriched and standard foldings goes indeed much further.

Lemma 5.8.17. *A pair of arcs $\widehat{e}_1, \widehat{e}_2$ in an enriched automaton $\widehat{\Gamma}_L$ admit an open (resp. closed) folding if and only if the corresponding arcs e_1, e_2 admit an open (resp. closed) folding in the X-skeleton Γ of $\widehat{\Gamma}_L$.*

Proof. The "only if" implication is obvious. For the converse, we distinguish between the two kinds of enriched folding: the implication is clear for closed foldings, since they apply to arbitrary abelian labels (see Figure 5.22). On the other hand, given a general pair of arcs in $\widehat{\Gamma}_L$ admitting an open folding in the skeleton Γ, consider the following sequence of abelian transformations:

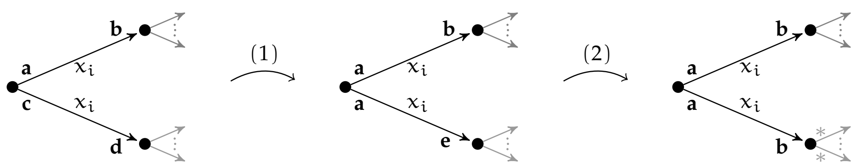

Fig. 5.23: Converting skeleton open folding into enriched open folding

where,

(1) is the abelian transformation (of type AII) consisting in adding $\mathbf{a} - \mathbf{c}$ and $(\mathbf{a} - \mathbf{c})\mathbf{A}_i$ to the abelian labels in the tail and head respectively of the bottom arc in the first automata (note that then, $\mathbf{e} = \mathbf{d} + (\mathbf{a} - \mathbf{c})\mathbf{A}_i$).

(2) is the abelian transformation (of type AIII) consisting in adding $\mathbf{b} - \mathbf{e}$ to the immediate abelian neighborhood of the final vertex in the bottom arc of the second automata in Figure 5.23.

(Recall that the asterisks in the last automata indicate possibly changed abelian labels.)

It is evident that, the transformed arcs in the final automaton (obtained using only abelian transformations) admit an open enriched folding.

Note that, since necessarily (al least) one of the folded vertices is not the basepoint, we can always perform a legal transformation of type AIII in (2) using a folded arc without the basepoint. This concludes the proof. □

5.9 Inverse (enriched) automata

The previous lemma has several important implications for us. One immediate consequence is that the property of an enriched automaton being folded is indeed a property of its skeleton.

Corollary 5.9.1. *Let $\widehat{\Gamma}_L$ be an enriched automaton, and Γ its skeleton. Then:*

$$\widehat{\Gamma}_L \text{ is folded } \iff \Gamma \text{ is folded}\,. \qquad \square$$

In the same vein, Lemma 5.8.17 allows us to adapt some definitions and procedures from the free case. For example, now it is clear that the notions of determinism and trimness (for enriched automata) are indeed properties of the skeleton.

Definition 5.9.2. An enriched automata $\widehat{\Gamma}_L$ is said to be *deterministic* (resp., *trim*) if its skeleton Γ is so. We define the trim, and total trim transformations in a completely analogous way as in the free case.

Moreover, restricting our subgroup representatives only to deterministic and trim enriched automata, not only fixes the third kind of looseness in Remark 5.7.14, but also fixes subgroup looseness of type II, since no parallel arcs with the same free label can appear in a folded automata (see Remark 5.8.7). Therefore, a closed base subgroup of a folded enriched automata recognizing $H \leqslant \mathbb{G}_{\mathbf{A}}$ must necessarily be the full base subgroup $(H \cap \mathbb{Z}^m)$ of H.

Lemma 5.9.3. *Let $\widehat{\Gamma}_L$ be a folded enriched automaton recognizing $H \leqslant \mathbb{G}_{\mathbf{A}}$. Then, $\overline{L^{\Gamma}} = H \cap \mathbb{Z}^m$. That is, a closed folded automata is necessarily full.*

Proof. The inclusion $\overline{L^{\Gamma}} \leqslant H \cap \mathbb{Z}^m$ is trivial. For the converse inclusion, suppose that $t^{\mathbf{a}} \in H \cap \mathbb{Z}^m$. This means that there exists a $\odot$-closed enriched walk $\widehat{\gamma}$ in $\widehat{\Gamma}_L$

that recognizes the element $1_{\mathbb{F}_X} t^a \in \mathbb{G}_\mathbf{A}$. Therefore γ must recognize the trivial element $1_{\mathbb{F}_X}$ in the skeleton Γ; but, since Γ is reduced by hypothesis, γ only can read the trivial element if it is a sequence of cancellations, i.e.,

$$\gamma = \gamma_1^{-1}\gamma_1\gamma_2^{-1}\gamma_2\cdots\gamma_k^{-1}\gamma_k\,,$$

where, for each $i = 1,\ldots,k$, γ_i is an elementary $\odot$-walk in the skeleton Γ. Hence, according Definition 5.7.9, the enriched walk $\widehat{\gamma}$ recognizes the set

$$\begin{aligned}
\ell(\widehat{\gamma}) &= t^L\,\ell(\widehat{\gamma}_1^{-1})\,t^L\,\ell(\widehat{\gamma}_1)\,t^L\,\ell(\widehat{\gamma}_2)^{-1}\,t^L\,\ell(\widehat{\gamma}_2)\,\cdots\,t^L\,\ell(\widehat{\gamma}_k)^{-1}\,t^L\,\ell(\widehat{\gamma}_k)\,t^L\\
&= t^L\,\ell(\gamma_1)^{-1}\,t^L\,\ell(\gamma_1)\,t^L\,\ell(\gamma_2)^{-1}\,t^L\,\ell(\gamma_2)\,\cdots\,t^L\,\ell(\gamma_k)^{-1}\,t^L\,\ell(\gamma_k)\,t^L\\
&= t^L\,t^{(L)\mathbf{A}_{\ell(\gamma_1)}}\,t^L\,t^{(L)\mathbf{A}_{\ell(\gamma_2)}}\,\ldots\,t^L\,t^{(L)\mathbf{A}_{\ell(\gamma_k)}}\,t^L\\
&= t^{L+(L)\mathbf{A}_{\ell(\gamma_1)}+L+(L)\mathbf{A}_{\ell(\gamma_2)}+\cdots+L+(L)\mathbf{A}_{\ell(\gamma_k)}+L} \leqslant \overline{L^\Gamma}\,.
\end{aligned}$$

Therefore $t^a \in \widehat{\gamma}$ must belong to the closure $\overline{L^\Gamma}$, and the proof is completed. □

Since nontrimness can be fixed in exactly the same way as in the free group case, we have reached a quite compact family of representatives of a given subgroup $H \leqslant \mathbb{G}_\mathbf{A}$; namely those being deterministic, closed, and trim.

Definition 5.9.4. An enriched automaton is said to be *inverse* if it is deterministic, basepoint closed, and trim.

Inverse (enriched) automata, although still not unique, are quite neat descriptions of subgroups of free-abelian by free groups. Namely, the factor decomposition of a given subgroup is explicitly encoded in them.

Proposition 5.9.5. *Let* $\widehat{\Gamma}_L$ *be an inverse automaton recognizing the subgroup* $H \leqslant \mathbb{G}_\mathbf{A}$ *with decomposition*

$$H = H\pi_\mathbb{F}\sigma \ltimes (H \cap \mathbb{Z}^m) \simeq H\pi_\mathbb{F} \ltimes (H \cap \mathbb{Z}^m)\,,$$

where σ *is a section as in* (5.38). *Then,*

(i) *The base subgroup* L *is full; i.e.,* $L = H \cap \mathbb{Z}^m$.

(ii) *The automata* $\widehat{\Gamma}$ *recognizes precisely the subgroup* $H\pi_\mathbb{F}/(H\cap\mathbb{Z}^m) \simeq H\pi_\mathbb{F}\sigma$.

(iii) *The skeleton* Γ *is the Stallings automaton of the projection* $H\pi_\mathbb{F} \simeq H\pi_\mathbb{F}\sigma$. □

Below, we see that for every involutive enriched automaton, there exist inverse automata equivalent to it. Note that a constructive proof of this fact (see Theorem 5.9.12) can not be given in the general setting.

Definition 5.9.6. Let $\widehat{\Gamma}_L$ be an involutive $\mathbb{Z}^m$-enriched X-automaton. We say that two vertices p, q in $\widehat{\Gamma}_L$ are *(graphically) equivalent*, denoted by $p \equiv q$, if they are equivalent in the skeleton Γ (i.e., if there exist a walk in the skeleton Γ from one vertex to the other reading the trivial element in $\mathbb{F}_X$).

Lemma 5.9.7. *Equivalence of vertices is an equivalence relation compatible with the enriched automaton structure; i.e., the map*

$$\begin{array}{rcl} V\Gamma & \twoheadrightarrow & V\Gamma/\equiv \\ p & \mapsto & [p] \end{array}$$

induces a well-defined epimorphism of enriched automata $\widehat{\Gamma}_L \twoheadrightarrow \widehat{\Gamma}_L/\equiv$, where the quotient automaton $\widehat{\Gamma}_L/\equiv$ is defined as follows:

1. *its vertex set is $V\Gamma/\equiv$ (the quotient of V modulo graphic equivalence);*
2. *its arcs are those in $\widehat{\Gamma}_L$ under the identification (i), with parallel arcs (both the original ones in $\widehat{\Gamma}_L$, and those created by the identification of vertices) identified using closed (enriched) foldings;*
3. *its basepoint is the closure of the result of adding to L the abelian contributions coming from closed foldings. (Note that some of these closed foldings will lay in the degenerated paths in $\widehat{\Gamma}_L/\equiv$ corresponding to paths between identified vertices in $\widehat{\Gamma}_L$ reading purely abelian elements.)*

(Note that, as it always happens in enriched automata, the quotient automaton $\widehat{\Gamma}_L/\equiv$ should be thought modulo the base subgroup.) □

Proposition 5.9.8. *Let $\widehat{\Gamma}_L$ be a $\mathbb{Z}^m$-enriched X-automaton. Then, the core of (any representative of) the quotient automaton $\widehat{\Gamma}_L/\equiv$ is an inverse $\mathbb{Z}^m$-enriched X-automaton recognizing the same subgroup as $\widehat{\Gamma}_L$.*

These (abelianly equivalent) automata are called *Stallings reductions* of $\widehat{\Gamma}_L$, and are denoted by $\widehat{\mathrm{St}}(\widehat{\Gamma}_L)$ when thought modulo abelian transformations, or some particular abelian labelling is distinguished.

Proof. We want to prove that $\langle \widehat{\Gamma}_L/\equiv \rangle = \langle \widehat{\Gamma}_L \rangle$. But this is clear by construction of the quotient automaton in Lemma 5.9.7: the quotient image of any $\odot$-walk $\widehat{\gamma}$ in $\widehat{\Gamma}_L$ is a $\odot$-walk in $\widehat{\Gamma}_L/\equiv$ reading exactly the same element in $\mathbb{G}_{\mathbf{A}}$; and, given a $\odot$-closed walk $[\widehat{\gamma}]$ in the quotient, either it corresponds exactly to a $\odot$-path in $\widehat{\Gamma}_L$, or the difference amounts to an abelian contribution by either a closed folding, or a vertex identification (both included in the basepoint of the quotient). □

Moreover, since Stallings reductions are inverse, we can associate them biunivocally (modulo graphic equivalence) with subgroups of $\mathbb{G}_{\mathbf{A}}$.

Corollary 5.9.9. *Two enriched automata $\widehat{\Gamma}_L, \widehat{\Delta}_M$ recognize the same subgroup of $\mathbb{G}_{\mathbf{A}}$ (i.e., they are equivalent) if and only if their quotient automata coincide (i.e., they are graphically equivalent).* □

Definition 5.9.10. Let H be a subgroup of $\mathbb{G}_{\mathbf{A}}$. Then, a *Stallings automaton* for H with respect to X, is any Stallings reduction of any $\mathbb{Z}^m$-enriched X-automata $\widehat{\Gamma}_L$ recognizing H. We will denote them by

$$\widehat{\mathsf{St}}(H, X) \coloneqq \operatorname{core}(\widehat{\mathsf{St}}(\widehat{\Gamma}_L)), \tag{5.49}$$

where, as before, (5.49) is thought modulo abelian transformations, or assuming some distinguished abelian labelling.

Remark 5.9.11. We have tightened the description in (5.46) of subgroups as enriched automata to the following (now bijective) map:

$$\begin{array}{rcl} \{\text{inverse } \mathbb{Z}^m\text{-enriched X-automata}\} / \equiv & \to & \{\text{subgroups of } \mathbb{G}_{\mathbf{A}}\} \\ \widehat{\Gamma}_L & \mapsto & \langle \widehat{\Gamma}_L \rangle_{\mathbb{G}_{\mathbf{A}}} \\ \widehat{\mathsf{St}}(H, X) & \leftarrow\!\shortmid & H\,. \end{array} \tag{5.50}$$

Note that the bijection (5.50) is in general not algorithmic, since we are not assuming the subgroups, or even the ambient group $\mathbb{G}_{\mathbf{A}}$ to be finitely generated. However, as it happened for free groups, we will see that graphical equivalence admits an effective description in terms of (enriched) foldings in the finitely generated case, which is an straightforward adaptation of the procedure given in Proposition 5.4.41.

Theorem 5.9.12. *Any finite enriched automaton $\widehat{\Gamma}_L$ recognizing a subgroup $H \leqslant \mathbb{G}_{\mathbf{A}}$ can be converted into a reduced automaton — recognizing the same subgroup H — after a finite number of (folding, abelian, or trim) transformations.*

Proof. By Proposition 5.4.41, we know that there exists a finite sequence of transformations

$$\Gamma \xrightarrow{\varphi^{(1)}} \Gamma^{(1)} \xrightarrow{\varphi^{(2)}} \cdots \xrightarrow{\varphi^{(p)}} \Gamma^{(p)} \xrightarrow{\tau} \Gamma^{(p+1)} = \mathsf{St}(\Gamma) \tag{5.51}$$

(where the $\varphi^{(i)}$s denote folding transformations, and τ denotes a total trim) converting the skeleton of $\widehat{\Gamma}_L$ into an inverse automaton.

Now it is just a matter of translating these sequence of standard transformations into a finite sequence of enriched (folding, abelian, or trim) transformations from $\widehat{\Gamma}_L$ to a reduced automaton. But we know from Lemma 5.8.17 that we can (using abelian transformations in the case of open foldings, and non basepoint transformations in the case of closed foldings) obtain respective enriched foldings

$\widehat{\varphi}^{(i)}$, for each $\varphi^{(p)}$ in the folding sequence (5.52). This immediately provides the desired finite sequence of (abelian, basepoint, and enriched folding) transformations converting $\widehat{\Gamma}_L$ into an equivalent deterministic enriched automata, say $\Gamma_{L_p}^{(p)}$.

Note that $\Gamma_{L_p}^{(p)}$ is still not necessarily basegroup closed nor trim. We finish our sequence of transformations fixing these two anomalies with the corresponding trim (τ), and basegroup closure (β) transformations, to finally obtain a reduced automaton for $\widehat{\Gamma}_L$ after a finite number of effective transformations.

$$\widehat{\Gamma}_L \overset{\widehat{\varphi}^{(1)}}{\rightsquigarrow} \Gamma_{L_1}^{(1)} \overset{\widehat{\varphi}^{(2)}}{\rightsquigarrow} \cdots \overset{\widehat{\varphi}^{(p)}}{\rightsquigarrow} \Gamma_{L_p}^{(p)} \overset{\tau}{\rightsquigarrow} \Gamma_{L_p}^{(p+1)} \overset{\beta}{\rightsquigarrow} \Gamma_{L_{p+1}}^{(p+1)} = \widehat{\mathsf{St}}(\widehat{\Gamma}_L) \tag{5.52}$$

Fig. 5.24: Computation of a reduced (enriched) automata

This concludes the proof. □

Corollary 5.9.13. *A reduced automaton recognizing a finitely generated subgroup* $H \leqslant \mathbb{G}_{\mathbf{A}}$ *is computable given any finite generating set for* H. □

For some practical issues, having a reduced automaton of a given finitely generated subgroup $H \leqslant \mathbb{G}_{\mathbf{A}}$ is enough to perform the desired computations and obtain results (e.g. to compute the enriched pullback).

Example 5.9.14. Let $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{Z}^m$, $H \leqslant \mathbb{Z}^m$, and consider the (parameterized) subgroup

$$H = \left\langle t^{\mathbf{L}}, x^3\, t^{\mathbf{a}}, yx\, t^{\mathbf{b}}, y^3xy^{-2}\, t^{\mathbf{c}} \right\rangle \leqslant \mathbb{F}_{\{x,y\}} \times \mathbb{Z}^m .$$

Then, the flower automaton of H is

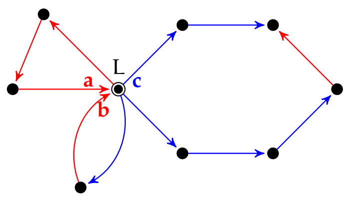

Now, after conveniently moving the label $\mathbf{c}$ (using AII and AIII transformations), we can successively fold (using FI transformations): first all three blue arcs incident to de basepoint, and then the two blue arcs departing from the end of the last folded arc; to get

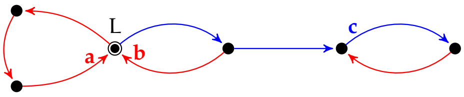

Finally, we can rearrange $\mathbf{a}$ as $\mathbf{a} + (\mathbf{b} - \mathbf{a})$ over the red triangle in order match the labelling of the two red arcs incident to the basepoint

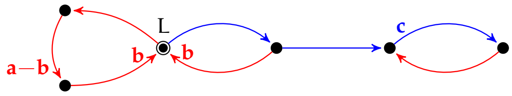

and finally, identify the two red arcs incident to the basepoint, to obtain the following reduced automaton for H:

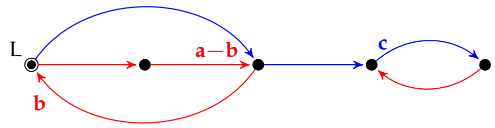

Fig. 5.25: Reduced automaton for $H = \langle t^{\mathbf{L}}, x^3\, t^{\mathbf{a}}, yx\, t^{\mathbf{b}}, y^3xy^{-2}\, t^{\mathbf{c}}\rangle$

Note, however, that this is not the only reduced automaton recognizing H. In order to have a neat bijection between subgroups and "some kind of automata" we still have to fix abelian equivalence (among reduced automata recognizing the same subgroup).

One possible shortcut would be simply looking at reduced automata modulo abelian transformations, but this coarse solution is not necessary because a much more tangible object can be considered, which, besides, provides some additional insight into the whole picture.

The drawback is that, in order to choose a unique reduced representative of a subgroup H, we will need to make some arbitrary choice. Namely, we need to arbitrarily choose a spanning tree in the Stallings automaton of $H\pi_{\mathbb{F}}$.

(Recall that the considerations in Theorem 5.3.10 must be taken into account whenever we consider existence and computability of spanning trees in graphs of arbitrary order.)

Definition 5.9.15. A reduced automaton $\widehat{\Gamma}_L$ is said to be *normalized* with respect to an spanning tree T (T-normalized, for short) if:

1. the abelian labelling of every arc in T is zero.
2. the initial (abelian) label of every arc outside T is zero.

Lemma 5.9.16. *Let* H *be a subgroup of* $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes \mathbb{Z}^m$. *Then, for every spanning tree* T *of the Stallings automaton of* $H\pi_{\mathbb{F}} \leqslant \mathbb{F}_X$, *there exists a normalized reduced automaton for* H *relative to* T. *Moreover, if the subgroup is given by a finite family of generators, then a normalized reduced automaton for* H *is computable.*

It is called an (enriched) Stallings automaton *of* H *with respect to* T, *and denoted by* $\mathsf{St}_{\mathsf{T}}(H)$. □

Proof. The existential claim is a consequence of the previous discussion.

For the algorithmic claim, suppose that H is a subgroup given by a finite family of generators, and therefore a reduced automaton $\widehat{\Gamma}_L$ recognizing H is computable. Let $\mathbf{T}$ be an spanning tree of $\widehat{\Gamma}_L$; we claim that we can normalize $\widehat{\Gamma}_L$ just pushing out the abelian labels, starting from the basepoint and through the spanning tree.

Indeed, suppose that we have removed all abelian labels from every arc at $\mathbf{T}$-distance at most k from the basepoint (i.e., from a vertex p_k at $\mathbf{T}$-distance k to a vertex p_{k+1} at $\mathbf{T}$-distance $k+1$ to the basepoint). Then, it is enough to, for every arc at $\mathbf{T}$-distance $k+1$ from the basepoint,

$$\widehat{e} \equiv p_{k+1} \xrightarrow[x]{\mathbf{a} \qquad \mathbf{b}} p_{k+2}$$

perform an edge transformation removing the label $\mathbf{a}$ from the tail of e, and then perform a non-basepoint transformation in p_{k+2} to remove the obtained label (concretely $\mathbf{b} - \mathbf{a}\mathbf{A}_x$) from the head of e. Note that this sequence of transformations can have only two possible outcomes:

(a) it modifies the abelian labelling of an arc at distance $k+1$ from the basepoint, or

(b) it modifies the abelian labelling of a cyclotomic arc (outside $\mathbf{T}$).

But since the automata $\widehat{\Gamma}_L$ is finite, after a finite number of steps there will be no arcs at distance $k+1$, and every non-zero abelian label will lie in a cyclotomic arc.

Finally, we can apply arc transformations to concentrate the abelian weight in the heads of cyclotomic arcs to get the $\mathbf{T}$-normalization of $\widehat{\Gamma}_L$, which is therefore computable. □

Once assumed the existence of reduced automata (and so AC), the existence and unicity of Stallings automata (modulo the basegroup) is immediate. Moreover, the nonzero abelian label in any arc outside the chosen tree is easily computable from the corresponding ◉-walk in any reduced automaton for H.

However, in order to properly define a bijection between subgroups and Stallings automata, we need an *a priori* well-defined uniform way of choosing the spanning tree for each reduced automaton $\widehat{\Gamma}_L$. This can be done using a well-order fixed a priori in the set X of free generators in $\widehat{\Gamma}_L$. (Note that a well-order in X automatically induces a well-order in the edge-neighborhood of any vertex in a deterministic X-automaton.)

We remark that the existence of well-orders for general sets is again equivalent to the axiom of choice (see [Jec73]),which we are assuming for the general setting.

Lemma 5.9.17. *Let Γ be an arbitrary inverse* X*-automaton, let* $\prec$ *be a well-order in* X*, and consider the family of trees* $\mathbf{T}_k$ *recursively defined by:*

- $\mathbf{T}_1 = \odot$
- $\mathbf{T}_{k+1}$ *is obtained by attaching to* $\mathbf{T}_k$ *every vertex* p *at distance* 1 *from it, using the arc from a vertex in* $\mathbf{T}_k$ *to* p *with* $\prec$*-minimum possible* X*-label. Formally:*

$$\begin{cases} \mathsf{V}\mathbf{T}_{k+1} = \mathsf{V}\mathbf{T}_k \cup \{\, p \in V : \; d(p, \mathbf{T}_k) = 1 \,\}\,, \\ \mathsf{E}\mathbf{T}_{k+1} = \mathsf{E}\mathbf{T}_k \cup \{ \min \mathsf{E}(p \leftarrow \mathsf{V}\mathbf{T}_k) \; : \; d(p, \mathbf{T}_k) = 1 \,\}\,. \end{cases}$$

Then, the union

$$\bigcup\nolimits_{k\geqslant 1} \mathbf{T}_k$$

is a spanning tree of Γ*. We call it the* radial spanning tree *of* Γ *induced by* $\prec$. □

If the generating set is finite, say $X = \{x_1, \ldots, x_p\}$, then the radial spanning tree of an Stallings automaton is always constructible, and will be considered by default. Also, if not state otherwise, we will assume the default ordering $x_1 \prec x_1^{-1} \prec \cdots \prec x_p \prec x_p^{-1}$.

Example 5.9.18. The Stallings automaton of the subgroup in Figure 5.25 (i.e., relative to the radial spanning tree induced by the ordering $a \prec a^{-1} \prec b \prec b^{-1}$) is

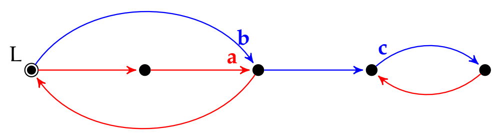

Fig. 5.26: Stallings automaton of $H = \langle t^{\mathbf{L}}, a^3\, t^{\mathbf{a}}, bat^{\mathbf{b}}, b^3ab^{-2}\, t^{\mathbf{c}} \rangle$

It is clear that the radial spanning tree only is algorithmically constructible when the label set X is finite.

Remark 5.9.19. Note, that if X is finite both the layer and the diagonal spanning trees are computable, but they are not necessarily the same. Thus, one of the construction schemes must be fixed a priori.

Therefore, a uniform choice of spanning trees for reduced automata can be done, and we can finally state our characterization theorem for subgroups of free-abelian by free groups.

Theorem 5.9.20. *Let* $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes_{\mathbf{A}_\bullet} \mathbb{Z}^m$ *be a free-abelian by free group as in* (5.32)*, and assume a particular well-order* $\prec$ *in* X*. Then, the map*

$$\begin{array}{rcl} \{\, \textit{Subgroups of}\ \mathbb{G}_{\mathbf{A}} \,\} & \to & \{\, \textit{Stallings}_{\prec}\ (X, \mathbb{Z}^m)\textit{-automata} \,\} \\ H & \mapsto & \mathrm{St}_{\prec}(H) \end{array} \tag{5.53}$$

(sending every subgroup $H \leqslant \mathbb{G}_{\mathbf{A}}$ to its Stallings automaton with respect to the spanning tree generated by $\prec$) is a bijection with inverse $\widehat{\Gamma}_L \mapsto \langle \widehat{\Gamma}_L \rangle_{\mathbb{G}_{\mathbf{A}}}$.

Moreover, if $\mathsf{St}_{\prec}(H) = \widehat{\Gamma}_L$, then the set

$$B_{\prec} = \left\{ u_{\prec}[\widehat{e}] : \widehat{e} \in E^{+}\left(\widehat{\Gamma} \setminus \mathbf{T}_{\prec}\right) \right\}, \tag{5.54}$$

is a free basis for the subgroup $H\pi_{\mathbb{F}}$ recognized by $\widehat{\Gamma}$. □

Corollary 5.9.21. *The restriction of the bijection (5.53) to finitely generated subgroups is computable. In particular, if a subgroup $H \leqslant \mathbb{G}_{\mathbf{A}}$ is given by a finite family of generators, then a Stallings automaton recognizing H is computable.* □

5.10 Algorithmic problems for free-abelian by free groups

We will start reviewing the status of the classic Dehn problems for the family of free-abelian by free groups (of the form $\mathbb{F}_n \ltimes \mathbb{Z}^m$). These results (only the word problem being solvable) makes it apparent the computational complicacy of this family.

Theorem 5.10.1.

(i) *The word problem is solvable for free-abelian by free groups.*

(ii) *The conjugacy problem is unsolvable for free-abelian by free groups* [BMV10].

(iii) *The isomorphism problem is unsolvable for free-abelian by free groups* [Zim85; Lev08]. □

The solvability of the word problem is, as often happens, a direct consequence of the computability of normal forms for the elements of groups in the family (see Remark 5.6.5).

The other two results are much more involved, and closely related. They are both based in the remarkable Theorem 3.1 in [BMV10] which links the conjugacy problem in a group extension with the orbit-decidability of the action subgroup. Since $\mathbb{F}_2$ embeds in $\mathrm{GL}_2(\mathbb{Z})$, and so $\mathbb{F}_2 \times \mathbb{F}_2$ (which is well known to have unsolvable MP) embeds in $\mathrm{GL}_4(\mathbb{Z})$, the authors deduce that $\mathrm{GL}_4(\mathbb{Z})$ contains finitely generated orbit undecidable subgroups. Then, the aforementioned Theorem 3.1 implies the existence of $\mathbb{Z}^4$-by-[f.g. free] groups with unsolvable conjugacy problem (see [BMV10, Corollary 7.6]).

Finally, in the unpublished note [Lev08], Levitt elaborates on the previous ideas to show that the isomorphism problem for $\mathbb{Z}^4$-by-[f.g. free] groups is unsolvable;

and recalls that the previous paper [Zim85] by Zimmermann contains a similar result for $\mathbb{Z}^4$-by-surface groups; whose argument also applies to $\mathbb{Z}^4$-by-[f.g. free] groups.

5.10.1 Membership problem

As in the free case, the first application of (now enriched) automata describing subgroups is that they naturally solve the MP within the class of free-abelian by free groups: namely a word w belongs to the subgroup generated by a finite set of words $w_1, \ldots, w_p$ in the generators if and only if it is readable in the Stallings automaton of the subgroup $\langle w_1, \ldots, w_p \rangle$.

Theorem 5.10.2. *The subgroup membership problem is solvable for free-abelian by free groups.*

Proof. It is clear that we can suppose any input word in normal form. So, given finitely many elements $w\, t^{\mathbf{a}}, w_1\, t^{\mathbf{a}_1}, \ldots, w_k\, t^{\mathbf{a}_k} \in \mathbb{G}_{\mathbf{A}}$, in order to decide whether $w\, t^{\mathbf{a}}$ belongs to the subgroup $H = \langle w_1\, t^{\mathbf{a}_1}, \ldots, w_k\, t^{\mathbf{a}_k} \rangle$, apply the following procedure:

1. Build the Stallings automata $\widehat{\Gamma}_L$ of H.

2. Try to read the free part w of the word in the skeleton Γ, keeping track of the global abelian contribution $\mathbf{c}_w \in \mathbb{Z}^m$ obtained in doing so. If this is not possible, return NO; otherwise continue.

3. If the final vertex (after reading w in Γ) is not the basepoint, then return NO; otherwise continue.

4. Check whether $t^{\mathbf{a}} \in \mathbf{c}_w + L$ (this can be easily done using linear algebra over $\mathbb{Z}$). In affirmative case return YES, otherwise return NO. □

5.10.2 Intersection problem

We have already shown that free-abelian by free groups are not Howson. Recall that even the simple case $\mathbb{Z}^m \times \mathbb{F}_2$ is not Howson (see Figure 5.49).

Therefore, it makes sense to consider the following problems for a finitely presented group G in this family.

Subgroup intersection (decision) problem, $\mathsf{SIP_d}(G)$**.** *Given a finite set of words* $u_1, \ldots, u_n, v_1, \ldots, v_m$ *in the generators of* G*, decide whether the subgroup intersection* $\langle u_1, \ldots, u_n \rangle_G \cap \langle v_1, \ldots, v_m \rangle_G$ *is finitely generated or not.*

Subgroup intersection (search) problem, $\mathsf{SIP_s}(G)$**.** *Given a finite set of words* $u_1, \ldots, u_n$, $v_1, \ldots, v_m$ *in the generators of* G, *compute a generating set for the subgroup intersection* $\langle u_1, \ldots, u_n \rangle_G \cap \langle v_1, \ldots, v_m \rangle_G$.

Subgroup intersection (full) problem, $\mathsf{SIP}(G)$**.** *Given a finite set of words* $u_1, \ldots, u_n$, $v_1, \ldots, v_m$ *in the generators of* G, *decide whether the subgroup intersection* $\langle u_1, \ldots, u_n \rangle_G \cap \langle v_1, \ldots, v_m \rangle_G$ *is finitely generated or not; and* in affirmative case, *compute a generating set for this intersection.*

In order to tackle these problems, we pretend to adapt the ideas in Section 5.5.1 to our enriched setting. Namely, given subgroups $H, K \leqslant \mathbb{G}_{\mathbf{A}}$, we want to determine a reduced (enriched) automaton associated to its intersection, and then derive from it the desired properties about the intersection. The construction below is introduced to this end.

Definition 5.10.3. The *product* of two enriched automata $\widehat{\Gamma}_L$ and $\widehat{\boldsymbol{\Delta}}_M$, denoted by $\widehat{\Gamma}_L \times \widehat{\boldsymbol{\Delta}}_M$, consists of the product $\boldsymbol{\Gamma} \times \boldsymbol{\Delta}$ of their respective skeletons, *doubly-enriched* with the abelian labelling detailed below:

1. To every arc $(e, e') \equiv (p, p') \xrightarrow{x_i} (q, q')$ in the product $\boldsymbol{\Gamma} \times \boldsymbol{\Delta}$ we add:
   1.1. the pair $(\mathbf{a}, \mathbf{a}')$ as initial label of (e, e'), where $\mathbf{a}$ and $\mathbf{a}'$ are the initial labels of the arcs e and e' (in $\widehat{\Gamma}_L$ and $\widehat{\boldsymbol{\Delta}}_M$), respectively; and
   1.2. the pair $(\mathbf{b}, \mathbf{b}')$ as final label of (e, e'), where $\mathbf{b}$ and $\mathbf{b}'$ are the final labels of e and e' (in $\widehat{\Gamma}_L$ and $\widehat{\boldsymbol{\Delta}}_M$), respectively.
2. We add the pair (L, L') as a label for the basepoint of the enriched product.

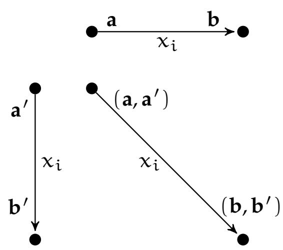

Fig. 5.27: Construction of an arc in the enriched product

Finiteness, connectivity, skeleton, core, etc. are defined in the natural way for doubly-enriched automata. In particular, the skeleton of the product is (by definition) the product of the skeletons of the factors.

Notation 5.10.4. We will denote doubly enriched automata by $\widehat{\widehat{\boldsymbol{\Upsilon}}}_{L,L'}$ where $\boldsymbol{\Upsilon}$ is its skeleton, and the double tilde stands for the double abelian labelling.

Generic double hat notation (for objects with double abelian labelling) is defined in the same way as generic single hat notation (see Notation 5.7.5). That is, if $\hat{\hat{\mathbf{o}}}$ is a doubly-enriched object, we will denote by $\hat{\mathbf{o}}$ (resp., by $\hat{\mathbf{o}}'$) the single-enriched object obtained by considering only the first (resp., second) abelian component in $\hat{\hat{\mathbf{o}}}$. For example, given an arbitrary doubly-enriched automaton $\hat{\hat{\Upsilon}}_{L,L'}$ we will denote by $\hat{\Upsilon}_L$ (resp., $\hat{\Upsilon}'_{L'}$) the enriched automata defined by its first (resp., second) components.

The following key result, clear by construction, is based on the same principles as the homologous Proposition 5.5.27 for free groups.

Proposition 5.10.5. *Let* H, H' *be subgroups of* $\mathbb{G}_{\mathbf{A}}$*, and let* $\hat{\Gamma}_L, \hat{\Gamma}'_{L'}$ *be Stallings automata recognizing* H *and* H' *respectively. Then, an element* $w\,\mathbf{t^a} \in \mathbb{G}_{\mathbf{A}}$ *belongs to the intersection* $H \cap H'$ *if and only if it is componentwise* $\odot$*-readable in the core of the enriched product* $\hat{\Gamma}_L \times \hat{\Gamma}'_{L'}$*. Formally,*

$$H \cap H' = \bigcup \left\{ \hat{w}_{\mathbb{G}_{\mathbf{A}}} \cap \hat{w}'_{\mathbb{G}_{\mathbf{A}}} \,:\, w \in \left\langle \mathsf{St}\left(H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}\right) \right\rangle \right\}, \tag{5.55}$$

where $\hat{w}_{\mathbb{G}_{\mathbf{A}}}$ *and* $\hat{w}'_{\mathbb{G}_{\mathbf{A}}}$ *are the completions of* w *in* H *and* H' *respectively.* □

Corollary 5.10.6. *Note that then,*

$$\begin{aligned}(H \cap H')\pi_{\mathbb{F}} &= \left\{ w \in \left\langle \mathsf{St}\left(H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}\right) \right\rangle \,:\, \hat{w}_{\mathbb{G}_{\mathbf{A}}} \cap \hat{w}'_{\mathbb{G}_{\mathbf{A}}} \neq \varnothing \right\} \\ &= \left\{ w \in \left\langle \mathsf{St}\left(H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}\right) \right\rangle \,:\, \mathcal{C}_L(w) \cap \mathcal{C}_{L'}(w) \neq \varnothing \right\} \qquad \square\end{aligned}$$

Hence, the core of the enriched product $\hat{\mathsf{St}}(H) \times \hat{\mathsf{St}}(H')$ encodes the intersection $H \cap H'$, and in particular whether it is finitely generated or not.

Definition 5.10.7. Let H, H' be subgroups of $\mathbb{G}_{\mathbf{A}}$, and let $\hat{\Gamma}_L, \hat{\Gamma}'_{L'}$ be Stallings automata recognizing H and H' respectively. Then, the core of $\hat{\mathsf{St}}(H) \times \hat{\mathsf{St}}(H')$ is called an *intersection scheme* for $H \cap H'$, and denoted $\hat{\mathsf{St}}(H) \wedge \hat{\mathsf{St}}(H')$; i.e.,

$$\hat{\mathsf{St}}(H) \wedge \hat{\mathsf{St}}(H') \coloneqq \operatorname{core}\left(\hat{\mathsf{St}}(H) \times \hat{\mathsf{St}}(H')\right).$$

We say that an intersection scheme is (**T**)-normalized, if we have distinguished a spanning tree **T** of the scheme, and have imposed the conditions in Definition 5.9.15 to both abelian components in the product automaton. The arcs outside a spanning tree (say **T**) are called (**T**-)*cyclomatic*.

Since for any enriched automata, the skeleton of the core is the core of the skeleton, we have:

$$\begin{aligned} \operatorname{sk}\operatorname{core}(\widehat{\mathsf{St}}(H) \times \widehat{\mathsf{St}}(H')) &= \operatorname{core}\operatorname{sk}(\widehat{\mathsf{St}}(H) \times \widehat{\mathsf{St}}(H')) \\ &= \operatorname{core}(\operatorname{sk}\widehat{\mathsf{St}}(H) \times \operatorname{sk}\widehat{\mathsf{St}}(H')) \\ &= \operatorname{core}(\mathsf{St}\,(H\pi_{\mathbb{F}}) \times (\mathsf{St}\,\big(H'\pi_{\mathbb{F}}\big)) \\ &= \mathsf{St}\,\big(H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}\big)\ . \end{aligned}$$

So, the free part of any element in the intersection $H \cap H'$ is recognized by the standard Stallings automaton $\mathsf{St}\,(H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}})$ (which is the skeleton of any intersection scheme for $H \cap H'$), and we recover the obvious inclusion $(H \cap H')\pi_{\mathbb{F}} \leqslant H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$. However, according Proposition 5.10.5, in order to belong to $(H \cap H')\pi_{\mathbb{F}}$, an element in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ needs to be componentwise readable in the enriched pullback, and this is something that not always happen (see Example 5.5.30 for a concrete case).

Lemma 5.10.8. *Let H, H' be subgroups of $\mathbb{G}_{\mathbf{A}} = \mathbb{F}_X \ltimes \mathbb{Z}^m$. Then,*

$$(H \cap H')\pi_{\mathbb{F}} \leqslant H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}\,, \tag{5.56}$$

and this inclusion can be strict. □

In some cases (see Section 5.10.3) the behaviour of this inclusion turns out to contain important information about the intersection problem (recall that in our context, it is enough to make the decision about the finite generability of the projection $(H \cap H')\pi_{\mathbb{F}}$, see Corollary 5.6.12).

Below we provide an algebraic description of the inclusion (5.63) in terms of intersection schemes. Before we fix some convenient notation.

Remark 5.10.9. Suppose that $\widehat{\widehat{\Upsilon}}_{L,L'}$ is an intersection scheme of rank r normalized with respect to certain spanning tree $\mathbf{T}$. Then every cyclomatic arc in $\widehat{\widehat{\Upsilon}}_{L,L'}$ has the form:

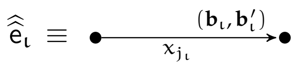

in $\widehat{\widehat{\Upsilon}}_{L,L'}$, the corresponding $\mathbf{T}$-petal has the form:

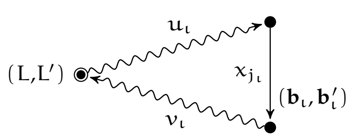

Fig. 5.28: A doubly-enriched petal

Then, using the multiplication rules in $\mathbb{G}_{\mathbf{A}}$, we componentwise read the petal as:

$$\begin{aligned} u_\iota x_{j_\iota} t^{\mathbf{b}_\iota} v_\iota t^{L} &=_{\mathbb{G}_{\mathbf{A}}} u_\iota x_{j_\iota} v_\iota t^{\mathbf{b}_\iota \mathbf{A}_{v_\iota} + L} = w_\iota t^{\mathbf{b}_\iota \mathbf{A}_{v_\iota} + L} =: \widehat{w}_\iota , \\ u_\iota x_{j_\iota} t^{\mathbf{b}'_\iota} v_\iota t^{L'} &=_{\mathbb{G}_{\mathbf{A}}} u_\iota x_{j_\iota} v_\iota t^{\mathbf{b}'_\iota \mathbf{A}_{v_\iota} + L'} = w_\iota t^{\mathbf{b}'_\iota \mathbf{A}_{v_\iota} + L'} =: \widehat{w}'_\iota , \end{aligned} \tag{5.57}$$

where $w_\iota := u_\iota x_{j_\iota} v_\iota$ is the element of $\mathbb{F}_X$ read by the skeleton of the petal in (5.28); and therefore $\{w_\iota\}_\iota$ is a free basis (of size r) of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$. Using these notations, the target subgroup $(H \cap H')\pi_{\mathbb{F}}$ admits a quite compact description.

Proposition 5.10.10. *Let H, H' be subgroups of $\mathbb{G}_{\mathbf{A}}$. Then, the projection of the intersection $H \cap H'$ to the free part is the preimage of the subgroup $L + L' \leqslant \mathbb{Z}^m$ under the map (5.58); i.e.,*

$$(H \cap H')\pi_{\mathbb{F}} = (L + L')\eta^{\leftarrow} ,$$

where

$$\begin{aligned} \eta \colon H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} &\to \mathbb{Z}^m \\ w(\overrightarrow{w_\iota}) &\mapsto \textstyle\sum_{k=1}^{|w|} (\mathbf{b}_{\iota_k} - \mathbf{b}'_{\iota_k})\, \mathbf{A}_{v_{\iota_k}} \prod_{l=k+1}^{|w|+1} \mathbf{A}_{w_{\iota_l}} . \end{aligned} \tag{5.58}$$

Proof. From (5.55), the subgroup $(H \cap H')\pi_{\mathbb{F}}$ consists precisely of those elements $w \in H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ that admit compatible abelian completions in H and H', that is such that $\widehat{w}_{\mathbb{G}_{\mathbf{A}}} \cap \widehat{w}'_{\mathbb{G}_{\mathbf{A}}} \neq \varnothing$.

Let $\boldsymbol{\Upsilon}$ be the Stallings automaton of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ with respect to certain spanning tree $\mathbf{T}$. Let $(w_\iota)_\iota$ be the canonical basis of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ relative to $(\boldsymbol{\Upsilon}, \mathbf{T})$; i.e., $w_\iota = u_\iota x_{j_\iota} v_\iota$, for each arc $e_\iota \in E(\boldsymbol{\Upsilon} \setminus \mathbf{T})$.

Now, since any $\circledbullet$-walk $\widehat{\widehat{\gamma}}$ in $\widehat{\widehat{\boldsymbol{\Upsilon}}}_{L,L'}$ decomposes as a product of (doubly-enriched) $\mathbf{T}$-petals, say

$$\widehat{\widehat{\gamma}} = \gamma_{\mathbf{T}}[\widehat{\widehat{e}}_{\iota_1}]\, \gamma_{\mathbf{T}}[\widehat{\widehat{e}}_{\iota_2}] \cdots \gamma_{\mathbf{T}}[\widehat{\widehat{e}}_{\iota_p}] = \prod_{k=1}^{p} \gamma_{\mathbf{T}}[\widehat{\widehat{e}}_{\iota_k}] ,$$

then, according to the notation in Figure 5.28, the object recognized by $\widehat{\widehat{\gamma}}$ is:

$$\begin{aligned} \widehat{\widehat{w}} &= \prod_{k=1}^{p} u_{\iota_k} x_{\iota_k} t^{(\mathbf{b}_{\iota_k}, \mathbf{b}'_{\iota_k})} v_{\iota_k} \\ &= \prod_{k=1}^{p} u_{\iota_k} x_{\iota_k} v_{\iota_k} t^{(\mathbf{b}_{\iota_k} \mathbf{A}_{v_{\iota_k}}, \mathbf{b}'_{\iota_k} \mathbf{A}_{v_{\iota_k}})} \\ &= \prod_{k=1}^{p} w_{\iota_k} t^{(\mathbf{b}_{\iota_k} \mathbf{A}_{v_{\iota_k}}, \mathbf{b}'_{\iota_k} \mathbf{A}_{v_{\iota_k}})} , \end{aligned} \tag{5.59}$$

where $w = w(\overrightarrow{w_\iota}) = w_{\iota_1} w_{\iota_2} \cdots w_{\iota_p}$ is any element in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ (expressed as a word in the basis $\{w_\iota\}_\iota$).

After breaking down the two components in (5.59), and applying the multiplication rules in $\mathbb{G}_{\mathbf{A}}$, we obtain the expressions:

$$\widehat{w} =_{\mathbb{G}_{\mathbf{A}}} \prod_{k=1}^{p} w_{\iota_k}\, t^{\mathbf{b}_{\iota_k} \mathbf{A}_{v_{\iota_k}} + L} \;=\; w\, t^{\sum_{k=1}^{p} \mathbf{b}_{\iota_k} \mathbf{A}_{v_{\iota_k}} \left(\prod_{l=k+1}^{p+1} \mathbf{A}_{w_{\iota_l}}\right) + L},$$

$$\widehat{w}' =_{\mathbb{G}_{\mathbf{A}}} \prod_{k=1}^{p} w_{\iota_k}\, t^{\mathbf{b}'_{\iota_k} \mathbf{A}_{v_{\iota_k}} + L'} \;=\; w\, t^{\sum_{k=1}^{p} \mathbf{b}'_{\iota_k} \mathbf{A}_{v_{\iota_k}} \left(\prod_{l=k+1}^{p+1} \mathbf{A}_{w_{\iota_l}}\right) + L'},$$

whose intersection is nonempty if and only if

$$\sum_{k=1}^{p} (\mathbf{b}_{\iota_k} - \mathbf{b}'_{\iota_k})\, \mathbf{A}_{v_{\iota_k}} \prod_{l=k+1}^{p+1} \mathbf{A}_{w_{\iota_l}} \in L + L'; \tag{5.60}$$

that is, if and only if

$$\sum_{k=1}^{|w|} \mathbf{d}_{\iota_k} \mathbf{W}_{\mathbf{k}} \in L + L', \tag{5.61}$$

where $\mathbf{d}_{\iota_k} = (\mathbf{b}_{\iota_k} - \mathbf{b}'_{\iota_k})\mathbf{A}_{v_{\iota_k}}$, and $\mathbf{W}_{\mathbf{k}} = \prod_{l=k+1}^{p+1} \mathbf{A}_{w_{\iota_l}}$.

This completes the proof. □

Remark 5.10.11. The map η in (5.58) is not (in general) a homomorphism.

When the intersecting subgroups $H, H' \leqslant \mathbb{G}_{\mathbf{A}}$ are finitely generated, then we can encode the intersection within a normalized intersecting scheme, which is *computable* and *finite* (much in the same form as the standard pullback encodes intersections of two finitely generated subgroups of the free group).

Corollary 5.10.12. *If H, H' are finitely generated subgroups of $\mathbb{G}_{\mathbf{A}}$ given by respective finite generating sets, then an intersection scheme for $H \cap H'$ is computable.*

Proof. This is immediate since the respective Stallings automata for H and H' are clearly computable (and finite) from the finite generating sets; and the product of two finite enriched automata — and its core — are again clearly computable. □

However, as we have already seen, the intersection of two finitely generated subgroups within $\mathbb{G}_{\mathbf{A}}$ does not need to be again finitely generated. So, the question is whether we can *algorithmically* extract certain information (e.g. whether the intersection is finitely generated or not) from the intersecting scheme or not. As we see in the next section, in the case of direct products

5.10.3 The direct product case

It is not surprising that the previous discussion takes a particularly simple (and much more algorithmic friendly) form if the semidirect action $\mathbf{A}_\bullet$ is trivial; namely for direct products of free-abelian and free groups. Moreover, when we consider the constructions in Section 5.10.2 for this case, a particularly appealing geometric structure arises that nicely complements the results in Section 2.3.

Let us rephrase the results from the last section in this setting; namely, for free-abelian times free (FATF) groups. Recall that, if H, H' are subgroups of $\mathbb{F}_X \times \mathbb{Z}^m$, then we denote by r the rank of the product automaton, that is:

$$\begin{aligned} r &:= \operatorname{rk}(\mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}}) \\ &= \operatorname{rk} \mathsf{St}\left(\mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}}\right) \\ &= \operatorname{rk}\left(\mathsf{St}\left(\mathsf{H}\pi_{\mathbb{F}}\right) \times \mathsf{St}\left(\mathsf{H}'\pi_{\mathbb{F}}\right)\right) ; \end{aligned}$$

and by $\{w_\iota\}_\iota$ the basis (of size r) of $\mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}}$ corresponding to some distinguished spanning tree in $\mathsf{St}\left(\mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}}\right)$.

Note that the rank r is not necessarily finite (even if both H and H' are finitely generated). We will denote by $\mathbb{Z}^{\oplus r} := \bigoplus_r \mathbb{Z}$ the free-abelian group of rank r (i.e., the abelianization of $\mathbb{F}_r$).

With the above conventions, Proposition 5.10.10 takes the following form over FATF groups.

Proposition 5.10.13. *Let* H, H' *be subgroups of* $\mathbb{F}_X \times \mathbb{Z}^m$. *Then, the projection of the intersection* $\mathsf{H} \cap \mathsf{H}'$ *to the free part is*

$$(\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}} = \left\{ \mathsf{w}(\overrightarrow{w_\iota}) \in \mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}} \,:\, \mathbf{w}\,\mathbf{B} \in \mathsf{L} + \mathsf{L}' \right\},$$

where $\mathbf{w} = \mathsf{w}^{\mathsf{ab}} \in \mathbb{Z}^{\oplus r}$ *denotes the abelianization of* w, *and* $\mathbf{B} \colon \mathbb{Z}^{\oplus r} \to \mathbb{Z}^m$ *is the linear map given by the* $r \times m$ *integer matrix* $\mathbf{B}$ *having as* i*-th row the element* $\mathbf{b}_i - \mathbf{b}_i' \in \mathbb{Z}^m$. *Equivalently:*

$$(\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}} = (\mathsf{L} + \mathsf{L}')\mathbf{B}^{\leftarrow}\rho^{\leftarrow} = \mathsf{M}\rho^{\leftarrow}, \tag{5.62}$$

where ρ *denotes the abelianization* $\mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}} \simeq \mathbb{F}_r \twoheadrightarrow \mathbb{Z}^{\oplus r}$, *and* $\mathsf{M} := (\mathsf{L} + \mathsf{L}')\mathbf{B}^{\leftarrow}$.

Proof. Taking $\mathbf{A}_w = \mathbf{I_m}$ for every $w \in \mathbb{F}_n$, the map $\eta \colon \mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}} \to \mathbb{Z}^m$ in (5.58) reduces to

$$\mathsf{w}(\overrightarrow{w_\iota}) \mapsto \textstyle\sum_{k=1}^{|\mathsf{w}|} (\mathbf{b}_{\iota_k} - \mathbf{b}'_{\iota_k}) = \sum_{i=1}^{r} |\mathsf{w}|_i (\mathbf{b}_i - \mathbf{b}'_i) = \mathbf{w}\mathbf{B},$$

where we have grouped summands with the same index ι_k, and we have used the notation in the statement (note that r denotes the rank of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$, i.e., the number of distinct w_i's in the basis). □

Observe that now, *contrary to what happens in the general case*, (5.62) describes our target subgroup $(H \cap H')\pi_{\mathbb{F}}$ as the preimage of a homomorphism. This fact is important for us, and has as an immediate consequence the following partial refinement of Lemma 5.10.8.

Corollary 5.10.14. *Let* H, H' *be subgroups of* $\mathbb{F}_n \times \mathbb{Z}^m$. *Then,*

$$(H \cap H')\pi_{\mathbb{F}} \trianglelefteqslant H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}\,, \tag{5.63}$$

and this inclusion can be strict.

Proof. The normality is an immediate consequence of $(H \cap H')\pi_{\mathbb{F}}$ being a full preimage of a (normal) subgroup of an abelian group by a homomorphism. □

The diagram below describes the current situation:

$$\begin{array}{ccccccccc}
\mathbb{F}_n \geqslant H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} & \simeq & \mathbb{F}_r & \overset{\rho}{\twoheadrightarrow} & \mathbb{Z}^{\oplus r} & \overset{\mathbf{B}}{\longrightarrow} & \mathbb{Z}^m \\
\triangledown\!\!/ & & \triangledown\!\!/ & & \triangledown\!\!/ & & \triangledown\!\!/ \\
(H \cap H')\pi_{\mathbb{F}} & \simeq & M\rho^{\leftarrow} & \longleftarrow\!\dashv & M & \longleftarrow\!\dashv & L + L'
\end{array} \tag{5.64}$$

Fig. 5.29: Projected intersection diagram for FATF groups

Now, since all the inclusions in Equation (5.64) are normal, we have the following chain of isomorphisms of (abelian) quotient groups:

$$H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}/(H \cap H')\pi_{\mathbb{F}} \simeq \mathbb{F}_r/M\rho^{\leftarrow} \simeq \mathbb{Z}^{\oplus r}/M \simeq \operatorname{im}\mathbf{B}/(L + L')\,, \tag{5.65}$$

which allows us to see skeletons of enriched Stallings automata, as Cayley digraphs of finitely generated abelian groups.

Lemma 5.10.15. *Let* H, H' *be subgroups of* $\mathbb{F}_X \times \mathbb{Z}^m$*. Then, the subgroup* $(\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}}$ *is either trivial (and hence has trivial Stallings automata), or*

$$\begin{aligned} \mathsf{St}\left((\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}}, \{w_i(X)\}_i\right) &\overset{(1)}{\simeq} \mathsf{St}\left(M\rho^{\leftarrow}, \{w_i\}_i\right) \\ &\overset{(2)}{=} \mathsf{core\,Sch}\left(M\rho^{\leftarrow}, \{w_i\}_i\right) \\ &\overset{(3)}{=} \mathsf{Sch}\left(M\rho^{\leftarrow}, \{w_i\}_i\right) \\ &\overset{(4)}{\simeq} \mathsf{Cay}\left(\mathbb{F}_{\{w_i\}_i}/M\rho^{\leftarrow}, \{[w_i]\}_i\right) \\ &\overset{(5)}{\simeq} \mathsf{Cay}\left(\mathbb{Z}^{\oplus r}/M, \{\mathbf{e}_i\}_i\right) \\ &\overset{(6)}{\simeq} \mathsf{Cay}\left(\operatorname{im}\mathbf{B}/(L + L'), \{\mathbf{e}_i\mathbf{B}\}_i\right), \end{aligned} \tag{5.66}$$

where $\{\mathbf{e}_i\}_i$ *denotes the canonical basis of* $\mathbb{Z}^{\oplus r}$*. (Note that the bases in* (5.66) *must be interpreted as ordered multisets in order to keep track of the link between generators in the corresponding automata.)*

Remark 5.10.16. Note that the subgroup $(\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}} \simeq M\rho^{\leftarrow}$ can only be trivial if $r = 0, 1$ (otherwise it would contain the nontrivial commutator subgroup of a nonabelian free group). More precisely:

$$(\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}} = \{1\} \;\Leftrightarrow\; \begin{cases} r = 0\,, \text{or} \\ r = 1 \text{ and } M = \{\mathbf{0}\}\,. \end{cases} \tag{5.67}$$

Proof. The isomorphism (1) (of automata) is clear from the group isomorphism $(\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}} \simeq M\rho^{\leftarrow}$, and Corollary 5.4.25; see Figure 5.29.

The equality (2) is just an instance of Proposition 5.4.30.

Since $M\rho^{\leftarrow} \simeq (\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}}$ is nontrivial and normal in $\mathbb{F}_r$, then the Schreier digraph of $M\rho^{\leftarrow}$ has no superfluous vertices, and coincides with the Cayley digraph of the quotient. Equality (3) and isomorphism (4) follow.

The isomorphism (5) is clear since the abelianization epimorphism $\rho\colon \mathbb{F}_r \twoheadrightarrow \mathbb{Z}^r$ induces an isomorphism $\mathbb{F}_r/M\rho^{\leftarrow} \simeq \mathbb{Z}^r/M$.

The isomorphism (6) is clear from the group isomorphism $\mathbb{Z}^{\oplus r}/M \simeq \operatorname{im}\mathbf{B}/(L + L')$ induced by $\mathbf{B}\colon \mathbb{Z}^{\oplus r} \twoheadrightarrow \operatorname{im}\mathbf{B}$. □

Now we claim that the previous discussion applied to finitely generated subgroups $\mathsf{H}, \mathsf{H}' \leqslant \mathbb{F}_n \times \mathbb{Z}^m$ allows us to solve the subgroup intersection problem. To this end, the normality of $(\mathsf{H} \cap \mathsf{H}')\pi_{\mathbb{F}}$ in $\mathsf{H}\pi_{\mathbb{F}} \cap \mathsf{H}'\pi_{\mathbb{F}}$ plays a key role, since it allows us to relate finite generability with finite index, which we will be able to check algorithmically.

Recall that if the subgroups H, H' are given by respective finite sets of generators, then we can compute a finite normalized intersection scheme for $H \cap H'$; and in particular, the (finite) rank r, the matrix $\mathbf{B}$, and the subgroup $L + L'$ are given and finite as well. Hence, the equation $\mathbf{wB} \in L + L'$ is solvable using standard linear algebra over $\mathbb{Z}$.

Moreover, suppose that we have already computed a basis $\{\mathbf{m_1}, \ldots, \mathbf{m_s}\}$ $(s \leqslant r)$ for the subgroup $M \leqslant \mathbb{Z}^r$ of solutions to the equation $\mathbf{wB} \in L + L'$; and denote by $\mathbf{M}$ the $s \times r$ integer matrix having as i-th row the element $\mathbf{m_i} \in \mathbb{Z}^m$. Then, we can apply the Smith normal form (SNF) decomposition to $\mathbf{M}$ and compute matrices $\mathbf{P} \in \mathrm{GL}_s(\mathbb{Z})$, $\mathbf{Q} \in \mathrm{GL}_r(\mathbb{Z})$, and an integer $s \times r$ diagonal matrix

$$\mathbf{D} = \mathbf{diag}(\delta_1, \ldots, \delta_s)\,,$$

where $s \leqslant r$, and $\delta_1 \mid \delta_2 \mid \cdots \mid \delta_s$ are strictly positive integers successively dividing each other; such that $\mathbf{PMQ} = \mathbf{D}$.

So, we can express the solution subgroup $M \leqslant \mathbb{Z}^r$ as the row space of any of the matrices $\mathbf{M}, \mathbf{PM}$, or $\mathbf{DQ}^{-1}$. That is,

$$M = \langle \mathbf{M} \rangle = \langle \mathbf{PM} \rangle = \langle \mathbf{DQ}^{-1} \rangle\,. \tag{5.68}$$

Note that this means that

$$M = \langle\, \delta_1 \mathbf{d_1}, \ldots, \delta_r \mathbf{d_r} \,\rangle\,,$$

where $\mathbf{d_i}$ is the i-th row of $\mathbf{Q}^{-1}$, and we have defined $\delta_j := 0$, for all $j = s+1, \ldots, r$.

Hence, in the finitely generated case we can rewrite Equation (5.65) more precisely as:

$$\begin{aligned}
H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} / (H \cap H')\pi_{\mathbb{F}} &\simeq \mathbb{Z}^r / \langle \mathbf{M} \rangle && (5.69)\\
&\simeq \mathbb{Z}^r / \langle \mathbf{DQ}^{-1} \rangle && (5.70)\\
&\simeq \mathbb{Z}^r / \langle \mathbf{D} \rangle && (5.71)\\
&= \textstyle\bigoplus_{i=1}^{s} \mathbb{Z}/\delta_i \oplus \mathbb{Z}^{r-s} && (5.72)\\
&= \textstyle\bigoplus_{i=1}^{r} \mathbb{Z}/\delta_i\,. && (5.73)
\end{aligned}$$

In particular, if both H and H' are finitely generated, then the index of $(H \cap H')\pi_{\mathbb{F}}$ in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ is:

$$\left[H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}} : (H \cap H')\pi_{\mathbb{F}} \right] = \textstyle\prod_{i=1}^{r} \left[\mathbb{Z} : \delta_i \mathbb{Z} \right] = \begin{cases} \prod_{i=1}^{s} \delta_i, & \text{if } s = r\,, \\ \aleph_0, & \text{if } s < r\,; \end{cases} \tag{5.74}$$

and Lemma 5.10.15 takes the form below.

Proposition 5.10.17. *Let* H, H' *be finitely generated subgroups of* $\mathbb{F}_X \times \mathbb{Z}^m$ *(given by respective finite generating sets). Then, the Stallings automaton of* $(H \cap H')\pi_{\mathbb{F}}$ *w.r.t. the basis* $\{w_i(X)\}_i$ *of* $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ *is either trivial, or:*

$$\mathsf{St}\left((H \cap H')\pi_{\mathbb{F}}, \{w_i(X)\}_i\right) = \mathsf{Cay}\left(\bigoplus_{i=1}^{r} \mathbb{Z}/\delta_i\mathbb{Z}, \{\mathbf{e}_i\mathbf{Q}\}_i\right), \tag{5.75}$$

where $\mathbf{PMQ} = \mathbf{diag}(\delta_1, \ldots, \delta_s)$ *is a SNF decomposition of the the matrix* $\mathbf{M}$ *having as rows the elements of a basis of the general solution of the equation* $\mathbf{wB} \in L + L'$*, and* $\{\mathbf{e}_i\mathbf{Q}\}_i$ *are the rows of* $\mathbf{Q}$.

Moreover, if the subgroup $(H \cap H')\pi_{\mathbb{F}}$ *turns out to be finitely generated (which is checkable algorithmically, see Proposition 5.10.20), then its Stallings automaton in (5.75) is computable.*

Proof. Note that $(H \cap H')\pi_{\mathbb{F}}$ can only be trivial if $r = 0, 1$ (otherwise it would contain the nontrivial commutator subgroup of $\mathbb{F}_r$), and this is always checkable algorithmically if H and H' are finitely generated. Namely, r is the rank of the (computable) Stallings automata of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$; and then, if $r = 0$, then $(H \cap H')\pi_{\mathbb{F}}$ is necessarily trivial; whereas if $r = 1$ (i.e., $\rho = \mathrm{id}_{\mathbb{Z}}$), then $(H \cap H')\pi_{\mathbb{F}}$ is trivial if and only if $M = \{\mathbf{0}\}$ (which is again a computable condition).

Otherwise (if $(H \cap H')\pi_{\mathbb{F}}$ is nontrivial) consider the following sequence of isomorphisms:

$$\begin{aligned}
\mathsf{St}\left((H \cap H')\pi_{\mathbb{F}}, \{w_i(X)\}_i\right) &\overset{(1)}{\simeq} \mathsf{Cay}\left(\mathbb{Z}^r/M, \{\mathbf{e}_i\}_i\right) \\
&\overset{(2)}{=} \mathsf{Cay}\left(\mathbb{Z}^r/\langle \mathbf{M} \rangle, \{\mathbf{e}_i\}_i\right) \\
&\overset{(3)}{=} \mathsf{Cay}\left(\mathbb{Z}^r/\langle \mathbf{D}\,\mathbf{Q}^{-1} \rangle, \{\mathbf{e}_i\}_i\right) \\
&\overset{(4)}{\simeq} \mathsf{Cay}\left(\mathbb{Z}^r/\langle \mathbf{D} \rangle, \{\mathbf{e}_i\mathbf{Q}\}_i\right) \\
&\overset{(5)}{=} \mathsf{Cay}\left(\bigoplus_{i=1}^{r} \mathbb{Z}/\delta_i\mathbb{Z}, \{\mathbf{e}_i\mathbf{Q}\}_i\right),
\end{aligned} \tag{5.76}$$

where $\{\mathbf{e}_i\}_i$ denotes the canonical basis of $\mathbb{Z}^r$, and hence $\{\mathbf{e}_i\mathbf{Q}\}_i$ are the rows of $\mathbf{Q}$.

Observe that all the steps in (5.76) are clear: the isomorphism (1) corresponds exactly to (5.66) in the finitely generated case, whereas the rest are easy consequences of the equalities in (5.68). In particular, the isomorphism (4) (of automata) corresponds to the isomorphism (of groups) $\mathbb{Z}^r/\langle \mathbf{DQ}^{-1} \rangle \simeq \mathbb{Z}^r/\langle \mathbf{D} \rangle$ induced by $\mathbf{Q} \in \mathrm{GL}_r(\mathbb{Z})$.

For the last claim, note that if the subgroups (H, H') are finitely generated, then we can compute an intersection scheme for $H \cap H'$, which contains all the data in the equation $\mathbf{w}\,\mathbf{B} \in L + L'$, from whose (computable) set of solutions we can obtain all the parameters in (5.75). And, furthermore, if the projection $(H \cap H')\pi_{\mathbb{F}}$

is finitely generated, then the Cayley digraph — and so the Stallings automaton — in Equation (5.76) is computable. □

Rewriting the edges in $\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, \{w_i(X)\}_i)$ in terms of the original generators X and reducing, we get the corresponding Stallings automata w.r.t. X.

Corollary 5.10.18. *Let H, H' be subgroups of $\mathbb{F}_X \times \mathbb{Z}^m$. Then, the Stallings automaton of the projection $(H \cap H')\pi_{\mathbb{F}}$ is either the trivial automaton (if $(H \cap H')\pi_{\mathbb{F}}$ is trivial); or otherwise, the Stallings reduction of the automaton obtained after replacing every arc labelled by $w_i(X)$ in $\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, \{w_i(X)\}_i)$, by the corresponding X-walk reading $w_i(X)$.*

Moreover, if the intersecting subgroups H, H' are finitely generated (i.e., given by respective finite *sets of generators), and the projection $(H \cap H')\pi_{\mathbb{F}}$ is also finitely generated, then the Stallings automaton $\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, X)$ is indeed computable.*

Proof. Since the resulting automaton is reduced by construction, it is enough to prove that the resulting inverse X-automaton recognizes the subgroup $(H \cap H')\pi_{\mathbb{F}}$. But this is obvious since it clearly recognizes exactly the same elements as the corresponding Stallings automaton w.r.t. $\{w_i(X)\}_i$). Finally, the unicity among inverse X-automata recognizing the same subgroup (Corollary 5.4.25).

For the second claim, recall that if H and H' are given by respective *finite* sets of generators, and the projection $(H \cap H')\pi_{\mathbb{F}}$ is also finitely generated, then the Stallings automaton $\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, \{w_i(X)\}_i)$ is finite and computable (Proposition 5.10.17); and both the substitutions $w_\iota \to w_\iota(X)$, and the final Stallings reduction can be performed algorithmically. □

Note also, that no closed folding is possible in the folding process in (2), since both inverse automata must recognize the same subgroup, and hence have the same rank, i.e., $\mathrm{rk}(\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, X)) = \mathrm{rk}(\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, \{w_i(X)\}_i))$.

Corollary 5.10.19. *Let H, H' be finitely generated subgroups of $\mathbb{F}_n \times \mathbb{Z}^m$, and let r denote the (finite) rank of $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$. Then,*

(a) *if $r = 0$, then $\mathrm{rk}((H \cap H')\pi_{\mathbb{F}}) = 0$;*

(b) *if $r = 1$, then $\mathrm{rk}((H \cap H')\pi_{\mathbb{F}}) = 0$ if $\delta_1 = 0$, and $\mathrm{rk}((H \cap H')\pi_{\mathbb{F}}) = 1$ otherwise.*

(c) *if $r \geqslant 2$, then*

$$\overline{\mathrm{rk}}\left((H \cap H')\pi_{\mathbb{F}}\right) = \overline{\mathrm{rk}}\left(H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}\right) \cdot \prod_{i=1}^{r} \left[\mathbb{Z} : \delta_i \mathbb{Z}\right], \tag{5.77}$$

where $\overline{\mathrm{rk}}(A) \coloneqq \max\{\mathrm{rk}(A) - 1, 0\}$ denotes the reduced rank of a subgroup A.

Proof. The case $r = 0$ is trivial.

In the case where $r = 1$, the (cyclic) subgroup $(H \cap H')\pi_{\mathbb{F}}$ is trivial if and only if $M = \{\mathbf{0}\}$, or equivalently if $\delta_1 = 0$. The claimed result follows.

For $r \geqslant 2$ the result follows easily from Equation (5.74). If the index of $(H \cap H')\pi_{\mathbb{F}}$ in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ is finite, then it corresponds precisely to the Schreier index formula (5.28). Otherwise, $(H \cap H')\pi_{\mathbb{F}}$ is a nontrivial normal subgroup of infinite index in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ and hence has infinite rank; and on the other hand, at least one of the δ_i's is zero, i.e., at least one of the indices $[\mathbb{Z} : \delta_i\mathbb{Z}]$ is infinite, and hence the right hand side of (5.77) is infinite as well. □

This corollary provides the following characterization, that allows us to immediately solve the intersection decision problem.

Proposition 5.10.20. *Let H, H' finitely generated subgroups of $\mathbb{F}_n \times \mathbb{Z}^m$. Then, the following conditions are equivalent:*

(a) *The intersection $H \cap H'$ is finitely generated.*

(b) *The projection subgroup $(H \cap H')\pi_{\mathbb{F}}$ is finitely generated.*

(c) *The subgroup $(H \cap H')\pi_{\mathbb{F}}$ is either trivial, or have finite index in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$.*

(d) *Either $r = 0, 1$ and the subgroup M is trivial; or M has finite index in $\mathbb{Z}^r$.*

(e) *Either $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ is trivial (i.e., $r = 0$), or all the multiplicities δ_i $(i = 1, \ldots, r)$ are strictly positive.*

(a) ⇔ (b). This is a particular case of Corollary 5.6.12.

[(b) ⇔ (c)] This is a direct consequence of Corollary 5.10.19: if $(H \cap H')\pi_{\mathbb{F}} = \{1\}$, then it is obviously finitely generated and there is is nothing to prove; otherwise the result follows immediately from Equation (5.77) (which still holds for $r = 1$, if M is nontrivial).

[(c) ⇔ (d)] On one side, from Remark 5.10.16, $(H \cap H')\pi_{\mathbb{F}} = \{1\}$ if and only if either $r = 0$, or $r = 1$ and $M = \{\mathbf{0}\}$ (note that in both cases the subgroup M is trivial). On the other side, since $(H \cap H')\pi_{\mathbb{F}}$ is a full homomorphism preimage of M, it has finite index in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ if and only if M has finite index in $\mathbb{Z}^r$.

[(d) ⇔ (e)] This is clear from Equation (5.77) (which still holds for $r = 1$, if M is nontrivial). □

So far we have described the skeleton $\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, X)$ of any Stallings automata for the intersection of two subgroups H, H' of a FATF group in terms of the Cayley graph of a certain finitely generated abelian group, and we have seen it to be computable if the intersecting subgroups and the intersection itself are finitely generated. Below we show that in order to enrich this skeleton to a genuine Stallings automaton for the intersection $H \cap H'$, it is enough to adapt the skeleton procedure to the *abelianly completed* generators given for the subgroups H and H'.

Theorem 5.10.21. *Let H, H' be subgroups of $\mathbb{F}_X \times \mathbb{Z}^m$. Then, a normalized Stallings automaton for the intersection $H \cap H'$ is the result of attaching to the basepoint of the skeleton $\mathsf{St}((H \cap H')\pi_{\mathbb{F}}, X)$ the subgroup $L \cap L' = H \cap H' \cap \mathbb{Z}^m$, and then (after distinguishing an spanning tree) labelling the end of every cyclomatic edge e_ι with any representative of the (nonempty) coset intersection $\mathcal{C}_L(\nu_\iota) \cap \mathcal{C}_{L'}(\nu_\iota)$, where ν_ι denotes the free word in X read by the ι-th petal in the skeleton.*

Proof. Since the resulting enriched automaton — say $\widehat{\Gamma}_{L \cap L'}$ — is reduced and normalized by construction, it is enough to show that it recognizes the subgroup $H \cap H'$.

But this is clear, since $\widehat{\Gamma}_{L \cap L'}$ has been built in such a way that:

(i) its basepoint subgroup $L \cap L'$ is the full basegroup of the intersection $H \cap H'$ (see Proposition 5.6.11).

(ii) its skeleton recognizes precisely the elements in the projection $(H \cap H')\pi_{\mathbb{F}}$ (that is, the elements in $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$ that admit compatible completions in H and H'), and

(iii) its abelian labelling provides to the elements in $(H \cap H')\pi_{\mathbb{F}}$ precisely with the completion that we know they have by the previous comment.

Finally, recall that these are precisely the conditions in Proposition 5.10.5 for an element to be in the intersection of the subgroups H and H'. □

Now the computability of finitely generated intersections is straightforward.

Corollary 5.10.22. *If H and H' are finitely generated subgroups of $\mathbb{F}_X \times \mathbb{Z}^m$ with finitely generated intersection, then a normalized Stallings automaton for $H \cap H'$ is computable.*

Proof. It is enough to realize that, if the intersecting groups H, H' and the intersection $H \cap H'$ are finitely generated, then every ingredient in Theorem 5.10.21 of the construction of the Stallings of an Stallings automaton for $H \cap H'$ is computable:

(i) since H and H' are given by finite families of generators, respective normalized Stallings automata for H and H' — including finite basis for their full basegroups L and L' — are computable (Corollary 5.9.21).

(ii) The skeleton $\operatorname{St}((H \cap H')\pi_{\mathbb{F}}, X)$ has already shown to be computable under the assumptions of the theorem (see Corollary 5.10.18).

(iii) A basis for the basepoint subgroup $L \cap L'$ is easily computable from the finite basis for H and H' in step (i) using linear algebra.

(iv) Representatives for the (necessarily nonempty) intersections $\mathcal{C}_L(\nu_\iota) \cap \mathcal{C}_{L'}(\nu_\iota)$ of abelian completions in every petal, are again computable using standard linear algebra.

(v) Since the resulting enriched Stallings automaton recognizing $H \cap H$ is finite, we can clearly normalize it algorithmically.

This concludes the proof. □

Our target result (below) follows easily from Proposition 5.10.20, and Corollary 5.10.22.

Theorem 5.10.23. *The (full) subgroup intersection problem is solvable for free-abelian times free groups.* □

Proof. The decision problem is solvable using Proposition 5.10.20: given two finitely generated subgroups $H, H' \leqslant \mathbb{F}_n \times \mathbb{Z}^m$ given by respective finite families of generators, one can decide whether the intersection is finitely generated or not computing an scheme for the intersection $H \cap H'$, and from it the data appearing in any of the last two conditions in Proposition 5.10.20. For example, computing the integers $r = \operatorname{rk}(H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}})$ and $s = \operatorname{rk} M$, and then:

- if $r = 0$, then answer YES; and otherwise
- if $s = r$ then answer YES, and if $s < r$ then answer NO.

In case that the intersection $H \cap H'$ is finitely generated (i.e., when the previous algorithm answers YES); then the (computable) Stallings automaton $\widehat{\Gamma}_{L \cap L'}$ for the intersection $H \cap H'$ provides a finite generating set in the usual way. Namely, a generating set for the intersection $H \cap H'$ consists of the following:

(i) A free-abelian basis A of the base subgroup $L \cap L'$; and

(ii) A free basis B for the section subgroup $(H \cap H')\pi_{\mathbb{F}}\sigma$ obtained from the words read by the petals in $\widehat{\Gamma}$ after distinguishing an spanning tree.

Since both A and B are clearly computable under our assumptions, then their (disjoint) union $A \sqcup B$ constitutes a (finite) computable generating set for the intersection $H \cap H'$, and the proof is concluded. □

Examples

As a first example of application of these graphical techniques, let us reconsider under this viewpoint the situation given in Lemma 2.3.1 as a counterexample for Howson's property in FATF groups.

Example 5.10.24. Let $H = \langle x, y \rangle$ and $H' = \langle xt, y \rangle$ be subgroups of the FATF group $\mathbb{F}_2 \times \mathbb{Z} = \langle x, y \mid - \rangle \times \langle t \mid - \rangle$. Then, $L = L' = L \cap L' = L + L' = \{0\}$, and enriched Stallings automata of H, and H′ are:

$\widehat{\mathsf{St}}(H) \equiv$ 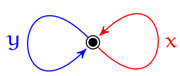

and $\widehat{\mathsf{St}}(H') \equiv$ 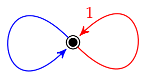

;

and therefore, $\widehat{\mathsf{St}}(H) \wedge \widehat{\mathsf{St}}(H') = \widehat{\mathsf{St}}(H) \times \widehat{\mathsf{St}}(H') \equiv$ 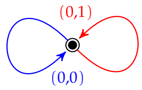

.

Note that from this intersection scheme it is clear that one can not equalize the two components of the red label (modulo 0) after any finite number of turns around the red loop (corresponding to the generator x), and therefore the intersection must be infinitely generated.

Formally, we have $\mathbf{B} = \left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right]$, and the matrix $\mathbf{M} = [\,0\ 1\,]$ has as row a basis for the set of solutions of the equation $\mathbf{wB} = \mathbf{0}$. Hence, the SNF of $\mathbf{M}$ is $\mathbf{D} = [\,1\ 0\,]$, with $\mathbf{P} = [\,1\,]$, and $\mathbf{Q} = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$.

Then, applying Proposition 5.10.17, we have that

$$\mathsf{St}\left((H \cap H')\pi_{\mathbb{F}}, \{x, y\}\right) \simeq \mathsf{Cay}\left(\mathbb{Z}/\mathbb{Z} \oplus \mathbb{Z}/0\mathbb{Z}, \{(0,1),(1,0)\}\right) \simeq \mathsf{Cay}\left(\mathbb{Z}, \{1, 0\}\right),$$

which, after substituting x by xt (and y by y) takes the form:

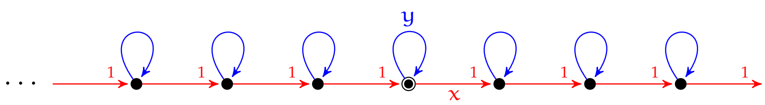

;

and after removing the abelian labels in the red arcs (which are bridges) constitutes the Stallings automaton of the intersection:

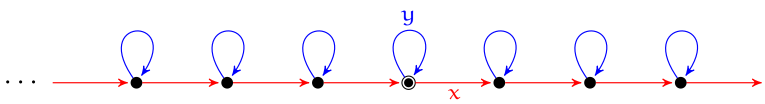

;

which obviously coincide with the one computed in Lemma 2.3.1.

Finally, we consider a parameterized example in order to show how very different situations can emerge from similar intersecting patterns.

Example 5.10.25. Consider the following two parameterized subgroups of the direct product $\mathbb{F}_{\{x,y\}} \times \mathbb{Z}^m$,

$$\begin{aligned} H &= \langle t^{\mathbf{L}}, x^3 t^{\mathbf{a}}, yx\, t^{\mathbf{b}}, y^3xy^{-2} t^{\mathbf{c}} \rangle, \\ H' &= \langle t^{\mathbf{L}'}, x^2 t^{\mathbf{d}}, yxy^{-1} \rangle, \end{aligned}$$

where $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d} \in \mathbb{Z}^m$, and L, L' are subgroups of $\mathbb{Z}^m$.

In order to compute (a basis for) the intersection $H \cap H'$ we first compute reduced automata for H and H', and then build its product (recall Examples 5.9.14 and 5.5.30).

(Note that we don't need a unique representative until reaching a scheme for the subgroup of the intersection, so reduced automata $\widehat{\Gamma}_L, \widehat{\Gamma}'_{L'}$ for $H, H'$are enough here.)

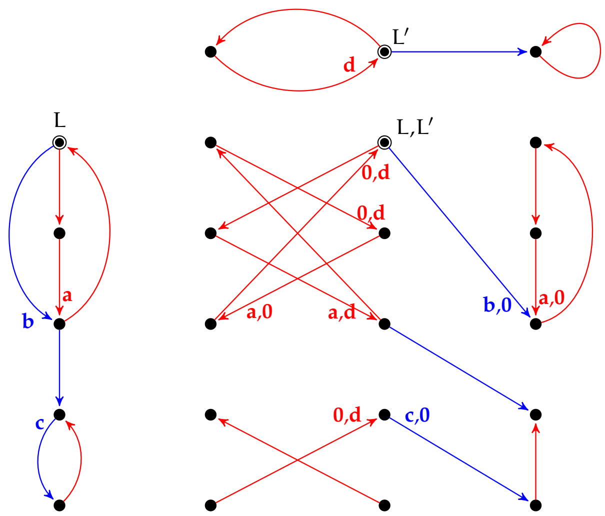

Fig. 5.30: Product of enriched automata

After removing non-basepoint components, performing a total trim, and rearranging, we obtain the core of the enriched product; namely:

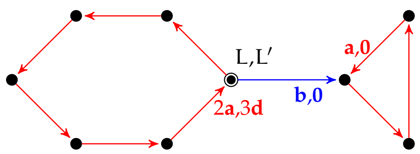

Fig. 5.31: Intersection scheme for $H \cap H'$

Finally, after normalizing (recall Lemma 5.8.13) we obtain a reduced automaton for $H \cap H'$:

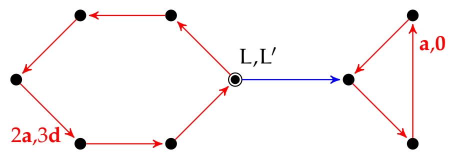

Fig. 5.32: Normalized intersection scheme for $H \cap H'$

Remark 5.10.26. Note that the abelian labels $\mathbf{b}$ and $\mathbf{c}$ in Figure 5.30 no longer appear in the normalized intersection scheme (Figure 5.32); the first one because it labels a bridge (Lemma 5.8.13), and the second one because it lies outside the core. Thus, neither $\mathbf{b}$ nor $\mathbf{c}$ will play any role in the intersection $H \cap H'$.

So, we have that

$$\begin{cases} w_1 = x^6, \\ w_2 = yx^3y^{-1} \end{cases} \tag{5.78}$$

is a basis for $H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$.

In particular, $x^6 \in H\pi_{\mathbb{F}} \cap H'\pi_{\mathbb{F}}$. Now, according to Proposition 5.10.5, $x^6 \in (H \cap H')\pi_{\mathbb{F}}$ if and only if it is componentwise readable in the intersection scheme in Figure 5.32; i.e., if there exists an abstract (reduced) word $w \in \mathbb{F}_{\{x,y\}}$, such that

$$\begin{aligned} x^6\, t^{\mathbf{0}} &\in w\big(\overrightarrow{x^6\, t^{2\mathbf{a}+L},\, yx^3y^{-1}\, t^{\mathbf{a}+L}}\big)_{G_A} \cap w\big(\overrightarrow{x^6\, t^{3\mathbf{d}+L'},\, yx^3y^{-1}\, t^{\mathbf{0}+L'}}\big)_{G_A} \\ &= w\big(\overrightarrow{x^6, yx^3y^{-1}}\big)\; t^{\lambda_1 2\mathbf{a}+\lambda_2\mathbf{a}+L} \cap w\big(\overrightarrow{x^6, yx^3y^{-1}}\big)\; t^{\lambda_1 3\mathbf{d}+\lambda_2\mathbf{0}+L'} \\ &= w\big(\overrightarrow{x^6, yx^3y^{-1}}\big)\; t^{(\lambda_1 2\mathbf{a}+\lambda_2\mathbf{a}+L)\,\cap\,(\lambda_1 3\mathbf{d}+\lambda_2\mathbf{0}+L')}\,, \end{aligned}$$

where $w^{ab} = (\lambda_1, \lambda_2)$. Now, since $x^6 \in w\big(\overrightarrow{x^6, yx^3b^{-1}}\big)$, we necessarily have that $w(w_1, w_2) = w_1$, and so $\lambda_1 = 1$ and $\lambda_2 = 0$. Thus, $x^6 \in (H \cap H')\pi_{\mathbb{F}}$ if and only if

$$\mathbf{0} \in (2\mathbf{a} + L) \cap (3\mathbf{d} + L')\,,$$

which is equivalent to the condition

$$2\mathbf{a} \in L \quad \text{and} \quad 3\mathbf{d} \in L'\,. \tag{5.79}$$

This example shows that, the inclusion $(H \cap H')\pi_{\mathbb{F}} \leqslant (H)\pi_{\mathbb{F}} \cap (H')\pi_{\mathbb{F}}$ can be strict, as claimed in Lemma 5.10.8. Indeed, as we have seen in Proposition 5.10.20, the index of this inclusion is directly related with the intersection problem.

Let's now study the intersection $H \cap H'$ in our example. According to Proposition 5.10.13, its projection $(H \cap H')\pi_{\mathbb{F}}$ is described by the words $w \in \mathbb{F}_2$, such that their abelianization $\mathbf{w} \in \mathbb{Z}^2$ satisfies the equation:

$$\mathbf{w}\,\mathbf{B} \in L + L'\,, \tag{5.80}$$

where $\mathbf{B}$ is the matrix having as i-th row the difference between the two enrichments in the i-th cyclotomic arc of the intersection scheme; in our example

$$\mathbf{B} = \begin{bmatrix} 2\mathbf{a} - 3\mathbf{d} \\ \mathbf{a} \end{bmatrix}.$$

We now distinguish different cases depending on the parameters $\mathbf{a}, \mathbf{d} \in \mathbb{Z}^2$, and the subgroups $L, L' \leqslant \mathbb{Z}^2$:

Case 1: Suppose $\mathbf{a} = (1,0)$, $\mathbf{d} = (0,1) \in \mathbb{Z}^2$, and $L_1 = \langle (0,6) \rangle$, $L_1' = \langle (3,-3) \rangle \leqslant \mathbb{Z}^2$. That is, consider

$$\begin{aligned} H_1 &= \langle t^{(0,6)}, x^3 t^{(1,0)}, yx\, t^{\mathbf{b}}, y^3 x y^{-2} t^{\mathbf{c}} \rangle, \\ H_1' &= \langle t^{(3,-3)}, x^2 t^{(0,1)}, yxy^{-1} \rangle. \end{aligned}$$

Then, $L_1 \cap L_1' = \{(0,0)\}$, $\mathbf{B} = \begin{bmatrix} 2\mathbf{a} - 3\mathbf{d} \\ \mathbf{a} \end{bmatrix} = \begin{bmatrix} 2 & -3 \\ 1 & 0 \end{bmatrix}$; and the set of solutions for

$$\mathbf{w} \begin{bmatrix} 2 & -3 \\ 1 & 0 \end{bmatrix} \in \langle (0,6), (3,-3) \rangle$$

is easily computable to be the subgroup M generated by the rows of the matrix $\mathbf{M} = \begin{bmatrix} -2 & 4 \\ 1 & 1 \end{bmatrix}$; which, in turn, admits the Smith normal form decomposition $\mathbf{PMQ} = \mathbf{D}$, where $\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix}$, $\mathbf{Q} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}$, and $\mathbf{D} = \begin{bmatrix} 1 & 0 \\ 0 & 6 \end{bmatrix}$.

Therefore, according to Proposition 5.10.17:

$$\begin{aligned} \mathsf{St}\left(M\rho^{\leftarrow}, \{w_1, w_2\}\right) &= \mathsf{Cay}\left(\mathbb{F}_{\{w_1,w_2\}}/M\rho^{\leftarrow}, \{[w_1], [w_2]\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}^2/\langle \mathbf{M} \rangle, \{(1,0),(0,1)\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}^2/\langle \mathbf{D}\mathbf{Q}^{-1} \rangle, \{(1,0),(0,1)\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}^2/\langle \mathbf{D} \rangle, \{(1,-1),(0,1)\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}/\mathbb{Z} \oplus \mathbb{Z}/6\mathbb{Z}, \{(1,-1),(0,1)\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}/6\mathbb{Z}, \{-1,1\}\right), \end{aligned}$$

which is easily constructible by inspection: denoting by a violet (resp., green) arc the action of the element -1 (resp., 1), we obtain the automaton:

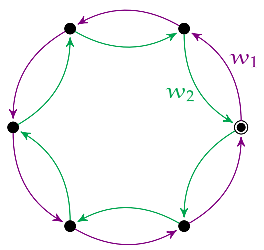

Fig. 5.33: Stallings automaton $\mathrm{St}\,(M\rho^{\leftarrow}, \{w_1, w_2\})$

which after replacing w_1 by x^6, w_2 by yx^3y^{-1}, and folding, becomes:

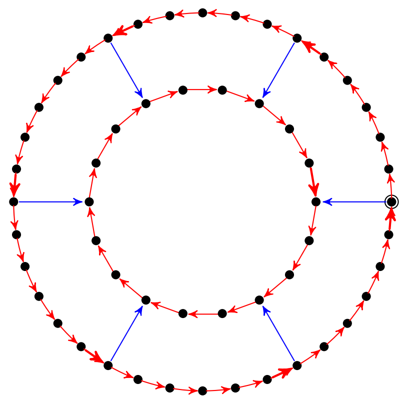

Fig. 5.34: Stallings automaton $\mathrm{St}\,\big((H_1 \cap H_1')\pi_{\mathbb{F}}, \{x, y\}\big)$

i.e., the Stallings automaton of the projection $H_1 \cap H_1')\pi_{\mathbb{F}}$ (which has rank equal to 7). Now we distinguish a maximal tree **T** (using thicker lines to denote its cyclotomic arcs), which provides the basis:

$$\begin{cases} v_1 = y\,x^3\,y^{-1}\,x^6\,, \\ v_2 = y\,x^6\,y^{-1}\,x^6\,y\,x^{-3}\,y^{-1}\,, \\ v_3 = y\,x^9\,y^{-1}\,x^6\,y\,x^{-6}\,y^{-1}\,, \\ v_4 = y\,x^{12}\,y^{-1}\,x^6\,y\,x^{-9}\,y^{-1}\,, \\ v_5 = y\,x^{15}\,y^{-1}\,x^6\,y\,x^{-12}\,y^{-1}\,, \\ v_6 = y\,x^{18}\,y^{-1}\,, \\ v_7 = x^6\,y\,x^{-12}\,y^{-1}\,, \end{cases} \tag{5.81}$$

for the projection $(H_1 \cap H_1')\pi_{\mathbb{F}}$.

In the same vein, after replacing w_1 by $x^6\,t^{(2\mathbf{a},3\mathbf{d})} = x^6\,t^{(2,0),(0,3)}$, and w_2 by $yx^3y^{-1}\,t^{(\mathbf{a},\mathbf{0})} = yx^3y^{-1}\,t^{(1,0),(0,0)}$, and folding we obtain (abbreviating $\mathbf{u} \coloneqq (2\mathbf{a}, 3\mathbf{d})$, and $\mathbf{v} \coloneqq (\mathbf{a}, \mathbf{0})$):

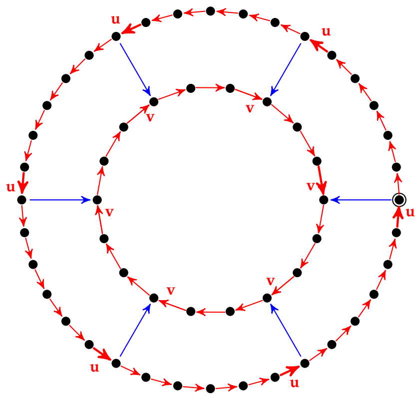

Fig. 5.35: Expanded intersection scheme for $H_1 \cap H_1'$

In order to normalize this doubly enriched automaton (w.r.t. $\mathbf{T}$), consider the following transformation between neighborhoods of clockwise successive blue arcs (note that these neighborhoods contain every abelian label) in Figure 5.35.

which after 5 applications (starting from the first blue neighborhood after the basepoint, and moving clockwise) provides the normalized automaton:

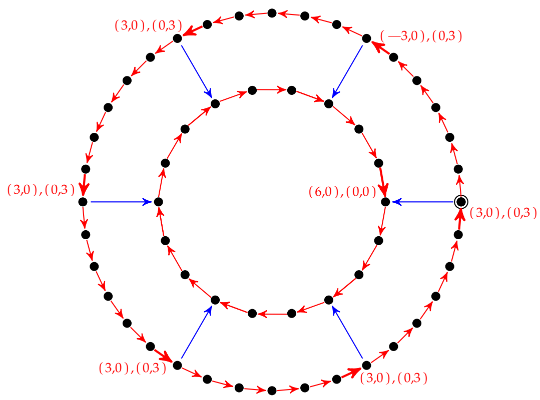

Fig. 5.36: Normalized expanded intersection scheme for $H_1 \cap H_1'$

Now, we know that every $\mathbf{T}$-petal in Figure 5.36 must have compatible abelian completions; i.e., must be readable in some way in both enriched components. Let us compute these intersections for every $\mathbf{T}$-petal:

$$\begin{aligned}
\mathcal{C}_{L_1}(v_6) \cap \mathcal{C}_{L_1'}(v_6) &= (6,0) + \langle (0,6) \rangle \cap (0,0) + \langle (3,-3) \rangle = \{(6,-6)\}, \\
\mathcal{C}_{L_1}(v_7) \cap \mathcal{C}_{L_1'}(v_7) &= (-3,0) + \langle (0,6) \rangle \cap (0,3) + \langle (3,-3) \rangle = \{(-3,6)\} \\
\mathcal{C}_{L_1}(v_i) \cap \mathcal{C}_{L_1'}(v_i) &= (3,0) + \langle (0,6) \rangle \cap (0,3) + \langle (3,-3) \rangle = \{(3,0)\},
\end{aligned} \tag{5.82}$$

for $i = 1, \ldots, 5$.

Finally, replacing each double abelian labelling in Figure 5.36 by one representative from the corresponding intersection class in (5.82), we obtain a Stallings automaton for the intersection $H_1 \cap H_1'$:

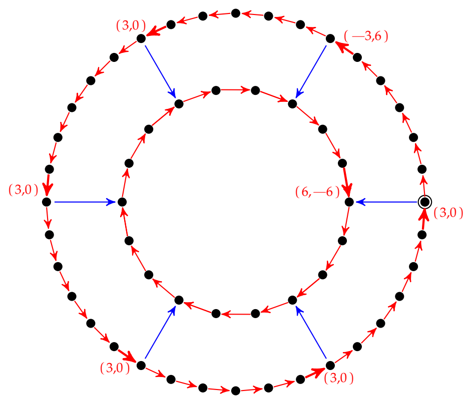

Fig. 5.37: Stallings automaton for $H_1 \cap H_1'$

which provides the basis:

$$\begin{aligned}
\widehat{v}_1 &= y\, x^3\, y^{-1}\, x^6\, t^{(3,0)}, \\
\widehat{v}_2 &= y\, x^6\, y^{-1}\, x^6\, y\, x^{-3}\, y^{-1}\, t^{(3,0)}, \\
\widehat{v}_3 &= y\, x^9\, y^{-1}\, x^6\, y\, x^{-6}\, y^{-1}\, t^{(3,0)}, \\
\widehat{v}_4 &= y\, x^{12}\, y^{-1}\, x^6\, y\, x^{-9}\, y^{-1}\, t^{(3,0)}, \\
\widehat{v}_5 &= y\, x^{15}\, y^{-1}\, x^6\, y\, x^{-12}\, y^{-1}\, t^{(3,0)}, \\
\widehat{v}_6 &= y\, x^{18}\, y^{-1}\, t^{(6,-6)}, \\
\widehat{v}_7 &= x^6\, y\, x^{-12}\, y^{-1}\, t^{(-3,6)},
\end{aligned} \tag{5.83}$$

for the intersection $H_1 \cap H_2'$.

Recall that since $L_1 \cap L_1'$ is trivial, the intersection does not have any nontrivial purely abelian part, and (5.83) is indeed a basis of the whole intersection $H_1 \cap H_1'$. In particular, in this case, the rank of the intersection is 7.

Case 2: Suppose $\mathbf{a} = (3,3)$, $\mathbf{d} = (2,2) \in \mathbb{Z}^2$, and $L_2 = \langle (1,2) \rangle$, $L_2' = \langle (0,0) \rangle \leqslant \mathbb{Z}^2$. That is, consider the subgroups:

$$\begin{aligned} H_2 &= \langle\, t^{(1,2)}, x^3\, t^{(3,3)}, yx\, t^{\mathbf{b}}, y^3 x y^{-2}\, t^{\mathbf{c}} \,\rangle\,, \\ H_2' &= \langle\, x^2\, t^{(2,2)}, yxy^{-1} \,\rangle\,. \end{aligned}$$

Then, $L_2 \cap L_2' = \{(0,0)\}$, $\mathbf{B} = \left[\begin{smallmatrix} 2\mathbf{a}-3\mathbf{d} \\ \mathbf{a} \end{smallmatrix}\right] = \left[\begin{smallmatrix} 0 & 0 \\ 3 & 3 \end{smallmatrix}\right]$; and the set of solutions for

$$\mathbf{w} \begin{bmatrix} 0 & 0 \\ 3 & 3 \end{bmatrix} \in \langle\, (1,2) \,\rangle$$

is clearly the subgroup M generated by the row of the matrix $\mathbf{M} = [1 \;\; 0]$. which is already in Smith normal form. That is, the change of basis matrices are trivial ($\mathbf{P} = \mathbf{I_1} = [1], \mathbf{Q} = \mathbf{I_2} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right]$); and $\mathbf{D} = [1\; 0]$.

Then, according to Proposition 5.10.17:

$$\begin{aligned} \mathsf{St}\,(M\rho^{\leftarrow}, \{w_1, w_2\}) &= \mathsf{Cay}\left(\mathbb{F}_{\{w_1,w_2\}}/M\rho^{\leftarrow}, \{[w_1],[w_2]\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}^2/\langle \mathbf{M} \rangle\,, \{(1,0),(0,1)\}\right) \\ &= \mathsf{Cay}\,(\mathbb{Z}/\mathbb{Z} \oplus \mathbb{Z}/0\mathbb{Z}\,, \{(1,0),(0,1)\}) \\ &= \mathsf{Cay}\,(\mathbb{Z}\,, \{0,1\})\ , \end{aligned}$$

which, denoting by a violet (resp., green) arc the action of the element 0 (resp., 1), takes the form:

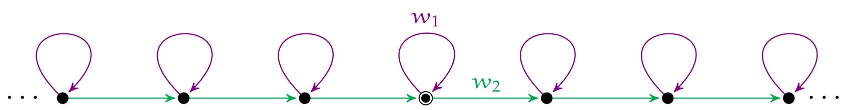

Fig. 5.38: Cayley graph of $\mathbb{Z}$ w.r.t. $\{0,1\}$

which after replacing w_1 by x^6, and w_2 by yx^3y^{-1}, and folding, becomes:

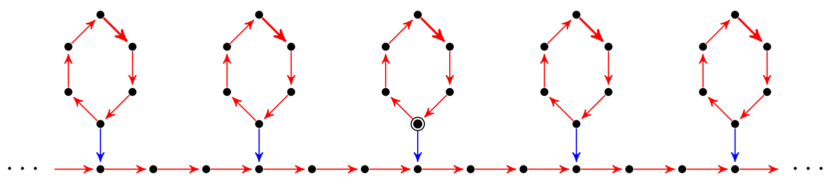

Fig. 5.39: Stallings automaton of $(H_2 \cap H_2')\pi_{\mathbb{F}}$ w.r.t. $\{x,y\}$

Note that it has infinite rank, and provides the basis:

$$v_k = y\,x^{3k}\,y^{-1}\,x^6\,y\,x^{-3k}\,y^{-1}, \quad \forall k \in \mathbb{Z},$$

for $(H_2 \cap H_2')\pi_{\mathbb{F}}$, with the corresponding cyclotomic arcs highlighted (thicker) in Figure 5.39.

According Theorem 5.10.21, if we instead substitute w_1 by $x^6\,t^{\mathbf{u}} = x^6\,t^{(2\mathbf{a},3\mathbf{d})} = x^6\,t^{(6,6),(6,6)}$, and w_2 by $yx^3y^{-1}\,t^{\mathbf{v}} = yx^3y^{-1}\,t^{(\mathbf{a},\mathbf{0})} = yx^3y^{-1}\,t^{(3,3),(0,0)}$, and fold, we obtain:

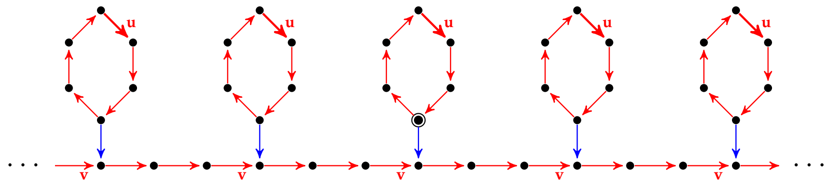

Fig. 5.40: Expanded intersection scheme for $(H_2 \cap H_2')\pi_{\mathbb{F}}$

which is already normalized if we remove all the $\mathbf{v}$'s (lying on bridges, see Lemma 5.8.13) from it. Therefore, taking the vector $(6,6)$ as a common representative in both completions, we obtain the following Stallings automaton:

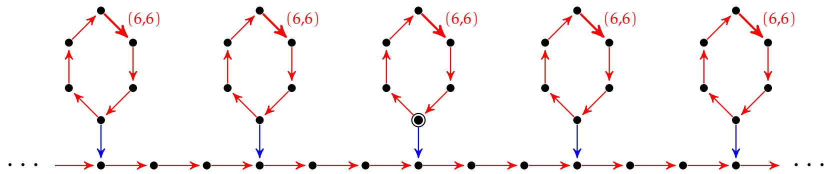

Fig. 5.41: Stallings automaton for $H_2 \cap H_2'$

The corresponding (infinite) basis for $H_2 \cap H_2'$ is:

$$\widehat{v}_k = v_k\,t^{(6,6)} = y\,x^{3k}\,y^{-1}\,x^6\,y\,x^{-3k}\,y^{-1}\,t^{(6,6)}, \quad \forall k \in \mathbb{Z},$$

(again, since $L_2 \cap L_2'$ is trivial, there is no abelian contribution to the basis), and hence the intersection $H_2 \cap H_2'$ has infinite rank.

Remark 5.10.27. Note that the fact that an element in one of the intersecting subgroups does not appear in the intersection does not mean that it does not affect it. For example, the cases 3 and 4 (below) correspond to case 2 after replacing the generator of L_i (from $t^{(1,2)}$ to $t^{(2,2)}$ and $t^{(1,1)}$ respectively). Observe that none of these elements belong to the corresponding intersection, but the intersections are different (even of different type).

Case 3: Suppose $\mathbf{a} = (3,3)$, $\mathbf{d} = (2,2) \in \mathbb{Z}^2$, and $L_3 = \langle (2,2) \rangle$, $L_3' = \langle (0,0) \rangle \leqslant \mathbb{Z}^2$. That is, consider the subgroups:

$$\begin{aligned} H_3 &= \langle\, t^{(2,2)}, x^3\, t^{(3,3)}, yx\, t^{\mathbf{b}}, y^3xy^{-2}\, t^{\mathbf{c}} \,\rangle\,, \\ H_3' &= \langle\, x^2\, t^{(2,2)}, yxy^{-1} \,\rangle\,. \end{aligned}$$

Note that the only difference with case 2 is that $L_3 \neq L_2$. Hence, as before, $\mathbf{B} = \left[\begin{smallmatrix} 0 & 0 \\ 3 & 3 \end{smallmatrix}\right]$, and $L_3 \cap L_3' = \{(0,0)\}$; and the set of solutions for the equation $\mathbf{w} \left[\begin{smallmatrix} 0 & 0 \\ 3 & 3 \end{smallmatrix}\right] \in \langle\, (2,2) \,\rangle$ is the subgroup generated by the rows of the matrix $\mathbf{M} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right]$, which is already in Smith normal form.

Then, according to Proposition 5.10.17:

$$\begin{aligned} \mathsf{St}\,(M\rho^{\leftarrow}, \{w_1, w_2\}) &= \mathsf{Cay}\,\left(\mathbb{F}_{\{w_1,w_2\}}/M\rho^{\leftarrow}, \{[w_1],[w_2]\}\right) \\ &= \mathsf{Cay}\,\left(\mathbb{Z}^2/\langle \mathbf{M} \rangle, \{(1,0),(0,1)\}\right) \\ &= \mathsf{Cay}\,(\mathbb{Z}/\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}\,, \{(1,0),(0,1)\}) \\ &= \mathsf{Cay}\,(\mathbb{Z}/2\mathbb{Z}\,, \{0,1\})\ , \end{aligned}$$

which, denoting by a violet (resp., green) arc the action of the element 0 (resp., 1), takes the form:

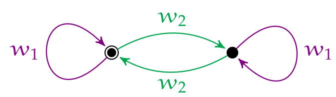

Fig. 5.42: Cayley digraph of $\mathbb{Z}/2\mathbb{Z}$ w.r.t. $\{0,1\}$

which after replacing w_1 by x^6, and w_2 by yx^3y^{-1}, and folding, becomes:

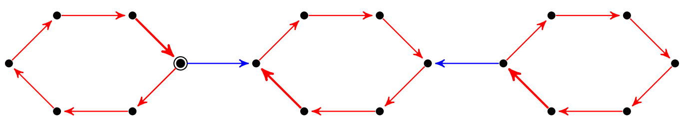

Fig. 5.43: Stallings automaton of $(H_3 \cap H_3')\pi_{\mathbb{F}}$ w.r.t. $\{x,y\}$

which provides the basis:

$$\begin{aligned} v_1 &= x^6\,, \\ v_2 &= y\,x^6\,y^{-1}\,, \\ v_3 &= y\,x^3\,y^{-1}\,x^6\,y\,x^{-3}\,y^{-1}\,. \end{aligned} \tag{5.84}$$

for $(H_3 \cap H_3')\pi_{\mathbb{F}}$, with the corresponding cyclotomic arcs highlighted (thicker) in Figure 5.43.

Now, after replacing w_1 by $x^6\, t^{(2\mathbf{a},3\mathbf{d})} = x^6\, t^{(6,6),(6,6)}$ and w_2 by $yx^3y^{-1}\, t^{(\mathbf{a},\mathbf{0})} = yx^3y^{-1}\, t^{(3,3),(0,0)}$, and fold, we obtain:

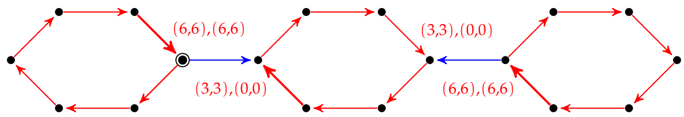

which we can normalize to get:

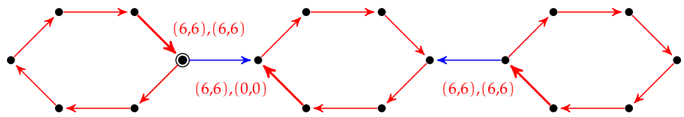

whose double labels we know that are equalizable. Taking a common representative for both abelian completions, for example:

$$(6,6) \in (6,6) + L_3 \cap (6,6) + L_3' = (6,6) + \langle (2,2) \rangle \cap (6,6) + \langle (0,0) \rangle, \text{ and}$$
$$(0,0) \in (6,6) + L_3 \cap (0,0) + L_3' = (6,6) + \langle (2,2) \rangle \cap (0,0) + \langle (0,0) \rangle$$

we finally obtain an Stallings automata for the intersection $H_3 \cap H_3'$:

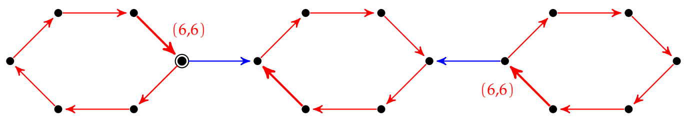

Fig. 5.44: Stallings automaton for $H_3 \cap H_3'$

which provides the basis:

$$\begin{aligned} \widehat{w}_1 &= x^6\, t^{(6,6)}, \\ \widehat{w}_2 &= y\, x^6\, y^{-1}, \\ \widehat{w}_3 &= y\, x^3\, y^{-1}\, x^6\, y\, x^{-3}\, y^{-1}\, t^{(6,6)}. \end{aligned} \tag{5.85}$$

for the subgroup $H_3 \cap H_3'$, which therefore has rank equal to 3.

Case 4: Suppose $\mathbf{a} = (3,3)$, $\mathbf{d} = (2,2) \in \mathbb{Z}^2$, and $L_4 = \langle (1,1) \rangle$, $L_4' = \langle (0,0) \rangle \leqslant \mathbb{Z}^2$. That is, consider the subgroups:

$$\begin{aligned} H_4 &= \langle\, t^{(1,1)}, x^3\, t^{(3,3)}, yx\, t^{\mathbf{b}}, y^3xy^{-2}\, t^{\mathbf{c}} \,\rangle, \\ H_4' &= \langle\, x^2\, t^{(2,2)}, yxy^{-1} \,\rangle. \end{aligned}$$

Again, the only difference with cases 2 and 3 lies in the abelian subgroup L_4. Hence, $L_4 \cap L_4' = \{(0,0)\}$, and $\mathbf{B} = \left[\begin{smallmatrix} 0 & 0 \\ 3 & 3 \end{smallmatrix}\right]$. In this case it is clear that the set of solutions of the equation $\mathbf{w} \left[\begin{smallmatrix} 0 & 0 \\ 3 & 3 \end{smallmatrix}\right] \in \langle\, (1,1)\, \rangle$ is the whole space $\mathbb{Z}^2$, namely the subgroup generated by the rows of the matrix $\mathbf{M} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right]$, which is already in Smith normal form.

Then, according to Proposition 5.10.17:

$$\begin{aligned} \mathsf{St}\left(M\rho^{\leftarrow}, \{w_1, w_2\}\right) &= \mathsf{Cay}\left(\mathbb{F}_{\{w_1,w_2\}}/M\rho^{\leftarrow}, \{[w_1],[w_2]\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}^2/\langle \mathbf{M}\rangle, \{(1,0),(0,1)\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}/\mathbb{Z} \oplus \mathbb{Z}/\mathbb{Z}, \{(1,0),(0,1)\}\right) \\ &= \mathsf{Cay}\left(\{0\}, \{0,0\}\right), \end{aligned}$$

(recall that we admit Cayley digraphs w.r.t. generating multisets) which has the form:

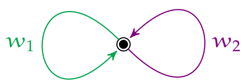

Fig. 5.45: Cayley digraph of the trivial (abelian) group $\{0\}$ w.r.t. $\{0,0\}$

and after replacing w_1 by x^6, and w_2 by yx^3y^{-1}, and folding, becomes:

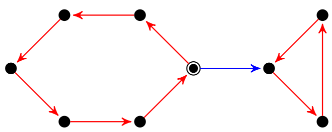

Fig. 5.46: Stallings automaton of $(H_4 \cap H_4')\pi_{\mathbb{F}}$ w.r.t. $\{x, y\}$

That is, in this case the Stallings automata of $(H_4 \cap H_4')\pi_{\mathbb{F}}$ and $H_4\pi_{\mathbb{F}} \cap H_4'\pi_{\mathbb{F}}$ do coincide. Hence $(H_4 \cap H_4')\pi_{\mathbb{F}} = H_4\pi_{\mathbb{F}} \cap H_4'\pi_{\mathbb{F}}$, with basis:

$$\begin{aligned} w_1 &= x^6, \\ w_2 &= y\,x^3\,y^{-1}. \end{aligned} \tag{5.86}$$

After replacing w_1 by $x^6\,t^{(2\mathbf{a},3\mathbf{d})} = x^6\,t^{(6,6),(6,6)}$ and w_2 by $yx^3y^{-1}\,t^{(\mathbf{a},\mathbf{0})} = yx^3y^{-1}\,t^{(3,3),(0,0)}$, and fold, we obtain:

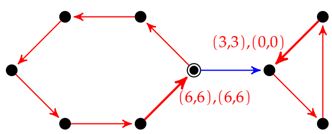

which is already normalized. Hence their double labels are equalizable, again taking $(6,6)$, and $(0,0)$ as representatives; and an Stallings automaton for $H_4 \cap H_4'$ is:

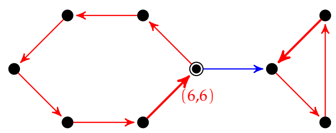

Fig. 5.47: Stallings automaton for $H_4 \cap H_4'$

which provides the basis:

$$\begin{aligned} \widehat{w}_1 &= x^6\, t^{(6,6)}\,, \\ \widehat{w}_2 &= y\, x^3\, y^{-1}\,. \end{aligned} \tag{5.87}$$

for the subgroup $H_4 \cap H_4'$, which therefore has rank equal to 2.

Case 5: Suppose $\mathbf{a} = (1,-1)$, $\mathbf{d} = (1,-2) \in \mathbb{Z}^2$, and $L_5 = \langle (12,0),(0,12)\rangle$, $L_5' = \langle (2,-2)\rangle \leqslant \mathbb{Z}^2$. That is, consider

$$\begin{aligned} H_1 &= \langle\, t^{(12,0)}\,,\, t^{(0,12)}\,,\, x^3\, t^{(1,-1)}\,,\, yx\, t^{\mathbf{b}}\,,\, y^3xy^{-2}\, t^{\mathbf{c}} \rangle\,, \\ H_1' &= \langle t^{(2,-2)}\,,\, x^2\, t^{(1,-2)}\,,\, yxy^{-1} \rangle\,. \end{aligned}$$

Then, $L_5 \cap L_5' = \{(12,-12)\}$, $\mathbf{B} = \left[\begin{smallmatrix} 2\mathbf{a}-3\mathbf{d} \\ \mathbf{a} \end{smallmatrix}\right] = \left[\begin{smallmatrix} -1 & 4 \\ 1 & -1 \end{smallmatrix}\right]$; and the set of solutions for

$$\mathbf{w} \begin{bmatrix} -1 & 4 \\ 1 & -1 \end{bmatrix} \in \langle\, (12,0),(0,12),(2,-2)\,\rangle = \langle\, (12,0),(2,-2)\,\rangle$$

is the subgroup M generated by the rows of the matrix $\mathbf{M} = \left[\begin{smallmatrix} 4 & 16 \\ 0 & 2 \end{smallmatrix}\right]$; with SNF decomposition $\mathbf{PMQ} = \mathbf{D}$, where $\mathbf{P} = \left[\begin{smallmatrix} 0 & 1 \\ 1 & -8 \end{smallmatrix}\right]$, $\mathbf{Q} = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$, and $\mathbf{D} = \left[\begin{smallmatrix} 2 & 0 \\ 0 & 4 \end{smallmatrix}\right]$. Therefore, according to Proposition 5.10.17

$$\begin{aligned} \mathsf{St}\,(M\rho^{\leftarrow},\{w_1,w_2\}) &= \mathsf{Cay}\left(\mathbb{F}_{\{w_1,w_2\}}/M\rho^{\leftarrow},\{[w_1],[w_2]\}\right) \\ &= \mathsf{Cay}\left(\mathbb{Z}^2/\langle\mathbf{D}\rangle\,,\{(0,1),(1,0)\}\right) \\ &= \mathsf{Cay}\,(\mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z}\,,\{(0,1),(1,0)\})\,, \end{aligned}$$

which, denoting by a violet (resp., green) arc the action of the element $(0,1)$ (resp., $(1,0)$), has the form:

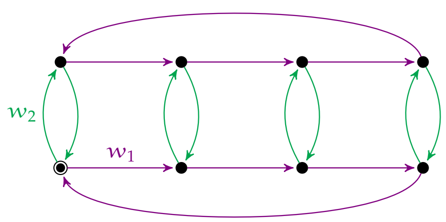

Fig. 5.48: Cayley digraph of $\mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z}$ w.r.t. $\{(0,1),(1,0)\}$

which after replacing w_1 by x^6, and w_2 by yx^3y^{-1}, and folding, becomes:

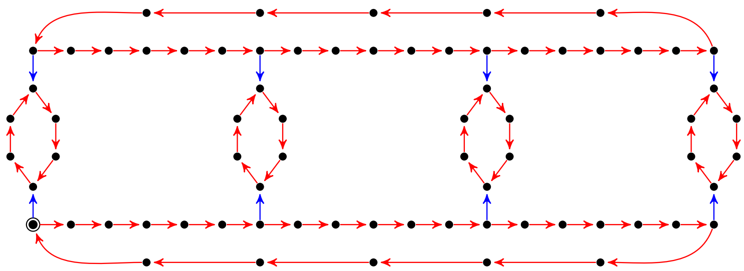

Fig. 5.49: Stallings automaton of $(H_5 \cap H_5')\pi_{\mathbb{F}}$ w.r.t. $\{x, y\}$

After replacing w_1 by $x^6\, t^{\mathbf{u}} = x^6\, t^{(2\mathbf{a},3\mathbf{d})} = x^6\, t^{(2,-2),(3,-6)}$ and w_2 by $yx^3y^{-1}\, t^{\mathbf{v}} = yx^3y^{-1}\, t^{(\mathbf{a},\mathbf{0})} = yx^3y^{-1}\, t^{(1,-1),(0,0)}$, and fold, we obtain:

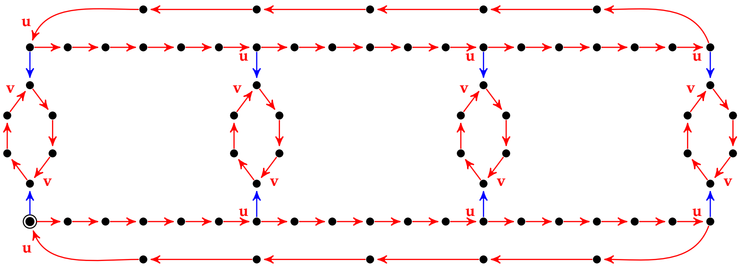

which is still not normalized (note that the double labels in it are not equalizable). However, after choosing an spanning tree and normalizing, we get

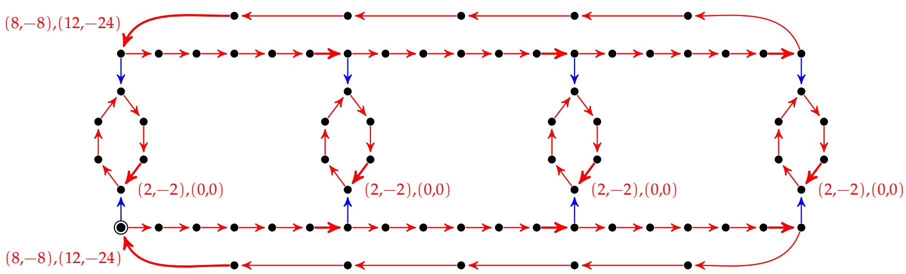

where we can finally equalize labels taking, for example:

$$(-4,-8) \in (8,-8) + \langle (12,0),(0,12) \rangle \cap (12,-24) + \langle (2,-2) \rangle \,, \text{ and}$$

$$(2,-2) \in (2,-2) + \langle (12,0),(0,12) \rangle \cap (0,0) + \langle (2,-2) \rangle \,;$$

and attach $\langle (12,-12) \rangle = L_5 \cap L_5'$ as a basepoint subgroup to obtain an Stallings automaton for the intersection $H_5 \cap H_5'$:

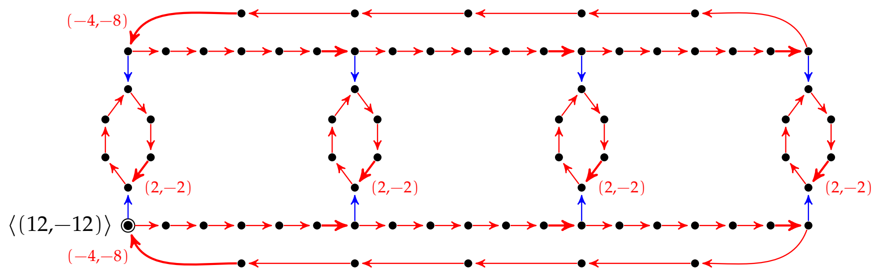

Fig. 5.50: Stallings automaton for $H_5 \cap H_5'$ w.r.t. $\{x, y\}$

which provides the basis:

$$
\begin{aligned}
v_1 &= yx^6y^{-1}\,t^{(2,-2)}\\
v_2 &= x^6yx^6y^{-1}x^{-6}\,t^{(2,-2)}\\
v_3 &= x^{12}yx^6y^{-1}x^{-12}\,t^{(2,-2)}\\
v_4 &= x^{18}yx^6y^{-1}x^{-18}\,t^{(2,-2)}\\
v_5 &= yx^3y^{-1}x^6yx^{-3}y^{-1}x^{-6}\\
v_6 &= yx^3y^{-1}x^{12}yx^{-3}y^{-1}x^{-12}\\
v_7 &= yx^3y^{-1}x^{18}yx^{-3}y^{-1}x^{-18}\\
v_8 &= x^{24}\,t^{(-4,-8)}\\
v_9 &= yx^3y^{-1}x^{24}yx^{-3}y^{-1}\,t^{(-4,-8)}\\
v_{10} &= t^{(12,-12)}
\end{aligned}
$$

for the subgroup $H_5 \cap H_5'$, which therefore has rank equal to 10.

Case 6: Suppose $\mathbf{a} = (6,6)$, $\mathbf{d} = (4,4) \in \mathbb{Z}^2$, $L_6 = \langle (6p, 6p) \rangle$, $L_6' = \langle (0,0) \rangle \leqslant \mathbb{Z}^2$, for some $p \in \mathbb{Z}$. That is, consider the subgroups:

$$
\begin{aligned}
H_6 &= \langle\, t^{(6p,6p)}, x^3\,t^{(6,6)}, yx\,t^{\mathbf{b}}, y^3xy^{-2}\,t^{\mathbf{c}} \,\rangle\,,\\
H_6' &= \langle\, x^2\,t^{(4,4)}\,,\ yxy^{-1} \,\rangle\,.
\end{aligned}
$$

So, $\mathbf{B} = \left[\begin{smallmatrix} 2\mathbf{a}-3\mathbf{d} \\ \mathbf{a} \end{smallmatrix}\right] = \left[\begin{smallmatrix} 0 & 0 \\ 6 & 6 \end{smallmatrix}\right]$, and $L_3 \cap L_3' = \{(0,0)\}$. Then, the set of solutions for $\mathbf{w}\left[\begin{smallmatrix} 0 & 0 \\ 6 & 6 \end{smallmatrix}\right] \in \langle\, (6p, 6p) \,\rangle$ is the subgroup generated by the rows of the matrix $\mathbf{M} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & p \end{smallmatrix}\right]$, which is already in Smith normal form, i.e., $\mathbf{PMQ} = \mathbf{D}$, where $\mathbf{P} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right]$, $\mathbf{Q} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right]$, and $\mathbf{D} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & p \end{smallmatrix}\right]$.

Then, according to Equation (5.66):

$$\begin{aligned}
\mathsf{St}\left(M\rho^{\leftarrow},\{w_1,w_2\}\right) &= \mathsf{Cay}\left(\mathbb{F}_{\{w_1,w_2\}}/M\rho^{\leftarrow},\{[w_1],[w_2]\}\right) \\
&= \mathsf{Cay}\left(\mathbb{Z}^2/\langle \mathbf{M}\rangle\,,\{(1,0),(0,1)\}\right) \\
&= \mathsf{Cay}\left(\mathbb{Z}/\mathbb{Z}\oplus\mathbb{Z}/p\mathbb{Z}\,,\{(1,0),(0,1)\}\right) \\
&= \mathsf{Cay}\left(\mathbb{Z}/p\mathbb{Z}\,,\{0,1\}\right)\,,
\end{aligned}$$

which, denoting by a violet arc the action of the element 0 (corresponding to w_1), and by a green arc the action of the element 1 (corresponding to w_2), takes the form:

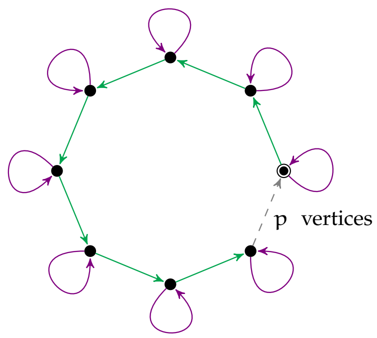

Fig. 5.51: Cayley digraph of $\mathbb{Z}/p\mathbb{Z}$ w.r.t. $\{0,1\}$

which after replacing w_1 by $x^6\,t^{\mathbf{u}} = x^6\,t^{(2\mathbf{a},3\mathbf{d})} = x^6\,t^{(12,12),(12,12)}$, and w_2 by $yx^3y^{-1}\,t^{\mathbf{v}} = yx^3y^{-1}\,t^{(\mathbf{a},\mathbf{0})} = yx^3y^{-1}\,t^{(6,6),(0,0)}$, and folding, becomes:

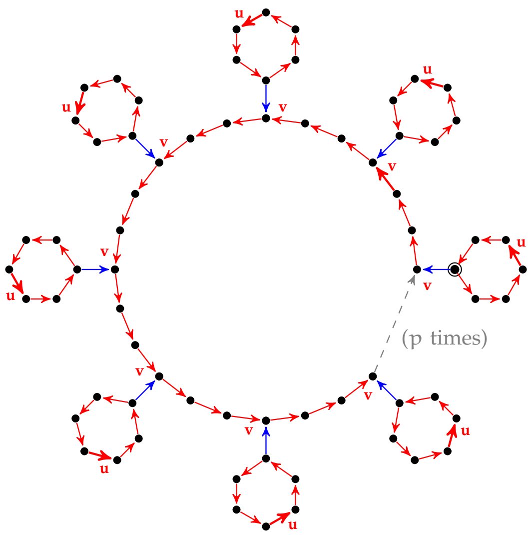

which after normalizing (w.r.t. the spanning tree having as cyclotomic arcs the thicker ones) takes the form:

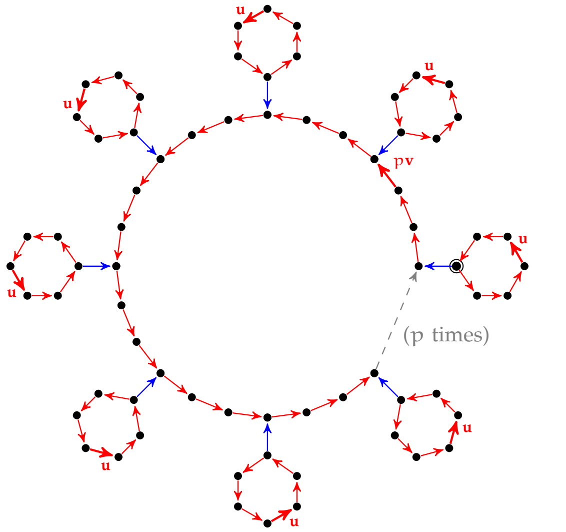

which allow us to equalize the double labelling, taking $(12,12)$ and $(0,0)$ as common representatives for **u** and $p\mathbf{v}$ respectively, and finally obtain an Stallings automaton for the intersection $H_6 \cap H_6'$:

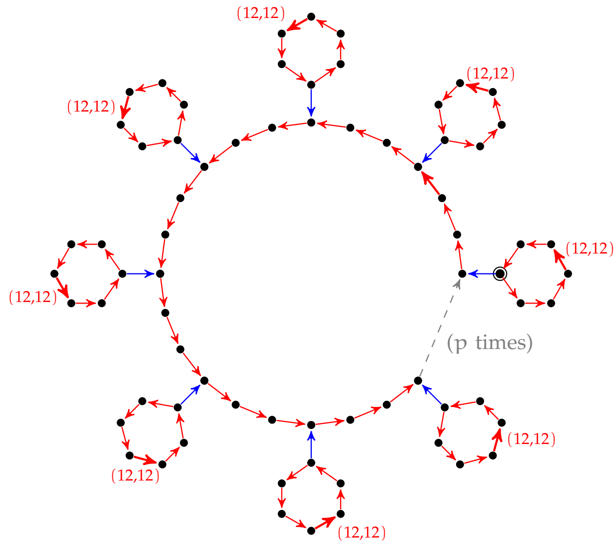

Fig. 5.52: Stallings automaton for $H_6 \cap H_6'$

and provides the basis:

$$\begin{aligned} \widehat{v}_0 &= yx^{3p}y^{-1}\,, \\ \widehat{v}_k &= yx^{3k}y^{-1}x^6yx^{-3k}y^{-1}\,t^{(12,12)}\,, \quad \forall k = 0,\ldots,p-1\,, \end{aligned} \tag{5.88}$$

for the intersection subgroup $H_6 \cap H_6',$, which therefore has rank equal to p.

This last case points out one last interesting consequence of the extended index formula (5.77): not only the intersection of two finitely generated subgroups

$H, H' \leqslant \mathbb{F}_n \times \mathbb{Z}^m$ can be infinitely generated (i.e., FATF groups are not Howson), but *even when it is finitely generated*, one can no longer bound the rank of $H \cap H'$ in terms of the ranks of the intersecting subgroups (moreover, as we see in the example, this can happen even keeping fixed the skeleton of the subgroups).

Proposition 5.10.28. *If for every pair of natural numbers* $n_1, n_2 \in \mathbb{N}^2$ *we define the* threshold function ζ *as:*

$$\zeta(n_1, n_2) \coloneqq \sup\{\, \mathrm{rk}(H_1 \cap H_2) < \infty : \ \mathrm{rk}(H_1) \leqslant n_1 \,,\, \mathrm{rk}(H_2) \leqslant n_2 \,\} \in \mathbb{N} \cup \{\infty\} \,,$$

then ζ *admits only the values* $0, 1, 2$ *and* ∞ *in FATF gropus. Namely:*

(a) *if* $\min\{n_1, n_2\} = 0$, *then* $\zeta(n_1, n_2) = 0$;

(b) *if* $\min\{n_1, n_2\} = 1$, *then* $\zeta(n_1, n_2) = 1$;

(c) *if* $n_1 = n_2 = 2$, *then* $\zeta(n_1, n_2) = 2$;

(d) *otherwise,* $\zeta(n_1, n_2) = \infty$.

Proof. Firstly note that, since we can always take $H_1 = H_2$ of rank equal to $\min\{n_1, n_2\}$, then for all $n_1, n_2 \in \mathbb{Z}^2$, we have that $\zeta(n_1, n_2) \geqslant \min\{n_1, n_2\}$. Now, we study the cases separately.

The case (a) is trivial.

The case (b) is also obvious, since the intersection of two cyclic subgroups must be cyclic, and we can take $H_1 = H_2$.

In case (c) (we can assume $\mathrm{rk}(H_1) = \mathrm{rk}(H_2) = 2$), we distinguish two subcases:

- *if either* $H_1\pi_{\mathbb{F}}$ *or* $H_2\pi_{\mathbb{F}}$ *has rank* 1, then both $H_1\pi_{\mathbb{F}} \cap H_2\pi_{\mathbb{F}}$ (and any of its subgroups, including $(H_1 \cap H_2)\pi_{\mathbb{F}}$) and $L_1 \cap L_2$ have rank at most 1. Therefore, from Corollary 5.6.13, $\mathrm{rk}(H_2 \cap H_2) \leqslant 2$.

- if $\mathrm{rk}(H_1\pi_{\mathbb{F}}) = \mathrm{rk}(H_2\pi_{\mathbb{F}}) = 2$, then $L = L' = L \cap L' = L + L' = \{\mathbf{0}\}$, and so M is the set of solutions of the equation $\mathbf{w}(B) = \mathbf{0}$, i.e., $M = \ker \mathbf{B} \leqslant \mathbb{Z}^m$. Now, from Proposition 5.10.20, the intersection $H_1 \cap H_2$ is finitely generated if and only if $M = \ker \mathbf{B}$ has finite index in $\mathbb{Z}^m$, that is if and only if $\mathbf{B} = \mathbf{0}$ (and the finite index is 1). Note that in this case $(H_1 \cap H_2)\pi_{\mathbb{F}} = H_1\pi_{\mathbb{F}} \cap H_2\pi_{\mathbb{F}}$ and hence:

$$\begin{aligned} \mathrm{rk}(H_1 \cap H_2) &= \mathrm{rk}((H_1 \cap H_2)\pi_{\mathbb{F}}) \\ &= \mathrm{rk}(H_1\pi_{\mathbb{F}} \cap H_2\pi_{\mathbb{F}}) \\ &\leqslant 1 + (\mathrm{rk}\, H_1 - 1)(\mathrm{rk}\, H_2 - 1) \;=\; 2\,, \end{aligned}$$

where we have applied the Hanna Neumann Theorem in the last inequality.

Again, since we can take $H_1 = H_2$, this proves that $\zeta(2,2) = 2$.

To prove (d) it is enough to adapt the case 6 in Example 5.10.25.

Concretely, according case 6, if we take $H_1 = \langle\, t^{(6p,6p)}, x^3\, t^{(6,6)}, yx \,\rangle$ (recall that, as we discussed in Example 5.5.30 the element y^3xy^{-2} plays no role in the intersection) and $H_2 = \langle\, x^2\, t^{(4,4)}\,,\, yxy^{-1} \,\rangle$, then we have $\operatorname{rk} H_1 = 3$, $\operatorname{rk} H_2 = 2$, and $\operatorname{rk}(H_1 \cap H_2) = p$; which proves that $\zeta(2,3) = \infty$.

To prove that $\zeta(3+n_1, 2+n_2) = \infty$, for all $n_1, n_2 \in \mathbb{N}$, it is enough to perform the following trick: consider $\{u_1, \ldots, u_{n_1}\}$, and $\{v_1, \ldots, v_{n_2}\}$ basis of respective subgroups of $\mathbb{F}_X$ such that

- $\langle u_1, \ldots, u_{n_1} \rangle \cap H_1 = \{1\}$,
- $\langle v_1, \ldots, v_{n_2} \rangle \cap H_2 = \{1\}$, and
- $\langle u_1, \ldots, u_{n_1} \rangle \cap \langle v_1, \ldots, v_{n_2} \rangle = \{1\}$,

and then take:

$$\begin{aligned} \widetilde{H}_1 &= \langle\, t^{(6p,6p)}, x^3\, t^{(6,6)}, yx, u_1, \ldots, u_{n_1} \,\rangle \\ \widetilde{H}_2 &= \langle\, x^2\, t^{(4,4)}\,,\, yxy^{-1}, v_1, \ldots, v_{n_2} \,\rangle\,. \end{aligned}$$

Now, it is clear by construction that $\operatorname{rk} \widetilde{H}_1 = 3 + n_1$, $\operatorname{rk} \widetilde{H}_2 = 2 + n_2$, and the rank of the intersection is $\operatorname{rk}(\widetilde{H}_1 \cap \widetilde{H}_2) = \operatorname{rk}(H_1 \cap H_2) = p$.

The claim $\zeta(3+n_1, 2+n_2) = \infty$ follows, and the proof is completed. □

This fact is relevant because it denies any possible extension of the recently proved Hanna Neumann conjecture (see Theorem 5.5.21) in this direction, thus providing clear limitations in the attempt to extend that celebrated result to broader families, which has become a very active research target since the proof of the original conjecture (see for example [DI10; AMS11; Zak14; ASS14]).

A topological approach

In this section, we present an alternative (more topological) description of the Stallings automaton for the intersection of two subgroups $H, H' \leqslant \mathbb{F}_n \times \mathbb{Z}^m$.

Theorem 5.10.29. *Let H, H' be subgroups of $\mathbb{F}_X \times \mathbb{Z}^m$. Then, a normalized Stallings automaton for the intersection $H \cap H'$ is the result of the following steps:*

1. *normalize the junction automaton $\widehat{\mathrm{St}}(H\pi_{\mathbb{F}}) \wedge \widehat{\mathrm{St}}(H'\pi_{\mathbb{F}})$ w.r.t. a certain distinguished spanning tree* $\mathbf{T}$*;*

2. *replace every label* w_i *— of the arc* e_i *in* $\mathsf{St}(M\rho^{\leftarrow}, \{w_i\}_i)$ *— with the (doubly-enriched) label of the corresponding* $\mathbf{T}$*-cyclotomic arc in the junction* $\widehat{\mathsf{St}}(H\pi_{\mathbb{F}}) \wedge \widehat{\mathsf{St}}(H'\pi_{\mathbb{F}})$*;*

3. *replace every vertex* p_j *in* $\mathsf{St}(M\rho^{\leftarrow}, \{w_i\}_i)$ *by a copy* $\mathbf{T}^{(j)}$ *of* $\mathbf{T}$*;*

4. *for every (now enriched* $X^{\pm}$*-labelled) arc* $e_i \equiv p_{j_1} \longrightarrow p_{j_2}$ *in* $\mathsf{St}(M\rho^{\leftarrow}, \{w_i\}_i)$*, join its tail and head with the leaves in* $\mathbf{T}^{(j_1)}$ *and* $\mathbf{T}^{(j_2)}$ *respectively, in such a way that the path from the basepoint in* $\mathbf{T}^{(j_1)}$ *to the basepoint in* $\mathbf{T}^{(j_2)}$ *reads precisely the word* $\widehat{w}_i$ *recognized by the* i*-th petal in* $\widehat{\mathsf{St}}(H\pi_{\mathbb{F}}) \wedge \widehat{\mathsf{St}}(H'\pi_{\mathbb{F}})$*;*

5. *normalize the resulting automaton, and equalize its abelian labels;*

6. *set as basepoint for the resulting automaton, the basepoint of the basepoint-tree* $\mathbf{T}^{(\circledast)}$*, and attach to it* $L \cap L'$ *as a basepoint subgroup.*

Proof. As in Theorem 5.9.12, it is clear by construction that the resulting enriched automaton is inverse, normalized, and recognizes exactly the subgroup $H \cap H'$. □

Part III

Partially commutative groups

Definition and generalities

6

In this part of the dissertation, we care about the second main generalization of free-abelian times free groups considered in Part I, namely that of *partially commutative groups* (PC-groups). In particular, we combine a generalization of the techniques developed in Section 2.3, with an algorithmic version of the Stallings-like automata theory for free products developed by Ivanov, to study several algorithmic intersection problems within the subfamily of *Droms groups* (those finitely generated PC-groups whose subgroups are again PC-groups).

6.1 Partially-commutative groups

Given any set of elements $U = \{u_i\}_i$ in a group G, one can always consider the (undirected simple) *commutation graph* of U in G, denoted by $\Gamma_{U,G}$, having as vertices the elements in U, and two different vertices u_i, u_j being adjacent if and only if they commute in G (i.e., if $[u_i, u_j] = 1$).

We call *partially commutative groups* (*PC-groups*, for short) the groups that admit presentations all whose relations are commutators between generators, i.e., presentations of the form $\langle X \mid R \rangle$, where R is a subset of $[X, X]$ (the *set* of commutators between elements in X).

We can represent this situation in a very natural way through the commutation graph $\Gamma = (X, E)$ having as vertices the generators in X, and two vertices $x, y \in X$ being adjacent if and only if its commutator $[x, y]$ belongs to R; then we say that the PC-group is *presented* by the graph Γ, and we denote it by $\mathbb{G}_\Gamma$.

Recall that a simple graph is undirected, and containing no loops or multiple edges; so, a simple graph $\Gamma = (X, E)$ is nothing more than a symmetric and irreflexive binary relation in X, and its edges can be represented as 2-subsets of X, that is $E \subseteq \binom{X}{2}$.

Definition 6.1.1. A group G is said to be *(free) partially commutative* (a PC-*group*, for short) if it admits the presentation

$$\mathbb{G}_\Gamma = \langle X \mid [x_1, x_2] = 1, \text{ whenever } \{x_1, x_2\} \in E \rangle , \tag{6.1}$$

for some (not necessarily finite) simple graph $\Gamma = (X, E)$ (i.e., if there exists an isomorphism $\varphi \colon \mathbb{G}_\Gamma \to G$). Then, we say that G is presented by the *commutation* *graph* Γ; that (6.1) is a *graphic presentation* for G; and — for each such φ — that $X\varphi$ is a *graphic set of generators* (a *basis*, for short) for G.

Below, we state an useful reformulation of the previous concepts.

Lemma 6.1.2. *Let* $U = \{u_j\}_j$ *be a subset of an arbitrary group* G. *Then, the following statements are equivalent:*

(a) U *is a (graphic) basis of* $\langle U \rangle_G$.

(b) *The natural epimorphism* $\mathbb{G}_{\Gamma_{U,G}} \twoheadrightarrow \langle U \rangle_G$ $(u_j \mapsto u_j)$ *is injective.*

Proof. The implication (b) $\Rightarrow$ (a) is obvious. For the converse, suppose that there exists a simple graph Γ such that $\varphi \colon \langle U \rangle_G \to \mathbb{G}_\Gamma$ is a group isomorphism. □

Partially commutative groups constitute a surprisingly rich family, and have been thoroughly studied — specially in the finitely generated case — during the last decades in different branches of mathematics and computer science. They appear in the literature also as *semi-free groups*, *graph groups*, or *right-angled Artin groups* (when they are finitely generated), among other names. Below, we recall some results about PC-groups we will need throughout the paper; we refer the reader to [Cha07; EKR05; Gre90; Kob13] for more detailed surveys, and further reference.

It is obvious that every simple graph Γ presents exactly one PC-group; more precisely, we have a surjective map $\Gamma \mapsto \mathbb{G}_\Gamma$ between (isomorphic classes of) simple graphs, and (isomorphic classes of) PC-groups.

Remark 6.1.3. The abelianization of a PC-group $\mathbb{G}_\Gamma$ is always the free-abelian group with rank equal to the number of vertices in Γ (and hence to the rank of $\mathbb{G}_\Gamma$). Namely, if $\Gamma = (X, E)$, then

$$(\mathbb{G}_\Gamma)^{\mathsf{ab}} = \bigoplus\nolimits_{x \in X} \mathbb{Z},$$

and thus,

$$\operatorname{rk} \mathbb{G}_\Gamma = \operatorname{rk} \mathbb{G}_\Gamma^{\mathsf{ab}} = |X|, \tag{6.2}$$

with the canonical basis consisting of the (classes of) vertices in Γ. The abelianization map is hence given by $w(X) \mapsto \mathbf{w} = (|w|_i)_i$, where $|w|_i$ denotes the total x_i-exponent in w (exactly in the same way as the abelianization works for free groups). In particular, the abelianization of any PC-group of finite rank m, is the finitely generated free-abelian group $\mathbb{Z}^m$.

So, it is clear that every pair of graphs presenting the same PC-group must have the same order (equal to the rank of the group). This, however, turns out to be a very weak statement since a key result proved by Droms in [Dro87b] states that the map $\Gamma \mapsto \mathbb{G}_\Gamma$ is indeed bijective.

Theorem 6.1.4 (Droms, 1987, [Dro87b]). *Let Γ_1, Γ_2 be simple graphs. Then, the groups $\mathbb{G}_{\Gamma_1}, \mathbb{G}_{\Gamma_2}$ are isomorphic if and only if the graphs Γ_1, Γ_2 are isomorphic, i.e.,*

$$\mathbb{G}_{\Gamma_1} \simeq \mathbb{G}_{\Gamma_2} \text{ (as groups) } \Leftrightarrow \Gamma_1 \simeq \Gamma_2 \text{ (as graphs)}. \qquad \square$$

Corollary 6.1.5. *For any PC-group $\mathbb{G}_\Gamma$, the following conditions are equivalent:* (a) *Γ is finite,* (b) *$\mathbb{G}_\Gamma$ is finitely presented,* (c) *$\mathbb{G}_\Gamma$ is finitely generated.* □

Definition 6.1.6. Finitely generated PC-groups are also called *right-angled Artin groups* (*RAAGs*, for short); we will use this last denomination when we want to emphasize the finitely generated character of a PC-group; whereas the term *PC-group* will refer to a general — possibly infinitely generated — partially commutative group.

Remark 6.1.7. Note that any local property (i.e., involving only finitely many vertices) holding for RAAGS, also holds for general PC-groups. For example an infinitely generated PC-group is Howson if and only if every emmbedded RAAG is Howson (see Section 6.2.1).

However, many important algebraic properties of RAAGs do no not generalize to infinitely generated PC-groups. Specially important for us is the case of Hopfianity, which, in the finitely generated case, can be deduced from residual finiteness (see [Kob13]).

Proposition 6.1.8. *RAAGs are Hopfian.* □

Of course infinitely generated PC-groups are not Hopfian in general (consider, for example, any infinitely generated free group $\mathbb{F}_X$, where any infinite proper subset of vertices $S \subseteq X$ *of the same cardinality* as X generate a proper quotient isomorphic to the starting group $\mathbb{F}_X$).

From Theorem 6.1.4, we have an absolutely transparent geometric characterization of isomorphic classes of PC-groups (resp., RAAGs): we can identify them with simple (resp., finite simple) graphs . Hence, several graph-related notions will be relevant to our discussion.

Examples 6.1.9. Some families of graphs having protagonism throughout the chapter are the following:

- *Complete graphs* (K_n): graphs with n vertices and all possible edges between them.

- *Edgeless graphs* (K_n^c): graphs with n vertices and no edges.
- *Path graphs* (P_n): graphs with n vertices $\{1,\ldots,n\}$, and every vertex except the last one, adjacent to the next one.
- *Cycle graphs* (C_n): graphs with n vertices $\{1,\ldots,n\}$, and every vertex except adjacent to the next one modulo n.

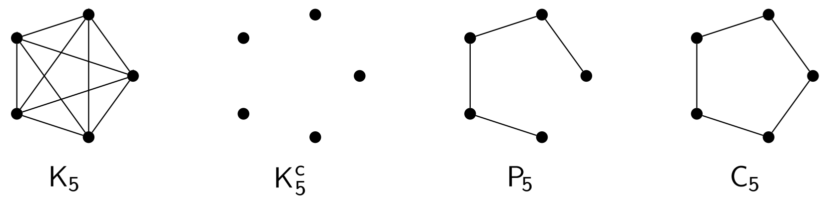

Fig. 6.1: Complete, edgeless, path, and cycle graphs of order 5

Definition 6.1.10. A *central vertex* (a.k.a. universal, or dominating vertex) is a vertex adjacent to any other vertex in the graph. The *center* of a graph Γ, denoted by $Z(\Gamma)$, is the set of central vertices in Γ. A graph having (at least) one central vertex is called a *cone*.

If there is no possible confusion, we will usually abuse notation and also call the center of Γ, and denote it by $Z(\Gamma)$, the subgraph induced in Γ by the set of central vertices.

Lemma 6.1.11. *The center of a graph is the intersection of all its maximal complete subgraphs.*

Proof. Let p be a vertex lying in the center of a graph Γ, and let K be a maximal complete subgraph of Γ; then p must belong to K, because otherwise K would be strictly included in the complete graph induced by $VK \cup \{p\}$, contradicting the maximality of K.

For the converse implication, suppose that p is a vertex belonging to every maximal complete subgraph of Γ. Then, for every other vertex $q \neq p$ in Γ (since every vertex in a graph belongs to some maximal complete subgraph), p belongs to some maximal complete subgraph of Γ containing q, and in particular, p is adjacent to q, as we wanted to prove. □

Remark 6.1.12. Note that the center of a graph Γ must be a complete graph contained in any (but not necessarily equal to any) largest complete subgraph in Γ.

A subgraph Δ of a graph $\Gamma = (X, E)$ is said to be *full* (or induced) if it has exactly the edges that appear in Γ over its vertex set, say $Y \subseteq X$. Then, Δ is called the *full subgraph of* Γ *spanned by* Y, and we denote it by $\Delta = \Gamma[Y]$. If Γ has a full subgraph isomorphic to a certain graph Δ, we will abuse the terminology and say that Δ is (or appears as) a full subgraph of Γ; we denote this situation by $\Delta \leqslant \Gamma$. When none of the graphs belonging to a certain family $\mathcal{F}$ appear as a full subgraph of Γ, we say that Γ is *$\mathcal{F}$-free*. In particular, a graph Γ is *Δ-free* if it does not have any full subgraph isomorphic to Δ.

Remark 6.1.13. Every subgraph of a $\mathcal{F}$-free graph is again $\mathcal{F}$-free.

Definition 6.1.14. A graph is said to be *chordal* if it contains no induced cycles of length stricly greater than 3; i.e., if it is $\{C_n, n \geqslant 4\}$-free.

In view of Theorem 6.1.4, it is natural to look for correspondences between algebraic properties (of PC-groups), and geometric properties (of the graphs presenting them). It turns out that several algebraic properties can be expressed in terms of forbidden subgraphs. The most relevant to our discussion will be:

- P_3-*free graphs*: these are disjoint unions of complete graphs (a.k.a. *cluster graphs*), corresponding to Howson PC-groups; see Theorem 6.2.2.
- *Finite chordal graphs*, corresponding to coherent RAAGs; see Theorem 6.1.24.
- *Finite* $\{P_4, C_4\}$-*free graphs* (i.e., *Droms graphs*) corresponding to Droms RAAGs; see Theorem 6.1.21.

Also, two types of operations between graphs will represent a prominent role throughout the chapter; namely, disjoint union, and graph join. Recall that the *disjoint union* of a family of graphs $\{\Gamma_i = (X_i, E_i)\}_i$ is the graph $\bigsqcup_i \Gamma_i := (\bigsqcup_i X_i, \bigsqcup_i E_i)$; whereas the *join* of two graphs Γ and Γ', denoted by $\Gamma \vee \Gamma'$, is the graph obtained by adding to $\Gamma \sqcup \Gamma'$ every edge joining a vertex in Γ to a vertex in Γ'.

For example, the PC-groups associated to the complete graph K_n, and its complementary (the edgeless graph) K_n^c, are respectively the free-abelian group $\mathbb{Z}^n$, and the free group $\mathbb{F}_n$. We can think of PC-groups as a generalization of these two extreme cases describing the intermediate commutativity situations between them.

So, for example, the finitely generated free-abelian times free group $\mathbb{Z}^m \times \mathbb{F}_n$ studied in Part I is presented by the join $K_m \vee K_n^c$ of a complete graph of order m, and an edgeless graph of order n.

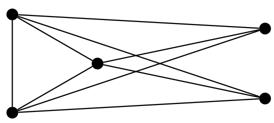

Fig. 6.2: The graph $K_3 \vee K_2^c$ presenting the group $\mathbb{Z}^3 \times \mathbb{F}_2$

All these facts are direct from definitions, and make the equivalence between the conditions in Lemma 6.1.15 almost immediate as well.

Lemma 6.1.15. *Let Γ be an arbitrary simple graph, and $\mathbb{G}_\Gamma$ the corresponding PC-group. Then, the following conditions are equivalent:*

(a) *The path on three vertices P_3 is not a full subgraph of Γ (i.e., Γ is P_3-free).*

(b) *The reflexive closure of Γ is a transitive binary relation.*

(c) *The graph Γ is a disjoint union of complete graphs.*

(d) *The group $\mathbb{G}_\Gamma$ is a free product of free-abelian groups.* □

More generally, disjoint unions and joins of graphs correspond to free products, and *weak* direct products of PC-groups respectively.

Proposition 6.1.16. *A PC-group $\mathbb{G}_\Gamma$ splits as a non-trivial free product if and only if its defining graph Γ is disconnected.*

Proof. The "if" part is immediate from the graph presentation: if $\Gamma = A \sqcup B$ with $A, B \neq \varnothing$, then $\mathbb{G}_\Gamma = \mathbb{G}_A * \mathbb{G}_B$ with $\mathbb{G}_A, \mathbb{G}_B \neq \{1\}$.

For the converse, suppose a connected graph Γ with a nontrivial decomposition $\mathbb{G}_\Gamma = A * B$. It is clear that $\Gamma \neq K_1$, so we can consider two adjacent vertices $p \sim q$ in Γ. Now, since $\langle p, q \rangle \simeq \mathbb{G}_{K_2} = \mathbb{Z}^2$ (which contains no non-abelian free subgroups), the Kurosh subgroup theorem implies that $\langle p, q \rangle$ is contained in certain conjugate of A or B. And, since Γ is connected, it follows that every vertex in Γ is conjugate to an element in one of the free factors, say A. But this is not possible because after abelianizing we would get $\operatorname{rk} G = \operatorname{rk} G^{\,ab} = \operatorname{rk} A^{\,ab} \leqslant \operatorname{rk} A < \operatorname{rk} A + \operatorname{rk} B$ (in contradiction with the Grushko-Neumann theorem). □

On the other hand, the corresponding relation between joins of graphs and direct products of PC-groups is an immediate corollary of the description of centralizers in PC-groups (see [Ser89] and [BC12]).

Theorem 6.1.17. *Let $1 \neq w \in \mathbb{G}_\Gamma$ be a cyclically reduced word. Then, the following statements are equivalent:*

(a) *The element w is contained in a join subgroup.*

(b) *The centralizer of w in $\mathbb{G}_\Gamma$ is not cyclic.*

(c) *The centralizer of w in $\mathbb{G}_\Gamma$ is contained in a join subgroup.*

(Recall that a word $w \in \mathbb{G}_\Gamma$ is called cyclically reduced *if it has minimal length among all reduced words given by cyclically permuting the letters of w.)* □

Corollary 6.1.18. *A PC-group $\mathbb{G}_\Gamma$ splits as a non-trivial direct product if and only if its defining graph Γ is a join.* $\square$

Corollary 6.1.19. *The center of a PC-group $\mathbb{G}_\Gamma$ is the (free-abelian) subgroup generated by the set of central vertices in Γ, i.e.,*

$$Z(\mathbb{G}_\Gamma) = \langle Z(\Gamma)\rangle \simeq \bigoplus_{p \in Z(\Gamma)} \mathbb{Z}. \qquad \square$$

Regarding subgroups, the first natural result is that those generated by subsets of vertices behave as one might expect: they are again PC-groups, and precisely presented by the corresponding induced subgraph. Below, we provide an elementary proof for this well known fact, which will be used later.

Lemma 6.1.20. *Let Γ be an arbitrary simple graph, and Y a subset of vertices of Γ. Then, the subgroup of $\mathbb{G}_\Gamma$ generated by Y is a PC-group presented by $\Gamma[Y]$.*

Proof. Let X be the set of vertices of Γ (then $Y \subseteq X$), and consider the following two homomorphisms:

$$\begin{array}{rcl} \mathbb{G}_{\Gamma[Y]} & \xrightarrow{\iota} & \mathbb{G}_\Gamma \\ y & \longmapsto & y \end{array} \quad , \qquad \begin{array}{rcl} \mathbb{G}_\Gamma & \xrightarrow{\rho} & \mathbb{G}_{\Gamma[Y]} \\ Y \ni y & \longmapsto & y \\ X \setminus Y \ni x & \longmapsto & 1 \end{array} .$$

It is clear that both ι and ρ are well defined homomorphisms (they obviously respect relations). Moreover, note that the composition $\iota\rho$ (ι followed by ρ) is the identity map on $\mathbb{G}_{\Gamma[Y]}$. Therefore, ι is a monomorphism, and thus $\mathbb{G}_{\Gamma[Y]}$ is isomorphic to its image under ι, which is exactly the subgroup of $\mathbb{G}_\Gamma$ generated by Y, as we wanted to prove. $\square$

Despite the naive appearance of the last result, it soon becomes clear that the full lattice of subgroups of a PC-group $\mathbb{G}_\Gamma$ can be rather complicated (even when Γ is finite). To start with, it turns out that not every subgroup of a PC-group is again a PC-group; as Droms proved in [Dro87c] through the neat result below.

Theorem 6.1.21 (Droms, 1987, [Dro87c])**.** *Let Γ be a finite graph. Then, every subgroup of $\mathbb{G}_\Gamma$ is again a (possibly infinitely generated) PC-group if and only if Γ is $\{P_4, C_4\}$-free.* $\square$

Since these families of groups and graphs are the main protagonists in our discussion, it is convenient to have a common name for them.

Definition 6.1.22. A *Droms graph* is a finite $\{P_4, C_4\}$-free graph. Accordingly, a *Droms group* is a PC-group (finitely) presented by a Droms graph.

So, Droms PC-groups, are precisely those *finitely generated* PC-groups $\mathbb{G}_\Gamma$ having all their subgroups again partially commutative (that is presented by some graph, *not necessarily a subgraph of* Γ). Indeed, a bit more can be deduced from the definition.

Corollary 6.1.23. *Every subgroup of a Droms group is again a (possibly infinitely generated) PC-group with* $\{P_4, C_4\}$*-free commutation graph. In particular, every finitely generated subgroup of a Droms group is again a Droms group.*

Proof. Let H be a subgroup of a Droms group $\mathbb{G}_\Gamma$. From theorem 6.1.21, we have that $H \simeq \mathbb{G}_\Delta$, where Δ is a possibly infinite graph. Now, again by theorem 6.1.21, if a graph $\Lambda \in \{P_4, C_4\}$ were a full subgraph of Δ, then there would exist a non-PC-group K, such that $K \leqslant \mathbb{G}_\Lambda \leqslant \mathbb{G}_\Delta \leqslant \mathbb{G}_\Gamma$, in contradiction with Γ being Droms. Finally, according corollary 6.1.5 the subgroup being finitely generated is equivalent to the defining graph being finite, and the final remark follows. □

From Corollary 6.1.5, it is clear that Droms groups are *coherent* (every finitely generated subgroup is finitely presented). However, this last class was proved to be bigger, again by Droms, who provided in [Dro87a] a characterization also in terms of forbidden graphs.

Theorem 6.1.24 (Droms, 1987, [Dro87a])**.** *Let* Γ *be a finite graph. Then, the PC-group* $\mathbb{G}_\Gamma$ *is coherent if and only if the graph* Γ *is chordal.* □

We remark the pertinacious absence of $\mathbb{G}_{C_4} = \mathbb{F}_2 \times \mathbb{F}_2$ from any collection of 'well-behaved' partially commutative groups (see also Mikhailova's Theorem 6.1.30 in the next section). Yet another argument in the same direction was given by Baumslag and Roseblade in [BR84], where they prove that there exist uncountably many nonisomorphic subgroups of $\mathbb{G}_{C_4} = \mathbb{F}_2 \times \mathbb{F}_2$ (see [DSS92] for more examples of PC-groups having this property).

6.1.1 First algorithmic properties

When considering algorithmic properties about partially commutative groups, we will often — depending on whether the studied property admits a reasonable algorithmic description — restrict ourselves to the finitely generated ones (that we will call RAAGs).

Notation 6.1.25. If there exists an algorithm that solves the problem $\mathsf{PROB}(G)$, then we say that the group G satisfies (property) PROB, or that property $\mathsf{PROB}(G)$ is satisfied. We extend this convention to families of groups; namely if $\mathsf{PROB}(G)$ is solvable for every group in a certain family $\mathcal{F}$, we say that $\mathcal{F}$ satisfies PROB, or that $\mathsf{PROB}(\mathcal{F})$ is satisfied.

For example, it is well known that free groups satisfy WP (*word problem*) and CP (*conjugacy problem*); and, as we will see, PC-groups do not satisfy MP (*subgroup membership problem*) since, for example $\mathsf{MP}(\mathbb{F}_2 \times \mathbb{F}_2) = \mathsf{MP}(\mathbb{G}_{C_4})$ is undecidable.

Remark 6.1.26. Note that, for RAAGs, one can always algorithmically obtain a graphic presentation from any given finite presentation, by brute force: just keep exploring the tree of all possible Tietze transformations applied to the initial (finite) presentation until getting one in the form of a PC-group (namely, with all relators being commutators or certain pairs of generators); this will be achieved in finite time because we know in advance that H is indeed a PC-group.

So, without lost of generality we can always assume (finite) graphic presentations for RAAGs.

For this family of groups — as often happens — the solvability of the *word problem* (WP) relies in the existence of "good normal forms" for its elements. Several different variants of normal forms for RAAGs have been introduced ([Die87; DK93; EKR05]) since Baudisch provided the first ones in [Bau77], allowing him to also give the following useful result.

Proposition 6.1.27 (Baudisch, 1981, [Bau81])**.** *RAAGs are torsion-free.* □

Accordingly, several solutions to the word problem have been found for RAAGs, some of them related to broader contexts — such as partially commutative monoids [Wra88; Wra89; LWZ90], or automatic groups [HM95; Van94] — and including also a (linear-time) solution for the *conjugacy problem* (CP). See a nice survey of both problems in [CGW09].

On the other hand, Theorem 6.1.4 (together with Remark 6.1.26) reduces the isomorphism problem (IP) for RAAGs to the graph isomorphism problem (which is obviously solvable since the involved graphs are always finite). We summarize these classic results in a single statement.

Theorem 6.1.28. *All three Dehn problems* (WP, CP *and* IP) *are solvable for RAAGs.* □

However, the algorithmic benignity of the properties stated so far does not always keep for problems that involve subgroups. For example, in [Mik58], Mikhailova uses a clever trick to translate the word problem of a finitely presented group into the membership problem of certain subgroup of $\mathbb{F}_2 \times \mathbb{F}_2$ (i.e., the RAAG presented by the square graph), to then — taking advantage of the existence of finitely presented groups with unsolvable word problem — prove that $\mathbb{F}_2 \times \mathbb{F}_2$ has unsolvable subgroup membership problem (MP), maybe the first natural algorithmic problem involving subgroups.

(Subgroup) membership problem, $\mathsf{MP}(G)$**.** *Given a finite set of words* $w, w_1, \ldots, w_n$, *in the generators of* G, *decide whether* w *represents an element in* $\langle w_1, \ldots, w_n \rangle_G$.

Remark 6.1.29. Note that from the computability point of view, for this particular problem, there is no difference between the ambient being (finitely generated) RAAGS, or (general) PC-groups: as far as we have the local (RAAG) description of the letters involved in the input, the problem remains the same.

Theorem 6.1.30 (Mikhailova, 1958, [Mik58])**.** *The group* $\mathbb{F}_2 \times \mathbb{F}_2$ *has unsolvable subgroup membership problem.* □

This result automatically establishes the family presented by C_4-free graphs as the lowest possible threshold for the class of PC-groups having solvable MP. Note that coherent (and thus Droms) PC-groups lie within this family of candidates; and indeed, in 2008, they were proved to have solvable MP.

Theorem 6.1.31 (Kapovich, Weidmann, and Myasnikov, 2005, [KWM05])**.** *Let* Γ *be a finite chordal graph (i.e.,* $\mathbb{G}_\Gamma$ *is a coherent RAAG) Then,*

(1) $\mathbb{G}_\Gamma$ *has decidable subgroup membership problem.*

(2) *given a finite subset* $S \subseteq \mathbb{G}_\Gamma$, *we can algorithmically find a presentation for the subgroup* $\langle S \rangle \leqslant \mathbb{G}_\Gamma$. □

Remark 6.1.32. Observe the current gap between the PC-groups which are known to have unsolvable MP (namely, those presented by graphs with an induced C_4), and those which are known to have solvable MP (namely, those presented by chordal graphs). This, of course, leaves open many natural questions, for example the ones below (in increasing order of generality).

Question 2. *Does the PC-group presented by the* 5*-cycle have solvable* MP*?*

Question 3. *For which* $n \geqslant 5$ *does the PC-group presented by the* n*-cycle have solvable* MP*?*

Question 4. *Characterize the RAAGs with solvable* MP.

Finally, we recall that for partially commutative monoids, the exact border for the corresponding membership problem is already known: Lohrey and Steinberg prove in [LS08], that the submonoid membership problem is solvable for a partially commutative monoid if and only if its commutation graph is Droms. Note that this implies, in particular, that the 4-path presents a group with solvable group membership problem (it is chordal), but unsolvable monoid membership problem (it is not Droms).

6.2 Intersections

In group theory, the study of intersections of subgroups has been recurrently considered in the literature. Roughly speaking, the problem is "*given* $H, K \leqslant G$, *find* $H \cap K$". However, in the context of Combinatorial Group Theory, where groups may be infinite, or even infinitely generated, one needs to be more precise about the word *find*, specially if one is interested in the computational point of view.

Definition 6.2.1. A group G is said to satisfy *Howson's property* — or to be *Howson*, for short — if the intersection of any two (and so, finitely many) finitely generated subgroups is again finitely generated.

For these groups, a natural meaning for the above sentence is '*given finite sets of generators* for two subgroups $H, K \leqslant G$, *compute a finite set of generators* for the intersection $H \cap K$'.

Classical examples of Howson PC-groups are the aforementioned families of free-abelian, and free groups (note that no rank restriction is needed: since Howson property only alludes to finitely generated subgroups, any instance of the property within a general free, or free-abelian group, can be reduced to a finitely generated ambient, which is again of the same kind). In $\mathbb{Q}^n$ (and $\mathbb{Z}^n$) Howson's property is granted, since *every* subgroup is again finitely generated; and one can use linear algebra (plus finite index considerations) in order to compute a finite set of generators for $H \cap K$, in terms of given finite sets of generators for H and K.

For the case of free groups, the latter is not true, but Howson himself proved in [How54] that the intersection of two finitely generated subgroups is always again finitely generated (here is where the name of the property comes from); and also gave an algorithm to compute generators for the intersection (see Figure 5.49, or [Sta83] for a later reformulation of this result in the nice language of pull-backs of finite automata).

However, not every PC-group is Howson: as we have seen in Part I, the subfamily of free-abelian times free groups turns out to be non-Howson in every non-degenerate case (i.e., they are Howson if and only if they do not have $\mathbb{Z} \times \mathbb{F}_2$ as a subgroup). So, two natural questions emerge:

- is it possible to characterize those PC-groups which are Howson?
- is it possible to decide, given two finitely generated subgroups in a PC-group, whether their intersection is finitely generated or not?

It turns out that the first question — considered below — is straightforward, and easily decidable algorithmically (in the finitely generated case); whereas the study of the — much more involved — second one is developed afterwards, and constitutes the main topic of this part of the dissertation.

6.2.1 Characterizations of Howson PC-groups

Along the following lines, we will see (Theorem 6.2.2) that the same condition as for free-abelian times free groups (namely, not containing $\mathbb{Z} \times \mathbb{F}_2$ as a subgroup) characterizes Howson property within PC-groups; and that it is equivalent to other important algebraic properties, such as being fully residually free.

For limit groups there are many different equivalent definitions. We shall use the one using fully residually freeness: a group G is said to be *fully residually free* if for every finite subset $S \subseteq G$ such that $1 \notin S$, there exist an homomorphism ϕ from G to a free group such that $1 \notin \phi(S)$. Then, a *limit group* can be defined as a finitely generated fully residually free group.

From this definition, it is not difficult to see that both free and free-abelian groups are fully residually free, and that subgroups and free products of fully residually free groups are again fully residually free. It is also straightforward to see that limit groups are *commutative-transitive* i.e., , for $1 \neq x, y, z \in G$, if x commutes with y and y commutes with z then x commutes with z (see [Wil05] for a gentle introduction to limit groups, and [CG05] for a more detailed — topologically oriented — survey focusing on its relation with the universal theory of free groups).

In [RSS13], the authors study the family of finitely generated partially commutative groups for which the fixed points subgroup of every endomorphism is finitely generated. Concretely, in Theorem 3.1 they characterize this family as those groups consisting in (finite) free products of finitely generated free-abelian groups.

Although it is an immediate consequence of well know facts, below we provide an elementary proof for two extra characterizations of this family, namely: being Howson, and being a limit group. Moreover, we observe that, for some of the properties, no restriction in the cardinal of the generating set is needed, and the result holds in full generality (i.e., for every — possibly non-finitely generated — PC-group).

As proved by Rodaro, Silva, and Sykiotis in [RSS13, Theorem 3.1], if we restrict to finitely generated PC-groups, (any of) the conditions in Lemma 6.1.15 describe exactly the family of those having finitely generated fixed point subgroup for every endomorphism — or equivalently, those having finitely generated periodic point subgroup for every endomorphism.

In the following theorem, we provide two extra characterizations for the PC-groups described in Lemma 6.1.15 (including the infinitely generated case). For completeness, we summarize them in a single statement together with the conditions discussed above.

Theorem 6.2.2. *Let Γ be an arbitrary (possibly infinite) simple graph, and $\mathbb{G}_\Gamma$ the PC-group presented by Γ. Then, the following conditions are equivalent:*

(a) *$\mathbb{G}_\Gamma$ is fully residually free,*

(b) *$\mathbb{G}_\Gamma$ is Howson,*

(c) *$\mathbb{G}_\Gamma$ does not contain $\mathbb{Z} \times \mathbb{F}_2$ as a subgroup,*

(d) *the graph Γ is P_3-free,*

(e) *$\mathbb{G}_\Gamma$ is a free product of free-abelian groups.*

Moreover, if the graph Γ is finite (i.e., $\mathbb{G}_\Gamma$ is finitely generated), then the following additional conditions are also equivalent:

(f) *For every $\phi \in \operatorname{End} \mathbb{G}_\Gamma$, the subgroup*

$$\operatorname{Fix} \phi = \{ g \in \mathbb{G}_\Gamma \,:\, \phi(g) = g \}$$

of fixed points of ϕ is finitely generated.

(g) *For every $\phi \in \operatorname{End} \mathbb{G}_\Gamma$, the subgroup*

$$\operatorname{Per} \phi = \{ g \in \mathbb{G}_\Gamma \,:\, \exists n \geqslant 1 \; \phi^n(g) = g \}$$

of periodic points of ϕ is finitely generated.

Proof. [(a) $\Rightarrow$ (b)]. Dahmani obtained this result for limit groups (i.e., assuming $\mathbb{G}_\Gamma$ finitely generated) as a consequence of them being hyperbolic relative to their maximal abelian non-cyclic subgroups (see Corollary 0.4 in [Dah03]). We note that the finitely generated condition is superfluous for this implication since the Howson property involves only finitely generated subgroups, and every subgroup of a fully residually free group is again fully residually free.

[(b) $\Rightarrow$ (c)]. Since Howson's property is subgroup-hereditary, it is enough to show that the group $\mathbb{Z} \times \mathbb{F}_2$ does not satisfy the Howson property (see Lemma 2.3.1).

[(c) $\Rightarrow$ (d) $\Rightarrow$ (e)]. From Lemma 6.1.20, if $\mathbb{G}_\Gamma$ does not contain the group $\mathbb{Z} \times \mathbb{F}_2$ (which is presented by P_3) as a subgroup, then P_3 is not a full subgraph of Γ. Equivalently (see Lemma 6.1.15), $\mathbb{G}_\Gamma$ is a free product of free-abelian groups.

[(e) $\Rightarrow$ (a)]. This is again clear, since both free-abelian groups and free products of fully residually free groups are again fully residually free. Note here, that no cardinal restriction is needed; neither for the rank of the free-abelian groups, nor for the number of factors in the free product, since the definition of fully residually freeness involves only finite families.

Finally, for the equivalence between (e), (f) and (g) under the finite generation hypothesis, see Theorem 3.1 in [RSS13]. □

Note that Remark 6.1.26, together with condition (d) in theorem 6.2.2, immediately provides the algorithmic recognizability of Howson RAAGs, in the finitely generated case.

An immediate corollary of Lemma 6.1.20 is that the PC-group presented by any full subgraph $\Delta \leqslant \Gamma$ is itself a subgroup of the PC-group presented by Γ, i.e., for every pair of graphs Γ, Δ, we have $\Delta \leqslant \Gamma \;\Rightarrow\; \mathbb{G}_\Delta \leqslant \mathbb{G}_\Gamma$.

The (eventual) converse implication suggests a distinguished family of PC-groups, that we will call explicit.

Definition 6.2.3. A PC-group $\mathbb{G}_\Delta$ (or the graph Δ presenting it) is said to be *explicit* if whenever you have it as a subgroup of a PC-group $\mathbb{G}_\Gamma$, then its commutation graph Δ is an induced subgraph of Γ; that is, if for every graph Γ,

$$\Delta \leqslant \Gamma \;\Leftrightarrow\; \mathbb{G}_\Delta \leqslant \mathbb{G}_\Gamma\,. \tag{6.3}$$

For example, it is straightforward to see that the only explicit edgeless graphs are those with zero, one, and two vertices: the first two cases are obvious, and for the third one, note that if $\mathbb{F}_2 \leqslant G$ then G can not be abelian. Finally, for $n \geqslant 3$, it is sufficient to note (again) that $\mathbb{F}_n$ is not an explicit subgroup of $\mathbb{F}_2$.

At the opposite extreme, the following is a well-known result relating the maximum abelian rank of a PC-group $\mathbb{G}_\Gamma$ with the order of a largest clique in Γ (see [KK13, Lemma 18]).

Lemma 6.2.4. *The maximum rank of a free-abelian subgroup of a RAAG $\mathbb{G}_\Gamma$ is the size of a largest clique in Γ.* □

As an immediate consequence, we have that every (finite) complete graph K_r is explicit:

$$\mathbb{Z}^r \leqslant \mathbb{G}_\Gamma \;\Leftrightarrow\; K_r \leqslant \Gamma\,, \tag{6.4}$$

and therefore, no infinitely generated free-abelian group can be embedded into a RAAG.

In the last years, embedability between PC-groups has been matter of growing interest and research (see [Kam09], [KK13] and [CDK13]) which has provided some new examples of explicit graphs, such as the square C_4 (proved by Kambites, in [Kam09]), or the path on four vertices P_4 (proved by Kim and Koberda in [KK13]).

To end with, we just remark that our characterization theorem (Theorem 6.2.2) immediately provides a new member of this family.

Corollary 6.2.5. *The path on three vertices* P_3 *is explicit.* □

Intersection problem for Droms groups 7

As we have already discussed, within any ambient containing non-Howson groups (e.g. that of PC-groups) it is natural to include the decision about finite generation of the intersection among the requirements of any reasonable definition for (algorithmic) *subgroup intersection problem*.

We formalize below this classical definition, together with the closely related *coset intersection problem*. Note, that every problem has a *decision version* consisting in checking some property, and the corresponding *search version* which, depending on the problem, could consist in finding a witness, or in characterizing all witnesses for the property. Of course, such a witness can only exist when the property holds. Thus, from the algorithmic point of view, the reasonable combination of both problems (that we call the full version of the problem) is defined to consist in first deciding whether the property holds or not, and if so, find a witness for it.

7.1 Algorithmic intersection problems

We recall the standard intersection problems for subgroups and cosets in a finitely presented group $G = \langle X \mid R \rangle$.

Subgroup intersection (decision) problem, $\mathsf{SIP_d}(G)$**.** *Given a finite set of words* $u_1, \ldots, u_n, v_1, \ldots, v_m$ *in the generators of* G*, decide whether the subgroup intersection* $\langle u_1, \ldots, u_n \rangle_G \cap \langle v_1, \ldots, v_m \rangle_G$ *is finitely generated or not.*

Subgroup intersection (search) problem, $\mathsf{SIP_s}(G)$**.** *Given a finite set of words* $u_1, \ldots, u_n$, $v_1, \ldots, v_m$ *in the generators of* G*, compute a generating set for the subgroup intersection* $\langle u_1, \ldots, u_n \rangle_G \cap \langle v_1, \ldots, v_m \rangle_G$.

Subgroup intersection (full) problem, $\mathsf{SIP}(G)$**.** *Given a finite set of words* $u_1, \ldots, u_n$, $v_1, \ldots, v_m$ *in the generators of* G*, decide whether the subgroup intersection* $\langle u_1, \ldots, u_n \rangle_G \cap \langle v_1, \ldots, v_m \rangle_G$ *is finitely generated or not; and* in affirmative case, *compute a generating set for this intersection.*

In Part I (see also [DV13]) we prove that direct products of free-abelian and free groups satisfy SIP. The goal of the present chapter is to extend and complement the algebraic arguments given there, in order to prove the same property (and some technically stronger variations) for a much wider family of groups, namely that of Droms PC-groups. In the course of doing so, we realized that SIP in free-like contexts is closely related to the also classical Coset Intersection Problem (CIP), stated below.

Coset intersection (decision) problem, $\mathsf{CIP_d}(G)$**.** *Given a finite set of elements* w, w', $u_1, \ldots, u_n, v_1, \ldots, v_m$ *in the generators of* G*, decide whether the coset intersection* $w\langle u_1, \ldots, u_n\rangle_G \cap w'\langle v_1, \ldots, v_m\rangle_G$ *is empty or not.*

Coset intersection (search) problem, $\mathsf{CIP_s}(G)$**.** *Given a finite set of elements* w, w', $u_1, \ldots, u_n, v_1, \ldots, v_m$ *in the generators of* G*, compute a representative for the coset intersection* $w\langle u_1, \ldots, u_n\rangle_G \cap w'\langle v_1, \ldots, v_m\rangle_G$.

Coset intersection (full) problem, $\mathsf{CIP}(G)$**.** *Given a finite set of elements* w, w', $u_1, \ldots, u_n, v_1, \ldots, v_m$ *in the generators of* G*, decide whether the coset intersection* $w\langle u_1, \ldots, u_n\rangle_G \cap w'\langle v_1, \ldots, v_m\rangle_G$ *is empty or not; and in negative case compute a coset representative.*

Remark 7.1.1. Note that, in terms of computability, the search problem in CIP is superfluous (i.e., $\mathsf{CIP} \Leftrightarrow \mathsf{CIP_d} \Rightarrow \mathsf{CIP_s}$). Namely, the computation of a coset representative $g'' \in gH \cap g'K$ once known that $gH \cap g'K \neq \varnothing$ can always be done by a brute force search along an enumeration of all the elements in G (even without assuming the word problem for G, since we only need its YES part, which is always computable).

Also, note that, since $(Hg)^{-1} = g^{-1}H$ and $g_1Hg_2 = g_1g_2(g_2^{-1}Hg_2) = g_1g_2H^{g_2}$, the variants of CIP for right and left cosets are equivalent problems.

Abelian groups clearly satisfy CIP, as do free groups as well — with a well-known modification of the pullback argument given in Stallings [Sta83] (see also Section 5.5.1). Moreover, in [DV13] CIP is also proved for groups of the form $\mathbb{Z}^m \times \mathbb{F}_n$.

Some convenient variations of the previous intersection problems are considered below for a general finitely presented group $G = \langle X \mid R\rangle$.

Twofold intersection problem, $\mathsf{TIP}(G)$**.** *Solve both* $\mathsf{SIP}(G)$ *and* $\mathsf{CIP}(G)$*. Mnemonically,*

$$\mathsf{TIP}(G) = \mathsf{SIP}(G) + \mathsf{CIP}(G).$$

Extended subgroup intersection problem, $\mathsf{ESIP}(G)$**.** *Given a finite set of elements* $w, w', u_1, \ldots, u_n, v_1, \ldots, v_m$ *in the generators of* G*, decide whether the intersection of the subgroups* $H = \langle u_1, \ldots, u_n \rangle_G$ *and* $K = \langle v_1, \ldots, v_m \rangle_G$ *is finitely generated or not; and in affirmative case, compute a generating set for the subgroup intersection* $H \cap K$*, and decide whether the coset intersection* $wH \cap w'K$ *is empty or not. Mnemonically,*

$$\mathsf{ESIP}(G) \;=\; \mathsf{SIP}(G) + \mathsf{CIP}_{\mathsf{fg}}(G)\,. \tag{7.1}$$

Notice that the difference between properties TIP and ESIP is that the second one says nothing about $gH \cap g'K$ in the case when $H \cap K$ is not finitely generated, while TIP is required to answer about emptyness even in this case; this is a subtlety that will become important along the chapter.

Remark 7.1.2. Note that all these properties (TIP, ESIP, CIP, SIP, MP and WP) clearly pass to subgroups.

Several obvious relations among already introduced problems are summarized in the diagram below:

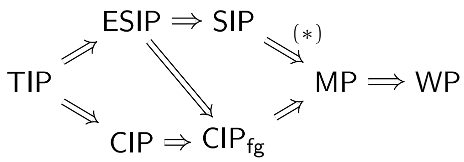

Fig. 7.1: Some dependencies between algorithmic problems

where the starred implication is true with the extra assumption that the involved group G is torsion-free, and has solvable word problem.

Lemma 7.1.3. *If a torsion-free group satisfies* SIP *and* WP*, then it also satisfies* MP*.*

Proof. Let $G = \langle X \mid R \rangle$. Given words $u, v_1, \ldots, v_m$ in X (call $h = [u]_G$, and $k_i = [v_i]_G$), apply SIP to $H = \langle u \rangle_G$, and $K = \langle v_1, \ldots, v_m \rangle_G$): since $H \cap K$ is cyclic (and so, finitely generated), SIP will always answer YES, and return a finite set of words $w_1, \ldots, w_p$ in X representing elements $g_1, \ldots, g_p \in G$, such that $H \cap K = \langle g_1, \ldots, g_p \rangle = \langle h^r \rangle$, for some $r \in \mathbb{Z}$.

Now, since each g_i must be a power of h (say $g_i = h^{r_i}$), we can compute the exponents $r_1, \ldots, r_p \in \mathbb{Z}$ by enumerating all powers of u and searching, by brute force, which of them represents each g_i, for each $i = 1, \ldots, r$. This can be done without using WP, since we already have the information that each g_i is indeed a power of h. That is, for each $i = 1, \ldots, r$, we can start checking — in parallel — whether $w_i u^{-r_i} \in \langle\langle R \rangle\rangle$, for $r = 0, \pm 1, \pm 2, \ldots$ until we find a suitable r_i.

Once we have obtained explicit integers $r_1, \ldots, r_p \in \mathbb{Z}$ such that $g_i = h^{r_i}$, we can effectively compute the greatest common divisor $r = \gcd(r_1, \ldots, r_p)$ satisfying $H \cap K = \langle g_1, \ldots, g_p \rangle = \langle h^r \rangle$.

Now, it is clear that $h \in K$ if and only if $h \in H \cap K = \langle h^r \rangle$; i.e., if and only if $h = h^{rs}$, for some $s \in \mathbb{Z}$.

To decide whether such an s exists, first apply WP to the input word u in order to decide whether $h = 1$ or not. In the affirmative case the answer is obviously YES; and otherwise (i.e., if $h \neq 1$) torsion-freeness of G tells us that $h \in K$, and the answer is YES, if and only if $r = \pm 1$. □

Remark 7.1.4. The torsion-freeness hypothesis is not necessary for the $\mathsf{CIP}_{fg} \Rightarrow \mathsf{MP}$ implication, since $g \in \langle h_1, \ldots, h_k \rangle \Leftrightarrow g \cdot \{1\} \cap 1 \cdot \langle h_1, \ldots, h_k \rangle \neq \varnothing$.

Corollary 7.1.5. *For PC-groups, both* SIP *and* CIP *imply* MP. *In particular,* SIP *and* CIP *are unsolvable for* $\mathbb{F}_2 \times \mathbb{F}_2 = \mathbb{G}_{C_4}$*, and hence for any PC-group containing it.*

Proof. It is enough to recall that PC-groups are torsion-free, that SIP, CIP and MP pass to subgroups; and that $\mathbb{F}_2 \times \mathbb{F}_2 = \mathbb{G}_{C_4}$ has unsolvable MP, and is explicit. □

So, we can not pretend to prove the solvability of SIP (or ESIP) and CIP for PC-groups in full generality: the maximum possible scope being — in principle — the same as that of MP, namely the family presented by C_4-free graphs. In the next section we discuss an important subfamily of that of C_4-free graphs for which we are able to prove the solvability of both problems.

7.2 Droms groups

The family of Droms groups has often established the threshold of positive algorithmic properties within both partially commutative groups (see e.g. [RSS13]), and partially commutative monoids (see [AH89; LS08]). Particularly suggestive for us is the result from Aalbersberg and Hoogeboom in [AH89] stating that the intersection problem for a partially commutative monoid is solvable if and only if its commutation graph is Droms. Moreover, since they have solvable membership problem (recall that they are coherent and Figure 7.4), they constitute a natural target where to study intersection problems.

Recall that Droms groups are the finitely generated PC-groups having all their subgroups being again PC-groups, and correspond exactly to those PC-groups presented by finite $\{C_4, P_4\}$-free graphs (which we call *Droms graphs*).

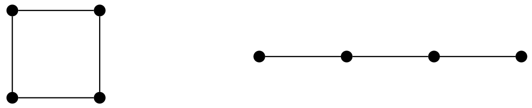

Fig. 7.2: Forbidden induced graphs in Droms graphs

This family of graphs has also been called *transitive forests*, *trivially perfect graphs*, *comparability graphs of trees*, or *quasi-threshold graphs*, and many equivalent characterizations are known for it (see [Wol62; Wol65; Gol78; JJC96; Rub15]).

Next, we detail one of such characterizations (providing a recursive construction which will allow us to use induction to prove the solvability of SIP and CIP for this family). It follows easily from the property stated below.

Lemma 7.2.1. *Every nonempty Droms graph is either disconnected, or contains a central vertex.*

Proof 1. By transposition, suppose that Γ is connected, Droms, and centerless. Note that, since Droms graphs are finite, every vertex has bounded degree.

Now, consider a vertex t non-adjacent to a vertex q of maximum degree in Γ (there must be at least one because otherwise q would be central), and a shortest path P from q to t. Let q, r, s be the starting vertices of P (note that s is not adjacent to q and may equal t). Since $\deg(q) \geqslant \deg(r)$, there is at least one neighbour p of q which is not adjacent to r. But then, the 4-path $pqrs$ have no chords (p is not adjacent to r, and s is not adjacent to q). Thus, either p is adjacent to s and the 4-cycle $pqrs$ is induced in Γ, or p is not adjacent to s and the 4-path $pqrs$ is induced in Γ. This completes the proof. □

Proof 2. By induction on n (the number of vertices of the graph).

$[n = 1, 2, 3]$ Trivial.

$[n = k \Rightarrow n = k+1]$ Let Γ_{k+1} be a connected Droms (i.e., $\{P_4, C_4\}$-free) graph with $k+1$ vertices. Then every connected component of any k-full subgraph of Γ_{k+1} is again a nonempty connected Droms graph. Let p be the vertex in $V(\Gamma_{k+1} \setminus \Gamma_k)$, for some k-full subgraph Γ_k of Γ_{k+1}.

Then, from the induction hypothesis (IH), we know that every connected component $\mathcal{C}_j$ of Γ_k has nontrivial center Z_j. Now we distinguish two cases:

(a) Γ_k is connected, and hence there is only one connected component in Γ_k with center Z. Then, we claim that V is adjacent to every vertex in Z, which is hence also central in Γ_{k+1}. This is so because otherwise, since Γ_{k+1} is connected, then there must exist a vertex $r \in \Gamma_k \setminus Z$ adjacent to q, and — since r is noncentral in Γ_k — another noncentral vertex s in Γ_k not adjacent to r. But then, if we take a central vertex $z \in Z$, we only have two possibilities:

- q is adjacent to s, and thus $q — r — z — s — q$ is an induced square in Γ_{k+1}.
- q is not adjacent to s, and thus $q — r — z — s$ is an induced 4-path in Γ_{k+1}.

Since none of this possibilities is possible (recall that Γ_{k+1} is $\{P_4, C_4\}$-free by hypothesis), we conclude that q is adjacent to every vertex in Z as claimed.

(b) Γ_k is disconnected, and hence Γ_k has at least two different connected components. Then, we claim that q is central in Γ_{k+1}. Again, we distinguish two cases:

- every connected component in Γ_k contains a single vertex. Then (since Γ_{k+1} is connected) q must be adjacent to every other vertex in Γ_{k+1}.
- some connected component in Γ_k, say $\mathcal{C}$, contains at least 2 vertices. Now suppose, by contraposition, that q is not central in Γ_{k+1}. Then, since (Γ_{k+1} is connected) there must exist a vertex r in $\mathcal{C}$ not adjacent to q, and adjacent to some vertex z central in Γ_k. But then, for any vertex z adjacent to q in any connected component different from $\mathcal{C}$ (which must exist since Γ_{k+1} is connected), we will have that $r — z — q — t$ is an induced 4-path in Γ_{k+1}, which is not possible since Γ_{k+1} is Droms.

Thus, in all cases, the graph Γ_{k+1} contains a central vertex, and the proof is concluded. □

Below, we warn that the last claim is no longer true if we remove the finiteness condition in Lemma 7.2.1.

Lemma 7.2.2. *Let Υ be the (infinite) graph with vertices the nonnegative integers, and every odd vertex adjacent to the rest of odd vertices and to the even vertices smaller than it (see Figure 7.3). Then, Υ is nonempty, centerless, and $\{P_4, C_4\}$-free.*

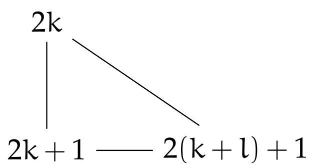

$(\forall k \geqslant 0,$ *and* $\forall l \geqslant 1)$

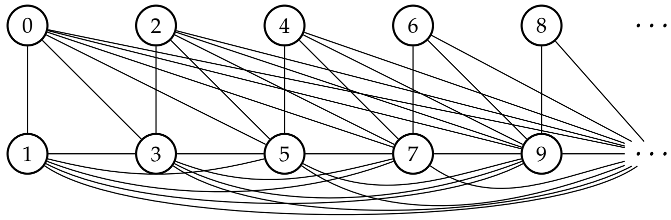

Fig. 7.3: The (infinite) centerless connected $\{P_4, C_4\}$-free graph Υ

Proof. Note that Υ is clearly connected and centerless (there is no odd integer greater than any even number). Finally let's see that every 4-path in Υ has a chord.

First, observe that every 4-path in Υ must contain at least two odd vertices: since the set of even vertices is independent in Υ, there exist no nontrivial connected subgraphs of even vertices, and every connected 4-subgraph Δ having only one odd vertex has a vertex (the odd one) of degree 3, and thus is not a path.

It is obvious that any 4-path P with three or more odd vertices has a chord (since they belong to a complete subgraph of Υ). So, there is only one case left, namely 4-paths with two even (say $p < q$), and two odd (say $r < s$) vertices. Note that for any such path the lowest in the path must be even, and so adjacent to both r and s, which are odd greater than p, and thus also adjacent. This provides a triangle in P, and concludes the proof. □

Indeed, since the higher vertex must be odd, there are only two possibilities, namely $p < q < r < s$, and $p < r < q < s$, corresponding to the following induced subgraphs of Υ.

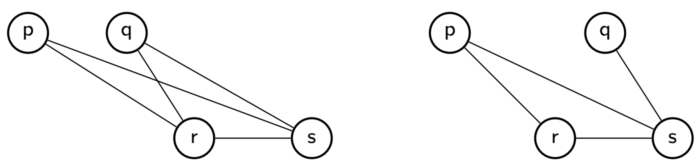

Fig. 7.4: Possible induced 4-vertex subgraphs of Υ containing P_4 (both chordal)

Remark 7.2.3. The previous example shows that the property in Lemma 7.2.1 is no longer true if we replace finite by infinite. Therefore the inductive techniques derived from this lemma (below) will be no longer applicable in the more general case of (maybe infinite) $\{P_4, C_4\}$-free PC-groups.

The key consequence of Lemma 7.2.1 is that if we remove the center from a Droms graph, we get either the empty graph, or a disconnected Droms graph.

Corollary 7.2.4. *Let Γ be a Droms graph, then*

$$\Gamma = Z(\Gamma) \vee (\Gamma \setminus Z(\Gamma)) = K_m \vee \Gamma_0 \, , \tag{7.2}$$

where K_m is a complete graph of rank $m \geqslant 0$, and $\Gamma_0 := \Gamma \setminus Z(\Gamma)$ is either empty or a disconnected Droms graph.

Proof. It is clear that $\Gamma = Z(\Gamma) \vee \Gamma_0$, and that $Z(\Gamma) = K_m$, for certain $m \geqslant 0$. So, it is enough to see that if Γ is Droms, then $\Gamma_0 = \Gamma \setminus Z(\Gamma)$ is either empty or disconnected. By transposition, suppose that Γ_0 is nonempty and connected. Then, since Γ_0 is a full subgraph of Γ, Γ_0 is again Droms, and there must exist (Lemma 7.2.1) a vertex $p \in \Gamma_0$ which is adjacent to every other vertex in Γ_0. But, by construction, every vertex in Γ_0 is also adjacent to every vertex outside Γ_0, therefore $p \in \Gamma_0 \cap Z(\Gamma)$, contradicting our definition of Γ_0. □

Remark 7.2.5. Let $\Gamma = K_m \vee \Gamma_0$ be a nonempty Droms graph decomposition, like (7.2). Then:

- $\mathbb{G}_\Gamma$ is free-abelian $\Leftrightarrow$ Γ is complete $\Leftrightarrow$ $Z(\Gamma) = \Gamma$ $\Leftrightarrow$ $\Gamma_0 = \varnothing$.
- $\mathbb{G}_\Gamma$ is *connected* $\overset{\text{def}}{\Leftrightarrow}$ Γ is connected $\Leftrightarrow$ $Z(\Gamma) \neq \varnothing$ $\Leftrightarrow$ $m \geqslant 1$ $\Leftrightarrow$ Γ is a cone .
- $\mathbb{G}_\Gamma$ is centerless $\Leftrightarrow$ Γ is disconnected $\Leftrightarrow$ $Z(\Gamma) = \varnothing$ $\Leftrightarrow$ $m = 0$.

The recursive definition follows immediately from the previous corollary. Below we state this key result in parallel with its algebraic counterpart.

Corollary 7.2.6. *The family of Droms graphs (resp., Droms groups) can be recursively defined as the smallest family $\mathfrak{D}$ (resp., $\mathcal{D}$) satisfying the following rules:*

[D1] $K_0 \in \mathfrak{D}$.	[D1] $\{1\} \in \mathcal{D}$.
[D2] $\Gamma_1, \Gamma_2 \in \mathfrak{D} \Rightarrow \Gamma_1 \sqcup \Gamma_2 \in \mathfrak{D}$.	[D2] $G_1, G_2 \in \mathcal{D} \Rightarrow G_1 * G_2 \in \mathcal{D}$.
[D3] $\Gamma \in \mathfrak{D} \Rightarrow K_1 \vee \Gamma \in \mathfrak{D}$.	[D3] $G \in \mathcal{D} \Rightarrow \mathbb{Z} \times G \in \mathcal{D}$.

Note that grouping together all successive applications of rules (2) and (3) in the construction, we can respectively substitute them by:

[D2]′ $(\forall k)\ \Gamma_1, \ldots, \Gamma_k \in \mathfrak{D} \Rightarrow \bigsqcup_{i=1}^{k} \Gamma_i \in \mathfrak{D}$.	[D2]′ $(\forall k)\ G_1, \ldots, G_k \in \mathcal{D} \Rightarrow \ast_{i=1}^{k} G_i \in \mathcal{D}$.
[D3]′ $(\forall m)\ \Gamma \in \mathfrak{D} \Rightarrow K_m \vee \Gamma \in \mathfrak{D}$.	[D3]′ $(\forall m)\ G \in \mathcal{D} \Rightarrow \mathbb{Z}^m \times G \in \mathcal{D}$.

Proof. It is obvious that the empty graph K_0 is Droms, and that the union of Droms graphs is again Droms. To see that if Γ is Droms, then $K_1 \vee \Gamma$ is Droms, suppose — by transposition — that P is an induced 4-path or square within $K_1 \vee \Gamma$. Since Γ is Droms, P must contain the joined vertex p. But then any edge joining p with the rest of vertices in P will be also in P. Therefore, P is neither a 4-path nor a square, in contradiction with the starting assumption.

This proves that every graph in $\mathcal{D}$ is Droms. The converse implication is an immediate consequence of Corollary 7.2.4. □

Remark 7.2.7. Observe that we can restrict step [D2]′) only to connected graphs; similarly, we can restrict step or [D3]′ only to disconnected (or trivial) graphs, and thus reduce the construction of any Droms graph to step [D1] followed by an alternate sequence of steps [D2]′ and [D3]′.

Note that this allows us to compactify the previous lemma into the one below.

Corollary 7.2.8. *Every Droms graph Γ is either complete, or the join of a (maybe empty) complete graph with a finite disjoint union of two or more nonempty connected Droms graphs.*

Moreover, this decomposition is unique; i.e., given a Droms graph Γ, there exist unique nonnegative integers $m \geqslant 0$, and $1 \neq k \geqslant 0$, and unique connected Droms graphs Λ_i $(i = 1, \ldots, k)$ such that

$$\Gamma = K_m \vee \bigsqcup_{i=1}^{k} \Lambda_i\,, \quad \left(\textit{equivalently, } \mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathop{\ast}_{i=1}^{k} \mathbb{G}_{\Lambda_i} \right). \tag{7.3}$$

Recall that $K_0 = \varnothing$, and we are using the convention that $\bigsqcup_{i=1}^{0} \Lambda_i = \varnothing$. □

Definition 7.2.9. The decomposition in (7.3) is called *primary decomposition* of the Droms graph Γ (or the Droms group $\mathbb{G}_\Gamma$). Recursively applying (7.3) to every non-complete connected graph appearing in (7.3), we can build a unique finite rooted tree (having complete graphs as leafs) describing Γ. This tree is called the *full decomposition* of Γ (or $\mathbb{G}_\Gamma$).

The following important result states that, not only finitely generated subgroups of Droms groups are again Droms, but that a graphic presentation for them is always computable from any finite set of generators, together with an explicit isomorphism expressing the original generators in terms of the new basis, and vice versa.

Proposition 7.2.10. *Let $\mathbb{G}_\Gamma = \langle X \mid R \rangle$ be a Droms group. Then, there exists an algorithm which, given words $w_1(X), \ldots, w_p(X)$ in the generators of $\mathbb{G}_\Gamma$:*

(i) *Computes a (basis/graphic presentation) for the subgroup $H = \langle w_1, \ldots, w_p \rangle \leqslant \mathbb{G}_\Gamma$.*

(ii) *Provides an effective isomorphism between the original generators and the new basis (and recursively enumerates all such isomorphisms).*

Proof. Since Droms graphs are chordal, by Theorem 6.1.31, we can effectively compute a finite presentation for H, say $H = \langle Y \mid S \rangle$.

Then, one can exhaustively explore the tree of all possible Tietze transformations applied to $\langle Y \mid S \rangle$ until getting one, say $\langle Z \mid Q \rangle$, in graphic form (namely, with all relators being commutators of certain pairs of generators); this will be achieved in finite time because we know in advance that H is indeed a finitely generated PC-group.

Note that at this point, we know that $H = \langle w_1, \ldots, w_p \rangle_{\mathbb{G}_\Gamma} \simeq \langle Y \mid S \rangle \simeq \langle Z \mid Q \rangle$; and even if we don't know the relation between the elements in the obtained basis $Z = \{z_1, \ldots, z_r\}$, and the original words $w_1(X), \ldots, w_p(X)$, since we know that such an isomorphism does indeed exist, we can start a brute force quest for it using the following two parallel procedures.

1. *Enumerate all homomorphisms* $z_j \mapsto \langle w_1, \ldots, w_p \rangle$.

   This can be done preforming , again in parallel, the following tasks:

   1.1. *Enumerate candidates*: enumerate all possible r-tuples $(v_j)_{j=1}^r$ of words in $\{w_1, \ldots, w_p\}^{\pm}$.

   (Note that each such k-tuple corresponds to a candidate homomorphism defined by $z_j \mapsto v_j$ for $j = 1, \ldots, r$.)

   1.2. *Filtering homomorphisms*: for each candidate from 1.1., check whether when substituting each occurrence of z_j by the corresponding v in each (commutation) relations in Q, the result is a relation in R.

   (Note that this procedure provides an enumeration of all homomorphisms $z_j \mapsto \langle w_1, \ldots, w_p \rangle$.)

2. *Filtering epimorphisms*: for each homomorphism $(z_j \mapsto v_j)_{j=1}^r$ from 1.2., enumerate all words in $\{v_1, \ldots, v_r\}^{\pm}$, and check whether every w_j $(j = 1, \ldots, p)$ appears in the list.

   (Note that, for every non-surjective homomorphisms, this process will continue forever, but if all then are preformed in parallel, every epimorphism will be detected in finite time.)

(Note also, that WP is available — since the involved groups are PC-groups — but not really necessary because we only need its YES part to perform the above checks.)

So, the previous combined algorithm outputs a list of all surjective homomorphisms $(z_j \mapsto v_j)_{j=1}^r$. Finally, recall that finitely generated PC-groups (and hence $\mathbb{G}_\Gamma$) are Hopfian (Proposition 6.1.8), and thus, this is indeed a list of all possible isomorphisms $(z_j \mapsto v_j)_{j=1}^r$. This completes the proof. □

Remark 7.2.11. This last result allows us to suppose any given generating set for a subgroup of a Droms group to be graphic (i.e., a basis).

7.3 Results

The main results of the chapter concern preservability of interesection properties through free and direct products, and the corresponding implications on the family of Droms groups. We summarize them below.

Theorem 7.3.1. *If two finitely presented groups* G_1 *and* G_2 *satisfy* TIP, *then their free product* $G_1 * G_2$ *also satisfies* TIP. □

Theorem 7.3.2. *If two finitely presented groups* G_1 *and* G_2 *satisfy* ESIP, *then their free product* $G_1 * G_2$ *also satisfies* ESIP. □

Theorem 7.3.3. *Let* G *be a Droms PC-group. If* G *satisfies* SIP, *then* $\mathbb{Z}^m \times G$ *also satisfies* SIP. □

Theorem 7.3.4. *Let* G *be a Droms PC-group. If* G *satisfies* ESIP, *then* $\mathbb{Z}^m \times G$ *also satisfies* ESIP. □

Note that similar preserving properties were studied for the Membership Problem (MP) by Mikhailova. In [Mik68] she proves that MP is preserved under free products; whereas in the already mentioned paper [Mik58], she shows that $\mathbb{F}_2 \times \mathbb{F}_2$ has unsolvable membership problem, proving that MP (and thus SIP and CIP) *do not* pass to direct products.

Since, by Theorems 7.3.2 and 7.3.4, the inductive steps [D2] and [D3]′ preserve ESIP, we deduce that all Droms PC-groups enjoy such property.

Theorem 7.3.5. *Every Droms PC-group satisfies* ESIP *(and, in particular,* SIP*).*

Proof. Let Γ be a Droms graph, and let $\mathbb{G}_\Gamma$ be the corresponding Droms PC-group. We will prove the result by induction on the number of vertices $|V\Gamma|$. If $|V\Gamma| = 0$, then $\mathbb{G}_\Gamma = 1$, and ESIP is obvious.

Now, consider a nonempty Droms graph Γ, and assume (I.H.) that every Droms PC-group with strictly less than $|V\Gamma|$ vertices does satisfy ESIP. Then, note that either Γ is complete (i.e., $\mathbb{G}_\Gamma$ is free-abelian, and thus satisfies ESIP); or otherwise, every graph Λ_i in the primary decomposition (7.3) is a Droms graph with strictly less that $|V\Gamma|$ vertices. So, applying the induction hypothesis, every Λ_i satisfies ESIP. Now, the solvability of ESIP for $\mathbb{G}_\Gamma$ follows from Theorems 7.3.2 and 7.3.4 and the decomposition in (7.3). This concludes the proof. □

This can be seen as a generalization of a result by Kapovich, Weidmann, and Myasnikov [KWM05] in the following sense. Since ESIP implies MP, our Theorem 7.3.5 proves a potentially stronger thesis than Theorem 6.1.31 for a smaller class of groups. It is interesting to ask the following question (open as far as we know):

Question 5. *Does the group* $\mathbb{G}_{P_4}$ *have solvable* SIP*?*

If the answer were negative then we would have a positive answer to the following open question:

Question 6. *Is it true that a PC-group satisfies* SIP *if and only if it is Droms?*

We devote Sections 7.4 and 7.5 to prove the inductive results regarding direct and free products, respectively.

7.4 The free product case

In this section, we shall prove the inductive intersection problems for free products, namely Theorems 7.3.1 and 7.3.2.

We follow the graph-theoretic approach of Ivanov in [Iva99], generalizing the classical Stallings machinery (see Section 5.4) for subgroups of free groups, to subgroups of free products. This allowed him to give an alternative proof that free products of Howson groups are Howson, as well as provide bounds for such an intersection (see [Iva99; Iva01; Iva08]). Namely, we use generalized folding techniques to algorithmically represent any finitely generated subgroup H of a free product, by a finite automaton providing a Kurosh decomposition for H. Then we analyze how the graph for the intersection $H \cap K$ is related to the graphs for H and K. This allows us, using ESIP for the factors, to decide whether $H \cap K$ is finitely generated or not; and if so, construct the graph of $H \cap K$ and from it compute the generators of $H \cap K$. Similarly, modified graphs representing cosets allow us to decide whether the intersection $Ha \cap Kb$ is empty or not, provided that $H \cap K$ is finitely generated.

Note that the use of this kind of automata to describe subgroups was later extended to the context of graphs of groups by Kapovich, Weidmann, and Myasnikov in [KWM05]. Their methods, which are based on a further generalization of the folding techniques, allow to construct algorithmically the subgroup automaton under some conditions on the edge groups, which automatically hold in the case of free products. However, they do not analyze subgroup intersections. On the contrary, Ivanov describe subgroup intersections, but does not consider their algorithmic behavior. For this reason we provide almost complete details.

So, we give algorithmic treatment to Ivanov's description for the intersection of subgroups in a free product, in order to solve the intersection problems stated in Section 7.1. Roughly speaking, the proof goes like this: given two finitely generated subgroups $H, K \leqslant G_1 * G_2$, construct the corresponding automata Γ_H and Γ_k which reflect the (Kurosh) free product structure of H, K respectively, and then adapt Ivanov's methods to either compute the full automaton recognizing the intersection $H \cap K$ (when it is finitely generated), or detecting when it is not. This allows us to solve various intersection problems on $G_1 * G_2$, assuming they hold in the factors.

The use of Stallings automata to understand the lattice of subgroups within different families is well-known, and we assume the reader is familiar with it; see Chapter 5 for the original version for free groups and a generalization to free-abelian by free groups, and [Iva99; KWM05; Mar07; SSV16] for other generalizations. In the following lines we present Ivanov's generalization of this technique to cover the case of free products $G = G_1 * G_2$.

For all the section, we fix two groups G_1 and G_2, for which we shall assume several properties in the sequel; the goal is to understand subgroups $H \leqslant G_1 * G_2$ via some objects, which we will call reduced *wedge* automata (by similarity with the free group case). In particular, they will encode the Kurosh subgroup decomposition of H. A result saying that, for any finitely generated H, one can always algorithmically construct a reduced automaton for it, will give as a corollary an algorithmic version of the classical Kurosh subgroup theorem.

Theorem 7.4.1. *Given words $w_1, \dots, w_n \in G_1 * G_2$, one can compute words $x_1, \dots, x_m$, $y_1, \dots, y_r$, $z_1, \dots, z_s \in G_1 * G_2$, and finitely many elements $a_{i,k} \in G_1$ and $b_{j,k} \in G_2$ such that*

$$H = \mathbb{F} * \mathop{\ast}_{i=1}^{r} (y_i^{-1} A_i y_i) * \mathop{\ast}_{j=1}^{s} (z_j^{-1} B_j z_j)$$

is a Kurosh subgroup decomposition for $H = \langle w_1, \dots, w_n \rangle$, where the set $\{x_1, \dots, x_m\}$ is a free basis for $\mathbb{F}$, the $a_{i,k}$'s are generators for $A_i \leqslant G_1$, $i = 1, \dots, r$, and the $b_{j,k}$'s are generators for $B_j \leqslant G_2$, $j = 1, \dots, s$. □

7.4.1 Wedge automata

The basic idea is the following: automata recognizing subgroups of free groups, say $\mathbb{F}_2 = \langle a \rangle * \langle b \rangle$, are directed, involutive, connected graphs with labels 'a' or 'b' attached to the arcs; this way, walks spell words on $\{a, b\}^{\pm}$, i.e., they represent elements from $\mathbb{F}_{\{a,b\}}$.

Now, to cover the more general situation of $G_1 * G_2$, we need to encode more information into the arcs: A classical (say a)-labelled arc would correspond to what we call a (say G_1)-wedge: an arc subdivided in two halves, admitting a (possibly trivial) label from G_1 on each side, and also admitting a (possibly trivial) subgroup $A \leqslant G_1$ as a label of the middle (special) vertex. Doing the same with the b-arcs (and subgroups of G_2) we get an automata with two types of vertices, *primary* (the original ones), and *secondary* (the new ones, dividing the original arcs).

Fig. 7.5: Wedges of first and second kind

In these new automata, walks are going to spell *subsets* of $G_1 * G_2$ (instead of *words* in $\{a, b\}^{\pm}$), by picking *all* elements from the label of a secondary vertex when traversing it. Allowing, in addition, vertices to have any degree, we get the new notion of wedge automaton.

Definition 7.4.2. Let G_1, G_2 be two groups. A *wedge automaton* on (G_1, G_2) (also called a (G_1, G_2)*-automaton*, for short) is a septuple $\boldsymbol{\Gamma} = (V\Gamma, E\Gamma, \iota, \tau, \ell, \circledast, ^{-1})$, where:

1. $\Gamma = (V\Gamma, E\Gamma, \iota, \tau, ^{-1})$ is an involutive digraph (called *underlying digraph* of $\boldsymbol{\Gamma}$) with:

   1.1. the set of vertices $V\Gamma$ consisting of three disjoint types,

   $$V\Gamma = V_0\Gamma \sqcup V_1\Gamma \sqcup V_2\Gamma .$$

   Namely,

   1.1.1. *primary vertices* in V_0, denoted by $\bullet$; except for a distinguished one called the *basepoint* of $\boldsymbol{\Gamma}$, denoted by $\circledast$.

   1.1.2. *secondary vertices*, of two different kinds: type 1 (or 1-secondary) vertices in V_1, denoted by $\circ$; and type 2 (or 2-secondary) vertices in V_2, denoted by $\circ$.

   1.2. Every arc in $E\boldsymbol{\Gamma}$ joining (in either direction) a primary vertex with either a 1-secondary vertex (called an arc of type 1, or 1-arc, in $E_1\boldsymbol{\Gamma}$), or a 2-secondary vertex (called an arc of type 2, or a 2-arc, in $E_2\boldsymbol{\Gamma}$), i.e.,

   $$E\boldsymbol{\Gamma} = E_1\boldsymbol{\Gamma} \sqcup E_2\boldsymbol{\Gamma} .$$

2. ℓ is a twofold *label map*, associating:

2.1. an element from G_ν to every ν-arc ($\nu = 1, 2$) of Γ,

$$\begin{aligned} \ell\colon E_\nu \Gamma &\to G_\nu \\ e &\mapsto \ell_e \end{aligned}$$

2.2. a subgroup of G_ν to every ν-secondary vertex ($\nu = 1, 2$) of Γ,

$$\begin{aligned} \ell\colon V_\nu \Gamma &\to \{\text{subgroups of } G_\nu\} \\ q &\mapsto \ell_q \end{aligned}$$

3. The involution $^{-1}\colon E\Gamma \to E\Gamma$ on the set of arcs extends naturally to an involution on the corresponding labels, i.e., $\iota e^{-1} = \tau e$, $\tau e^{-1} = \iota e$, and $\ell_{e^{-1}} = \ell_e^{-1}$.

Remark 7.4.3. Since wedge automata are assumed to be involutive by definition, on pictures we will only represent one arc (usually the one departing a secondary vertex) from each pair of inverses $\{e, e^{-1}\}$, understanding implicitly the existence of the other. In the same vein, we will understand the *degree* of a vertex in a wedge automata as the number of arcs leaving (or arriving) a vertex.

Remark 7.4.4. Note that the underlying digraph of a wedge automata is, by definition, primary-secondary bipartite, with both types of vertices admitting, in principle, any degree. We will usually represent them in bipartite form with the arcs leaving secondary vertices (see Figure 7.6).

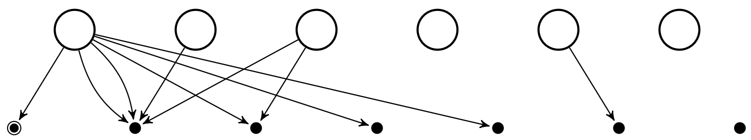

Fig. 7.6: Underlying digraph of certain wedge automaton

Definition 7.4.5. We say that a wedge automaton Γ is *connected* (resp., *finite*) if the underlying digraph $(V\Gamma, E\Gamma, \iota, \tau)$ is so. We will also say that that a wedge automaton Γ is of *finite type* if the underlying digraph is finite, and the subgroups labelling the secondary vertices are all finitely generated. This will always be the situation when we consider computational issues; in this case, the labels of vertices will be given by finite sets of generators. We say that a vertex label is *trivial* if $\ell_q = \{1\}$.

If not stated otherwise all the wedged automata appearing from this point will be assumed to be finite.

Definition 7.4.6. Let Γ be a (G_1, G_2)-wedge automaton. A *walk* in Γ is a sequence of alternating and successively incident vertices and arcs, starting and ending at primary vertices. The *length* of a walk is the number of arcs in the sequence defining it. A walk (of length 0) consisting of only a primary vertex is called a *trivial walk*.

Note that, since Γ is bipartite by definition, any walk γ in a wedge automaton Γ must alternate visits to primary and secondary vertices; so it has the form:

$$\gamma = p_0(e_1^{-1}q_1e_1')p_1(e_2^{-1}q_2e_2')p_2\cdots p_{r-1}(e_r^{-1}q_re_r')p_r \tag{7.4}$$

where $p_0,\ldots,p_r$ are (not necessarily distinct) primary vertices, $q_1,\ldots,q_r$ are (not necessarily distinct) secondary vertices, and for every $k = 1,\ldots,r$, all three of e_k, q_k, e_k' are simultaneously of the same type ν_k ($\nu_k = 1,2$); see Figure 7.7.

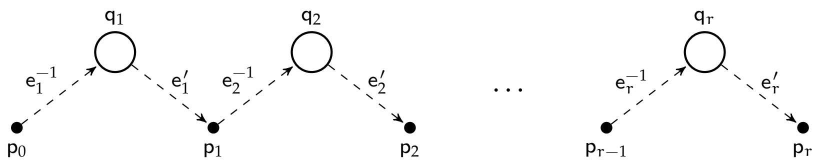

Fig. 7.7: Elementary decomposition of the walk γ in (7.4)

So, the length of a walk γ in a wedge automaton is always even ($l(\gamma) = 2r$), and any walk γ can be subdivided into a sequence of r consecutive adjacent subwalks of length 2.

Definition 7.4.7. An *elementary walk* is a walk of length 2 (hence visiting only one secondary vertex). An elementary walk is called of type ν (ν-*elementary*) if its secondary vertex is of type ν ($\nu = 1,2$).

Remark 7.4.8. An elementary walk is *degenerate* (or *backtracking*) if it consists of consecutive inverse arcs, otherwise it can be open (if it visits three different vertices, i.e., it crosses exactly one wedge in the enriched automaton), or closed (if it uses twice the same primary vertex, using reverse — but not inverse — parallel arcs). See Figure 7.8.

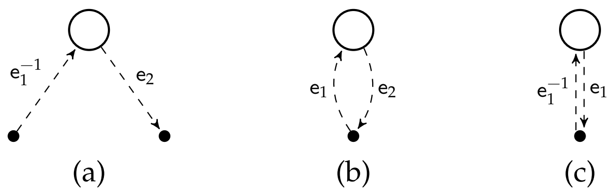

Fig. 7.8: Open (a), closed (b), and backtracking (c) elementary walks

Definition 7.4.9. Observe that every walk γ in Γ (beginning and ending at primary vertices) decomposes as a product of elementary walks (either degenerate or nondegenerate, and with possible repetitions) in a unique way

$$\gamma = \gamma_1\gamma_2\cdots\gamma_r\,; \tag{7.5}$$

(for convention, we take $r = 0$ when the walk γ is trivial). This is called the *elementary decomposition* of γ; accordingly, the number of elementary walks in the elementary decomposition of γ is called the *elementary length* of γ (see Figure 7.7). Hereinafter, unless stated otherwise, any walk decomposition of the form (7.5) will be assumed to be elementary.

Remark 7.4.10. Note that γ is a *nondegenerate walk* (i.e., it involves no backtracking) if and only if the γ_i's on its elementary decomposition are all nondegenerate, and there is no backtracking in the products $\gamma_i \cdot \gamma_{i+1}$.

Definition 7.4.11. We say that a walk γ is *alternating* if its elementary decomposition sequence $\gamma_1, \gamma_2, \ldots, \gamma_r$ alternates between types 1 and 2.

Remark 7.4.12. We emphasize that (in a general wedge automaton) the elementary decomposition of a walk *is not* necessarily alternating; i.e., it is possible to have in (7.5) consecutive elementary walks of the same type.

Definition 7.4.13. The (total) *label of a walk* γ, denoted by ℓ_γ, is defined to be the subset

$$\begin{aligned} \ell_\gamma \;&=\; \left\{ (\ell_{e_1}^{-1} h_1 \ell_{e_1'})(\ell_{e_2}^{-1} h_2 \ell_{e_2'}) \cdots (\ell_{e_r}^{-1} h_r \ell_{e_r'}) \;:\; h_k \in \ell_{q_k} \right\} \\ &=\; (g_1^{-1}\, H_1\, g')\,(g_2^{-1}\, H_2\, g_2')\; \cdots\; (g_r^{-1}\, H_r\, g_r') \;\subseteq\; G_1 * G_2\,. \end{aligned}$$

(Note that here brackets indicate just the labelling of the elementary decomposition.) The label of an elementary walk is called *elementary label*.

That is, while travelling along γ, when we traverse an arc e, we *pick* its label ℓ_e, and when we traverse a secondary vertex q we take all labels $h \in \ell_q$ (primary vertices have no contribution to ℓ_γ, and can be thought as trivially labelled: $\ell_p = 1$, for all $p \in V_0$). Collecting labels in all such possible ways, we form the label ℓ_γ.

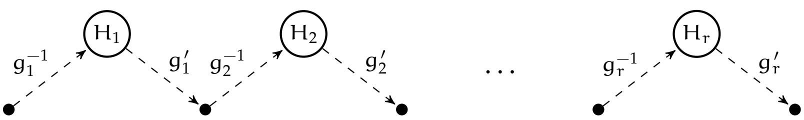

Fig. 7.9: The label of a walk with elementary length r

Two variations (subsets) of the total label will be used throughout the paper:

- Picking only the trivial element in every secondary vertex of γ we obtain what we know as the *basic label* of γ,

  $$\ell_\gamma^\bullet \;=\; (\ell_{e_1}^{-1} \ell_{e_1'})(\ell_{e_2}^{-1} \ell_{e_2'}) \cdots (\ell_{e_r}^{-1} \ell_{e_r'}) \;\in\; G_1 * G_2\,;$$

  of course, $\ell_\gamma^\bullet \in \ell_\gamma$.

- Picking only the trivial element in every secondary vertex in γ *except one*, say q_j we obtain what we call the *label of* γ *centered in* q_i,

$$\ell_{\gamma}^{q_j} = (\ell_{e_1}^{-1}\ell_{e_1'})\cdots(\ell_{e_j}^{-1}\ell_{q_j}\ell_{e_j'})\cdots(\ell_{e_r}^{-1}\ell_{e_r'}) \in G_1 * G_2\,;$$

of course, $\ell_{\gamma}^{q_j} \subseteq \ell_{\gamma}$, for all $j = 1, \ldots, r$.

Remark 7.4.14. Note that if the secondary vertices $q_1, \ldots, q_r$ successively visited by the walk γ do alternate between *type* 1 and *type* 2, and the choices of h_k's are such that $\ell_{e_i}^{-1}h_k\ell_{e_i'} \neq 1$ for all $i = 1, \ldots, r$, then the brackets in the expression

$$(\ell_{e_1}^{-1}h_1\ell_{e_1'})(\ell_{e_2}^{-1}h_2\ell_{e_2'})\cdots(\ell_{e_r}^{-1}h_r\ell_{e_r'})$$

indicate, precisely, the syllable decomposition of the element in $G_1 * G_2$ read by γ; otherwise, some consecutive pairs of brackets may merge into the same syllable. So, in general, the syllable length of every element in ℓ_γ is less than or equal to its elementary length $r = l(\gamma)/2$, with equality if and only if the previous two conditions are satisfied.

Definition 7.4.15. Let Γ be a (G_1, G_2)-wedge automaton and let p, p' be two connected primary vertices. We define the *coset recognized* by Γ relative to (p, p') to be the set:

$$\langle \Gamma \rangle_{(p,p')} \coloneqq \bigcup\nolimits_{\gamma} \ell_{\gamma}\,,$$

where the union runs over all walks in Γ from p to p'. When $p = p'$, then we abbreviate $\langle \Gamma \rangle_p \coloneqq \langle \Gamma \rangle_{(p,p)}$. Moreover, if $p = p' = \circledast$ (the basepoint of Γ) then we simply write $\langle \Gamma \rangle = \langle \Gamma \rangle_{\circledast}$, and we call it the *subgroup recognized* by Γ.

The lemma below justifies the previous terminology.

Lemma 7.4.16. *Let* Γ *be a* (G_1, G_2)*-wedge automaton, and let* $p, p' \in V_0\,\Gamma$*. Then,*

(i) *the set* $\langle \Gamma \rangle_p$ *is a subgroup of* $G_1 * G_2$*;*

(ii) *the sets* $\langle \Gamma \rangle_p$ *and* $\langle \Gamma \rangle_{p'}$ *are conjugate to each other; more concretely,*

$$\langle \Gamma \rangle_{p'} = \langle \Gamma \rangle_p^g\,, \textit{ for every } g \in \langle \Gamma \rangle_{(p,p')}\,;$$

(iii) *the set* $\langle \Gamma \rangle_{(p,p')}$ *is a right coset of* $\langle \Gamma \rangle_{(p,p)}$*; more concretely,*

$$\langle \Gamma \rangle_{(p,p')} = \langle \Gamma \rangle_{(p,p)} \cdot g\,, \textit{ for every } g \in \langle \Gamma \rangle_{(p,p')}\,.$$

Proof. For subsets $A, B \subseteq G_1 * G_2$, define the inverse and the product in the natural way $A^{-1} = \{a^{-1} \mid a \in A\}$, and $A \cdot B = \{ab \mid a \in A, b \in B\}$. Then, for α and β two incident walks in Γ, it is clear that $\ell_{\alpha^{-1}} = \ell_{\alpha}^{-1}$ and $\ell_\alpha \cdot \ell_\beta = \ell_{\alpha\beta}$. The lemma follows straightforward from Definition 7.4.13. □

Proposition 7.4.17. *For every subgroup* $H \leqslant G_1 * G_2$*, there exists a* (G_1, G_2)*-wedge automaton* Γ *recognizing* H*. Furthermore, if* H *is finitely generated, one such* Γ *is algorithmically constructible from a finite set of generators for* H *given in normal form.*

Proof. Let $H = \langle W \rangle$, where $W = \{w_1, w_2, \ldots\}$ is a (finite or infinite) set of generators for H. For every non-trivial generator in W, say w, consider its normal form as an element of $G_1 * G_2$, say

$$w = a_1 b_1 \cdots a_s b_s \,, \tag{7.6}$$

with $s \geqslant 1$, $a_i \in G_1$, $b_i \in G_2$, $a_i \neq 1$ for all $i = 2, \ldots, s$, and $b_i \neq 1$ for all $i = 1, \ldots, s-1$. Let $\mathsf{Fl}(w)$ denote the (G_1, G_2)-wedge automaton depicted in Fig. 7.10, and called the *(wedge) petal* corresponding to w. Clearly, $\langle \mathsf{Fl}(w) \rangle = \langle w \rangle$.

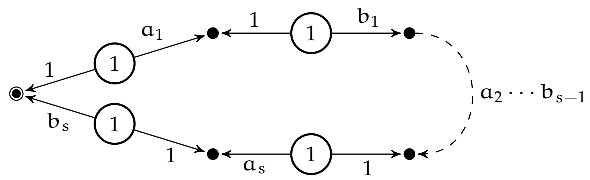

Fig. 7.10: A wedge petal

Now consider $\mathsf{Fl}(W)$ the disjoint union of all the $\mathsf{Fl}(w_i)$'s identifying the basepoints into a single primary vertex (declared as basepoint); the resulting object is a (G_1, G_2)-wedge automaton called the *flower (wedge) graph* corresponding to W and clearly satisfying $\langle \mathsf{Fl}(W) \rangle = \langle W \rangle = H$. Moreover, $\mathsf{Fl}(W)$ is of finite type if and only if W has finite cardinal; and, in this case, it is clearly constructible from the given w_i's. □

So, we have a well defined surjective map

$$\begin{array}{rcl} \{\, (G_1, G_2)\text{-wedge automata} \,\} & \to & \{\, \text{subgroups of } G_1 * G_2 \,\} \\ \Gamma & \mapsto & \langle \Gamma \rangle \,. \end{array} \tag{7.7}$$

However, in order to describe appropriately the behavior of the subgroups of $G_1 * G_2$ (according our target), we need to restrict the class of wedge graphs representing a given subgroup. This is what we do in the next section through the concept of *reduced* wedge automata.

7.4.2 Reduced wedge automata

Wedge automata are too general, and just a first step in order to understand subgroups of free products $G_1 * G_2$ (they are only geometric realizations of sets of generators for subgroups, as the previous proposition illustrates).

The really useful objects to understand subgroups of $G_1 * G_2$ will be the so-called *reduced* wedge graphs. These are wedge graphs enjoying some extra regularity properties, good enough to extract from them enough algebraic information to characterize the subgroup they represent.

Essentially, we will ask our wedge graphs to be 'deterministic' in a sense that will be precisely specified in Definition 7.4.18, and naturally extends that of classical Stallings automata for free groups.

On the other hand, we shall detail how to algorithmically build such a reduced automaton for a subgroup $H \leqslant G_1 * G_2$, from any given finite set of generators for H. Similar constructions are called 'irreducible graphs' by Ivanov in [Iva99], and are particular cases of the so-called 'folded graphs' in Kapovich–Weidman–Miasnikov [KWM05].

Definition 7.4.18. Let G_1, G_2 be two groups, and let Γ be a finite (G_1, G_2)-wedge automaton. We say that Γ is *reduced* if the following conditions are satisfied:

(i) Γ is connected;

(ii) every primary vertex of Γ is incident with at most one arc from $E_1\Gamma$, and at most one arc from $E_2\Gamma$. That is, it is adjacent to at most one 1-secondary vertex and at most one 2-secondary vertex, being connected to each of them by at most one arc;

Fig. 7.11: Forbidden "nondeterminstic" situations in a reduced automaton

(iii) no nondegenerate elementary walk reads the trivial element. That is, for $\nu = 1, 2$, and every ν-secondary vertex $q \in V_\nu\Gamma$, and for every pair of *different* ν-edges e_1, e_2 with $\iota e_1 = \iota e_2 = q$, we have that $1 \notin \ell_{e_1}^{-1}\ell_q\ell_{e_2}$ (or, equivalently, $\ell_{e_1}\ell_{e_2}^{-1} \notin \ell_q$).

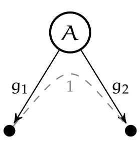

$(1 \in g_1^{-1}Ag_2)$.

Fig. 7.12: Forbidden nondegenerate wedge reading the trivial element

Remark 7.4.19. Note that in a reduced wedge automaton, the degree of any primary vertex is either 1 (corresponding to having only one adjacent secondary vertex) or 2 (corresponding to having an adjacent secondary vertex of each type). Therefore, the elementary decomposition of any walk in a reduced automaton is always alternating, and spell the syllable decomposition in $G_1 * G_2$ of the read element. Secondary vertices, on its turn, can have any degree in a reduced reduced wedge automaton.

Assume now that Γ is a reduced wedge automaton. Then, Remark 7.4.10 can be restated in the following way: γ is nondegenerate if and only if it is alternating and the elementary walks on its elementary decomposition are all nondegenerate. In this case, additionally, property (iii) from Definition 7.4.18 ensures that the elementary decomposition of γ gives the syllable decomposition of every $g \in \ell_\gamma$ as element from $G_1 * G_2$ (because nondegenerate elementary walks do not admit the trivial element as a label). However, this is not the whole story: even with some of the γ_i's being degenerate, we can still get the syllable decomposition of $g \in \ell_\gamma$ assuming that the elements picked from the labels of the backtracking vertices are all non-trivial. This motivates the following definition and the subsequent important technical lemma.

Definition 7.4.20. Let γ be be a walk in a wedge automaton Γ, with elementary decomposition $\gamma = \gamma_1\gamma_2\cdots\gamma_r = p_0(e_1^{-1}q_1e_1')p_1(e_2^{-1}q_2e_2')p_2\cdots p_{r-1}(e_r^{-1}q_re_r')p_r$. We define the *reduced label* of γ as

$$\widetilde{\ell}_\gamma = \left\{ (\ell_{e_1}^{-1}c_1\ell_{e_1'})(\ell_{e_2}^{-1}c_2\ell_{e_2'})\cdots(\ell_{e_r}^{-1}c_r\ell_{e_r'}) \ : \ \begin{array}{l} c_i \in \ell_{q_i} \\ c_i \neq 1 \text{ if } \gamma_i \text{ is degenerate} \end{array} \right\}$$

Of course, $\widetilde{\ell}_\gamma \subseteq \ell_\gamma$.

Lemma 7.4.21. *For a reduced* (G_1, G_2)*- automaton* Γ*, we have*

$$\langle\Gamma\rangle = \bigcup_\gamma \ell_\gamma = \bigcup_{\widehat{\gamma}} \widetilde{\ell}_{\widehat{\gamma}},$$

where the first union runs over all ⊙-walks γ *of* Γ*, and the second one only over the alternating ⊙-walks* $\widehat{\gamma}$ *of* Γ*.*

Proof. The inclusion '$\supseteq$' is clear, since the first union is over more sets than the second one, and $\ell_\gamma \supseteq \widetilde{\ell}_\gamma$.

To see the inclusion '$\subseteq$', fix a ⊙-walk $\gamma = \gamma_1\cdots\gamma_r$, and an element $g \in \ell_\gamma$, and let us find an alternating ⊙-walk $\widehat{\gamma}$ such that $g \in \widetilde{\ell}_{\widehat{\gamma}}$.

In fact, if γ is not alternating there is $\nu = 1, 2$ and $i = 1, \ldots, r-1$ such that $\gamma_i = p_{i-1}e_i^{-1}q_ie_i'p_i$ and $\gamma_{i+1} = p_ie_{i+1}^{-1}q_{i+1}e_{i+1}'p_{i+1}$ are both of type ν; so, by condition (ii) in Definition 7.4.18, $q_i = q_{i+1}$ and $e_i' = e_{i+1}$. Replacing $\gamma_i\gamma_{i+1}$ by $= p_{i-1}e_i^{-1}q_ie_{i+1}'p_{i+1}$, we get a new ⊙-walk $\gamma_{(1)}$, with shorter elementary decomposition and such that $g \in \ell_{\gamma_{(1)}}$ as well, since $(\ell_{e_i}^{-1}c_i\ell_{e_i'})(\ell_{e_{i+1}}^{-1}c_{i+1}\ell_{e_{i+1}'}) = \ell_{e_i}^{-1}(c_ic_{i+1})\ell_{e_{i+1}'}$, for all $c_i, c_{i+1} \in \ell_{q_i} \leqslant G_\nu$).

Repeating this operation a finite number of times, say k ,we can assume that $\widehat{\gamma} = \gamma_{(k)}$ is alternating (and $g \in \ell_{\widehat{\gamma}}$).

It remains to prove that $g \in \widetilde{\ell}_{\widehat{\gamma}}$, i.e., that it can be obtained by picking always non-trivial elements at all the backtracking vertices of $\widehat{\gamma} = \gamma_1 \cdots \gamma_s$: if q_j is the secondary vertex in the degenerate elementary walk $\gamma_j = p_{j-1} e_j^{-1} q_j e_j p_{j-1}$, and the corresponding c_j picked in the formation of g is trivial, then just ignore $(\ell_{e_j}^{-1} c_j \ell_{e_j}) = 1$, and realize g in the label of $\widehat{\gamma}_{(1)} = \gamma_1 \cdots \gamma_{j-1} \gamma_{j+1} \cdots \gamma_s$, a ⊙-walk with shorter elementary decomposition, which will be again alternating after repeating the operation in the third paragraph. Repeating this operations (a finite number of times) until no trivial choices are made at the degenerate vertices, we obtain the desired result. □

Remark 7.4.22. The usefulness of the previous lemma is the following: when realizing an element $g \in \langle \Gamma \rangle$ from the recognized subgroup of a reduced automaton X as $g \in \widetilde{\ell}_{\widehat{\gamma}}$ for some alternating ⊙-walk $\widehat{\gamma}$, the elementary decomposition of $\widehat{\gamma}$ automatically provides the syllable decomposition of g as an element of $G_1 * G_2$. This is a crucial bridge between the algebraic and the geometric aspects of the theory.

One of the most useful facets of reduced (G_1, G_2)-automata is that they naturally encode the Kurosh free product decomposition (the induced splitting) of their fundamental group as subgroup of $G_1 * G_2$. With some technical differences, our exposition follows [Iva99; KWM05], but with an special emphasis made on the algorithmic point of view.

The theorem below appears as Lemma 4 in Ivanov's [Iva08] and, in a more general setting, as Proposition 4.3 in Kapovich–Weidmann–Myasnikov [KWM05]; we highlight here the algorithmic nature of the proof, for Γ of finite type. Note that, after Theorem 7.4.30, realizing every finitely generated subgroup $H \leqslant G_1 * G_2$ as recognized by some finite reduced (G_1, G_2)-automaton, we obtain an alternative *constructive* proof of the classical Kurosh subgroup theorem, for finitely generated subgroups.

Notation 7.4.23. Let Γ be a reduced (G_1, G_2)-wedge automaton, and let $H = \langle \Gamma \rangle \leqslant G_1 * G_2$ be its recognized subgroup. Fix a maximal subtree $\mathbf{T}$ in Γ, let $E = E\Gamma \setminus E\mathbf{T}$ be the set of edges of Γ outside $\mathbf{T}$, and let E^+ be a subset of E containing exactly one element of each pair of mutually inverses $e, e^{-1} \in E$.

For every arc $e \in E^+$, let q be the secondary vertex in Γ incident with e. Then:

(a) the subautomaton of Γ consisting of e together with the tree segments from the basepoint to the head and tail of e is called the *(primary)* e*-petal of* $\mathbf{T}$, and denoted by $\mathbf{T}[e]$.

(b) the subautomaton of Γ consisting of the tree segment from the basepoint to q is called the *(secondary)* q*-petal of* $\mathbf{T}$, and denoted by $\mathbf{T}[q]$.

Slightly abusing terminology and notation, we also call (respectively primary and secondary) $\mathbf{T}$*-petals* the closed walks, denoted by $\gamma_{\mathbf{T}}[e]$ and $\gamma_{\mathbf{T}}[q]$, respectively touring the petals $\mathbf{T}[e]$ and $\mathbf{T}[q]$, namely:

$$\gamma_{\mathbf{T}}[e] \coloneqq \gamma_{\mathbf{T}}[\circledbullet, \iota e]\, e\, \gamma_{\mathbf{T}}[\tau e, \circledbullet]\,, \quad \text{and} \quad \gamma_{\mathbf{T}}[q] \coloneqq \gamma_{\mathbf{T}}[\circledbullet, \iota q]\, \gamma_{\mathbf{T}}[q, \circledbullet]\,.$$

Definition 7.4.24. We define the *petal label* x_e of a nondegenerate petal $\gamma_{\mathbf{T}}[e]$ to be its basic label; whereas the *petal label* $(H_q)^{z_q}$ of a degenerate petal $\gamma_{\mathbf{T}}[q]$ is defined to be the label of $\gamma_{\mathbf{T}}[q]$ centered in q. That is:

$$x_e \coloneqq \ell^{\bullet}_{\gamma_e} = z_{\iota e}^{-1} \ell_e\, z_{\tau e}\,, \quad \text{and} \quad (H_q)^{z_q} \coloneqq \ell^{q}_{\gamma_{\mathbf{T}}[q]} = (\ell_{\gamma_{\mathbf{T}}[q]})^{z_q} = z_q^{-1} \ell_q\, z_q\,,$$

where, for any (primary or secondary) vertex p in Γ we have denoted by z_q the basic label $\ell^{\bullet}_{\mathbf{T}}[q, \circledbullet] \in G_1 * G_2$.

Finally, a *reduced secondary-petal label* is

$$\ell^{q*}_{\gamma_{\mathbf{T}}[q]} \coloneqq z_q^{-1} y_q\, z_q\,,$$

where $y \in \ell_q \setminus \{1\} = \ell_q^*$. (See Figure 7.13.)

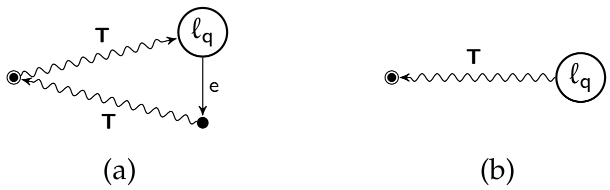

Fig. 7.13: (a) Primary $\mathbf{T}$-petal $\mathbf{T}[e]$ with petal label $x_e = z_{\iota e}^{-1} \ell_e\, z_{\tau e} \in G_1 * G_2$ (b) Secondary $\mathbf{T}$-petal $\mathbf{T}[q]$ with petal label $z_q^{-1} \ell_q\, z_q \leqslant G_1 * G_2$

Theorem 7.4.25. *Let Γ be a reduced (G_1, G_2)-automaton. Then, with the above notations, the subgroup recognized by Γ is*

$$\langle \Gamma \rangle = \mathbb{F} * \left(\mathop{\ast}_{q \in V_1\Gamma} z_q^{-1}\, \ell_q\, z_q \right) * \left(\mathop{\ast}_{q \in V_2\Gamma} z_q^{-1}\, \ell_q\, z_q \right), \tag{7.8}$$

where $\mathbb{F}$ is the free subgroup of $\langle \Gamma \rangle$ freely generated by the set $\{z_{\iota e}^{-1} \ell_e\, z_{\tau e} \mid e \in E^+\}$, and $\ell_q \leqslant G_\nu$ are the labels of the corresponding secondary vertices (of type $\nu = 1, 2$) in Γ.

Moreover, if Γ is of finite type, then the subgroup $\langle \Gamma \rangle$ is finitely generated, and we can algorithmically compute a Kurosh decomposition like (7.8).

Proof. Adapting the characterization in [Mil04, Theorem 3.2] to our case, we need to prove the following two facts:

(i) *The elements in*

$$\{ z_{\iota e}^{-1} \ell_e z_{\tau e} \mid e \in E^+ \} \sqcup \bigsqcup_{q \in V_1 \Gamma} z_q^{-1} \ell_q z_q \sqcup \bigsqcup_{q \in V_2 \Gamma} z_q^{-1} \ell_q z_q \tag{7.9}$$

generate the recognized subgroup $\langle \Gamma \rangle \leqslant G_1 * G_2$. *(Note that indeed, this claim holds even if* Γ *is not reduced.)*

This is clear because since for any ◉-walk

$$\gamma = p_0 (e_1^{-1} q_1 e_1') p_1 (e_2^{-1} q_2 e_2') p_2 \cdots \cdots p_{r-1} (e_r^{-1} q_r e_r') p_r$$

in Γ (where e_i, e_i' are arcs in $E\Gamma$), the associated walk γ_T (going to the basepoint and back — through T — after visiting each element in γ), i.e.,

$$\gamma_T = \gamma_T[e_1^{-1}]\, \gamma_T[q_1]\, \gamma_T[e_1']\, \gamma_T[e_2^{-1}]\, \gamma_T[q_2]\, \gamma_T[e_2'] \;\cdots\; \gamma_T[e_r^{-1}]\, \gamma_T[q_r]\, \gamma_T[e_r'] \,,$$

has the same label as γ. So, any element recognized by γ is of the form:

$$x_{e_1}^{-1}\, (z_{q_1}^{-1} y_{q_1} z_{q_1})\, x_{e_1'}\, x_{e_2^{-1}}\, (z_{q_2}^{-1} y_{q_2} z_{q_2})\, x_{e_2'} \cdots x_{e_r^{-1}}\, (z_{q_r}^{-1} y_{q_r} z_{q_r})\, x_{e_r'} \,, \tag{7.10}$$

where $y_{q_i} \in \ell_{q_i}$, for $i = 1, \ldots, r$. Finally note that the only nontrivial elements $x_{e_i}, x_{e_i'}$ in (7.10), are those corresponding to arcs outside T. Thus, (7.10) is a product of elements in (7.9) and their inverses, as we wanted to prove.

(ii) *If* $w = \prod_{i=1}^{r} g_i$ *is an alternating word in the generators in* (7.9) *representing the trivial element in* $\langle \Gamma \rangle$*, then there exists a syllable* g_j *in* w *which is trivial in its corresponding factor in* (7.8).

By transposition, suppose that every syllable g_j $(i = 1, \ldots, r)$ in w is nontrivial in its corresponding factor. That is, suppose:

1. $w = \prod_{i=1}^{r} g_i$;
2. every syllable g_i of w is either
   (a) equal to a nontrivial power of $z_{\iota e}^{-1} \ell_e z_{\tau e}$, for some $e \in E^+$, or
   (b) equal to a reduced petal label $z_q^{-1} y_q z_q$, where $y_q \in \ell_q \setminus \{1\} = \ell_q^*$, for certain secondary vertex q in Γ;

Then, the set

$$\{ \langle z_{\iota e}^{-1} \ell_e z_{\tau e} \rangle \,:\, e \in E^+ \} \cup \{ z_q^{-1} \ell_q^* z_q \,:\, q \in V_1 \cup V_2 \} \tag{7.11}$$

is called the set of *syllable types*.

3. no consecutive syllables in w are of the same type in (7.11).

So, disregarding order (which does not affect our next argumentation) we have the following three cases for petal labels separating consecutive syllables in w:

(a) Consecutive labels of *different-type* ‘primary-petals’:

$$(z_{\iota e_1}^{-1}\ell_{e_1}\, z_{\tau e_1})(z_{\iota e_2}^{-1}\ell_{e_2}\, z_{\tau e_2})\,, \tag{7.12}$$

where $e_1 \neq e_2^{\pm 1}$. This automatically implies that (7.12) is the basic label of two *distinct* consecutive primary petals in Γ (see Figure 7.14). In particular, every possible cancellation in (7.12) occurs within the tree $\mathbf{T}$.

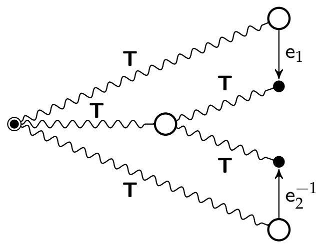

Fig. 7.14: Distinct consecutive primary petals separating syllables in w

(b) Consecutive reduced labels of *different-type* secondary-petals:

$$(z_{q_1}^{-1}\, y_{q_1}\, z_{q_1})(z_{q_2}^{-1}\, y_{q_2}\, z_{q_2})\,, \tag{7.13}$$

where $y_{q_1} \in \ell_{q_1}^*$, $y_{q_2} \in \ell_{q_2}^*$, and, $q_1 \neq q_2$.

Here we distinguish two subcases depending on whether: the bifurcation vertex is primary (b1) (note that in this case it must be the basepoint, since Γ is reduced by hypothesis), or secondary (b2);

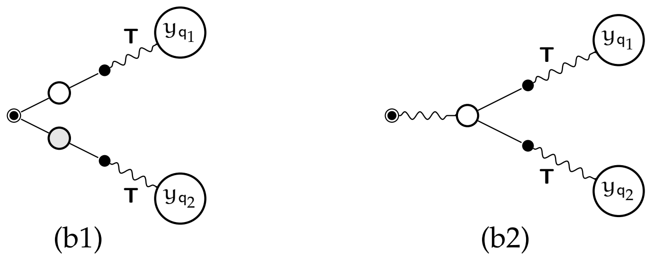

Fig. 7.15: Consecutive primary petals separating alternating syllables in w, with secondary (b2), and primary (b1) bifurcation vertex

Now we claim that in (7.13), the product $z_{q_1} z_{q_2}^{-1}$ is always nontrivial; and hence every possible cancellation in (7.13) occurs within the tree $\mathbf{T}$.

If the bifurcation vertex is primary, we know — since Γ is reduced — that the two secondary vertices adjacent to the basepoint must be of different type (see case (b1) in Figure 7.15). But, since any walk in a reduced wedge automaton is alternating, this means that the first syllables of the labels of z_{q_1} and z_{q_2} must belong to different G_i's. Hence, $z_{q_1} \neq z_{q_2}$, as claimed.

If the bifurcation vertex is secondary (case (b2) in Figure 7.15), suppose by contraposition, that $z_{q_1} = z_{q_2}$. Then, in particular, the basic labels z'_{q_1}, z'_{q_2} of the subwalks from the bifurcation vertex to q_1 and q_2 would coincide (i.e., $z'_{q_1}(z'_{q_2})^{-1} = 1$). But, again, since any walk in a reduced automata must be alternating, this would imply that the elementary walk from q_1 to q_2 containing the bifurcation vertex should be trivial as well, contradicting condition (iii) in the definition of reduced wedge automaton.

(c) (Reduced) secondary petal label followed by primary petal label:

$$(z_{\iota e}^{-1} \ell_e \, z_{\tau e})(z_q^{-1} \, y_q \, z_q)\,, \tag{7.14}$$

where $y_q \in \ell_q^*$.

Here we distinguish three cases depending on whether the bifurcation vertex is primary (note that then it must be the basepoint and the two petals must be overlapping (see (c1)); or secondary, where — in turn — we distinguish between overlapping petals (c2), and non-overlapping petals (c3).

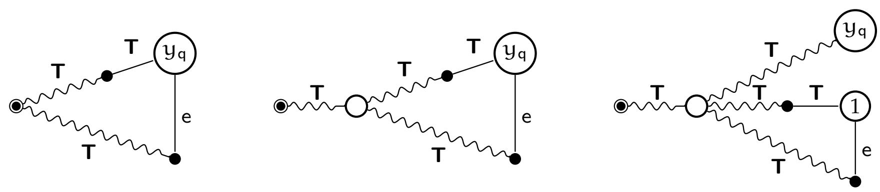

(c1) (c2) (c3)

Again, we claim that in all three cases, all cancellations in (7.14) occur within the tree $\mathbf{T}$.

For cases (c1) and (c2) the same argument used for (b2) works: namely, if (7.14) were trivial, then we would have an elementary walk (centered in q) reading the trivial element, which is not possible in a reduced automaton.

Finally, case (c3) easily reduces to case (b2), since all the petal syllables are assumed to be nontrivial by hypothesis.

Thus, since any possible cancellation between syllables in w occurs within the tree $\mathbf{T}$, a fragment of every syllable survives in w, which therefore is nontrivial, as we wanted to prove. This concludes the proofs of (ii), and the theorem. □

Corollary 7.4.26. *For a reduced (G_1, G_2)-automaton Γ, the group $\langle\Gamma\rangle$ is finitely generated if and only if Γ is of finite type (i.e., the underlying graph of Γ has finite rank, and all vertex labels of Γ are finitely generated).*

Proof. From (7.8) in Theorem 7.4.25, the recognized subgroup $\langle\Gamma\rangle$ is finitely generated if and only if both $\mathbb{F}$ and all the ℓ_q's are finitely generated, which are precisely the two conditions for the wedge automata Γ to be of finite type. □

7.4.3 Effective construction of reduced automata

Theorem 7.4.25 and Corollary 7.4.26 are easily seen to be false if we substitute reduced automata by general wedge automata.

This is why the conditions in the definition of reduced automata are important. The next step is to show that every finitely generated subgroup $H \leqslant G_1 * G_2$ can be realized as the fundamental group of some reduced (G_1, G_2)-automaton Γ of finite type. Therefore, the contents of Theorem 7.4.25 and Corollary 7.4.26 applies to *every* finitely generated subgroup of $G_1 * G_2$. Additionally, and more importantly, the construction of Γ will be made algorithmic from any given finite set of generators for H.

To achieve this goal, we introduce several elementary operations on wedge graphs which will not change their fundamental group. All the arc-transformations are stated only for one of the orientations in Γ, assuming the corresponding transformation on its involutive counterpart.

Definition 7.4.27. Let us consider the following elementary transformations ($\Gamma \rightsquigarrow \Gamma'$) on a (G_1, G_2)-wedge automaton:

(i) *Adjustment*: consists in replacing the label of any arc e leaving a secondary vertex q, by $h \cdot \ell_e$, for any $h \in \ell_q = H$; see Figure 7.16.

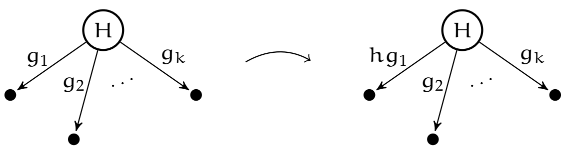

Fig. 7.16: Adjustment

(Note, in particular, that any arc-label belonging to the subgroup it reaches is negligible.)

(ii) *Conjugation*: (for $\nu = 1, 2$) consists in replacing, given $g \in G_\nu$, the label ℓ_p of a ν-secondary vertex q, by $g\ell_p g^{-1}$; and replacing the label ℓ_{e_i} of every arc e_i incident from p, by the respective $g\ell_{e_i}$ (see Figure 7.17).

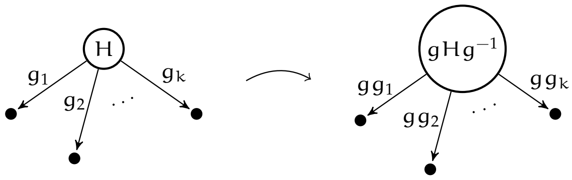

Fig. 7.17: Conjugation

(iii) *Isolation*: consisting in removing from Γ all the connected components not containing the basepoint. That is, we define Γ' as the connected component of Γ containing the basepoint (see Figure 7.17).

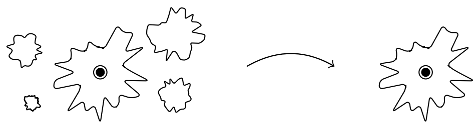

Fig. 7.18: Isolation

(iv) *Primary open folding*: for $\nu = 1, 2$, given two ν-secondary vertices q_1, q_2 adjacent to the same primary vertex through respective arcs e_1, e_2 with the same label $g \in G_\nu$, consists in identifying q_1 and q_2 into a new secondary vertex with label $\langle \ell_{q_1}, \ell_{q_2} \rangle$, and identifying the arcs e_1, e_2 into a new arc with the same label g (see Figure 7.19).

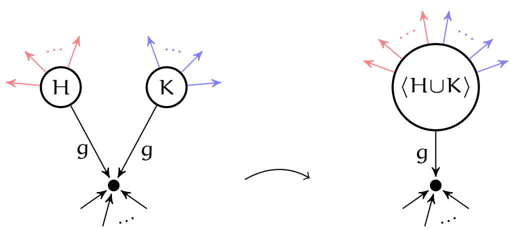

Fig. 7.19: Primary open folding

(v) *Secondary open folding:* for $\nu = 1, 2$, given a ν-secondary vertex q adjacent to two different primary vertices p_1, p_2 through arcs e_1, e_2 having the same label $g \in G_\nu$, consists in identifying the vertices p_1 and p_2, and the arcs e_1, e_2 into an arc maintaining the label g (see Figure 7.20).

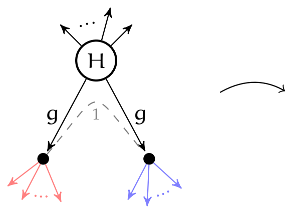

Fig. 7.20: Secondary open folding

(vi) *Closed folding*: for $\nu = 1, 2$, given a primary vertex p adjacent to a ν-secondary vertex q by two different (parallel) edges e_1, e_2, consists in identifying e_1 and e_2 into a single arc with label ℓ_{e_1}, and change the label of q from ℓ_q to $\langle \ell_q, \ell_{e_1}\ell_{e_2}^{-1} \rangle \leqslant G_\nu$ (see Figure 7.21).

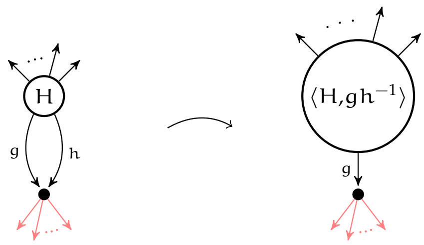

Fig. 7.21: Closed folding

(Of course, in any transformation involving arcs, the corresponding transformation preserving involution is assumed in the respective inverse arcs.)

Remark 7.4.28. The (folding) transformations (iv), (v), and (vi) in Definition 7.4.27 decrease the number of arcs in the automata exactly by 1.

The following result is straightforward to check, and we leave the details to the reader.

Lemma 7.4.29. *Let* Γ *be a reduced* (G_1, G_2)*-automaton, and let* $\Gamma \rightsquigarrow \Gamma'$ *be any of the elementary transformations in Definition 7.4.27. Then, the recognized subgroups of* Γ *and* Γ' *do coincide, i.e.,* $\langle \Gamma \rangle = \langle \Gamma' \rangle \leqslant G_1 * G_2$. □

Using these elementary operations appropriately, we can already give a constructive proof of the existence of reduced (G_1, G_2)-automata representing any given finitely generated subgroup $H \leqslant G_1 * G_2$.

Theorem 7.4.30 (Ivanov, 1999, [Iva99])**.** *For any groups* G_1, G_2 *groups, and any finitely generated* $H \leqslant G_1 * G_2$*, there exists a reduced* (G_1, G_2)*-wedge automaton recognizing* H.

Moreover, if both G_1 *and* G_2 *have solvable membership problem, then given a finite set of generators of a (finitely generated) subgroup* $H \leqslant G_1 * G_2$*, one can algorithmically construct a reduced* (G_1, G_2)*-automaton of finite type recognizing* H.

Proof. It is clear that the (trivial) wedge automata consisting only of a primary vertex (being the basepoint) and no secondary vertices nor edges is reduced (G_1, G_2)-automaton of finite type recognizing the trivial group. So, for the rest of the proof, we can assume $H \neq 1$.

Given a finite set of generators $W = \{w_1, \ldots, w_n\}$ for H, consider the flower automaton $\mathsf{Fl}(W)$; as seen above, this is a connected (G_1, G_2)-wedge automaton of finite type recognizing H. In order to gain properties (ii) and (iii) from Definition 7.4.18, we now combine elementary transformations to clean Γ in the following three ways:

(I) if, for $\nu = 1, 2$, Γ has a primary vertex p adjacent to two different ν-secondary vertices, then we apply a *corrected primary open folding* at p: namely, we first apply a suitable conjugation transformation to (say) the second ν-vertex in order to equalize the involved arc labels, an then apply an primary open folding to identify both arcs.

(II) if, for $\nu = 1, 2$, Γ has a primary vertex p and a ν-secondary vertex q connected by two different arcs e_1, e_2. then we apply a *closed folding* to identify the two arcs into a single one.

(III) if Γ has a secondary vertex q adjacent to two different primary vertices through arcs e_1 and e_2 (of the same type), say with $\iota e_1 = \iota e_2 = q$, and such that $1 \in \ell_{e_1}^{-1} \ell_q \ell_{e_2}$ (i.e., $\ell_{e_1} \ell_{e_2}^{-1} \in \ell_q$), then we apply a *corrected secondary open folding* at q: i.e., we first apply a suitable adjustment transformation to (say) the second arc e_2 in order to equalize the involved arc labels, an then apply an elementary open folding to identify e_1 and e_2 into an arc with label ℓ_{e_1}.

Observe that we can algorithmically recognize whether each of these situations occur: (I) and (II) are clear, whereas (III) can be detected using the solution to the membership problem for G_1 and G_2 which we are assuming in the hypothesis.

Therefore, the cleaning procedure is straightforward: successively detect any of these situations, and then apply the corresponding cleaning move. Note that although this could create new instances of any of the three situations to be fixed, the total number of arcs always decrease after applying any of the cleaning operations. So, no matter in which order do we perform the operations, since the starting flower automaton is finite, the process will terminate after a finite number of steps, giving as a final output a (G_1, G_2)-wedge automaton Γ' where the situations (I)-(III) do not occur anymore. This cleaning procedure is called a *folding process*.

We claim that the result of this process (Γ') is a reduced (G_1, G_2)-automaton of finite type recognizing the same subgroup $H = \langle \Gamma \rangle$, so proving the result.

Indeed, by construction, Γ' is connected. By (I) and (II), for $\nu = 1, 2$, any primary vertex of Γ' is adjacent to at most one ν-secondary vertex through at most one arc; therefore, Γ' satisfies (ii). On the other hand, by (III), Γ' satisfies property (iii). Hence, Γ' is a reduced (G_1, G_2)-automaton.

Moreover, the underlying graph of Γ' is clearly finite, and the labels of vertices are finitely generated because they started being trivial, and passed finitely many times through the two operations 'adding a generator', and 'merging two sets of generators into one'. Hence, Γ' is of finite type.

Since elementary operations preserve the recognized subgroup (by Lemma 7.4.29), we have $\langle \Gamma' \rangle = \langle \Gamma \rangle = H$. The proof is completed. □

Remark 7.4.31. The reduced (G_1, G_2)-automaton constructed in the proof of Theorem 7.4.30 can depend, in principle, on the order in which we perform the operations to $\mathsf{Fl}(W)$, and on the initial set of generators W for H. With a bit more of technical work, it is possible to canonically associate to H a *unique* reduced (G_1, G_2)-automaton Γ_H such that $\langle \Gamma_H \rangle = H$ (which will be the common result of the above sequence of operations, no matter in which order we performed them, and independently from the initial W as well). We do not develop these details here because we shall not need this uniqueness along the present paper.

7.4.4 A reduced automaton for the intersection

We now move on the main topic of this second part of the chapter; namely, intersections of subgroups of free products. The goal of this subsection is to detail an algorithmic procedure to build an reduced automaton recognizing the intersection of two finitely generated subgroups of a free product $G = G_1 * G_2$.

Recall that if G_1 and G_2 are Howson, then $G_1 * G_2$ is Howson as well (this was first proved by Baumslag in [Bau66], whereas later, Ivanov [Iva99] gave an alternative proof, which is essentially the one we present here). However, a free product $G_1 * G_2$ can very well contain finitely generated subgroups $H, K \leqslant G_1 * G_2$ such that $H \cap K$ is *not* finitely generated.

Below, we algorithmize the argument given in [Iva99] (generalizing the classical "pullback" technique for free groups) to describe a reduced (G_1, G_2)-automaton $\Gamma_H \wedge \Gamma_K$ such that $\langle \Gamma_H \wedge \Gamma_K \rangle = H \cap K$, in terms of given reduced (G_1, G_2)-automata Γ_H, with $\langle \Gamma_H \rangle = H$; and Γ_K, with $\langle \Gamma_K \rangle = K$. Note that this construction is not algorithmic since $\Gamma_H \wedge \Gamma_K$ may very well be not of finite type, even when both Γ_H and Γ_K are so (this situation corresponds to the case where H, K are both finitely generated but $H \cap K$ is infinitely generated, see Corollary 7.4.26).

Later, we shall restrict ourselves to automata of finite type (Γ_H and Γ_K) and, for this case, we shall give an effective procedure which starts constructing, locally, the aforementioned junction automaton $\Gamma_H \wedge \Gamma_K$ recognizing the intersection. While running, there will be an alert observing the construction: if some (algorithmically checkable) specific situation occurs, then the intersection $H \cap K$ is infinitely generated. We shall achieve our goal by proving that, in finite time, either the alert sounds or the procedure terminates providing the finite type reduced automaton $\Gamma_H \wedge \Gamma_K$ as output.

So, suppose that $H, K \leqslant G_1 * G_2$ are arbitrary finitely generated subgroups, and Γ_H and Γ_K are finite type reduced (G_1, G_2)-wedge graphs satisfying $\langle \Gamma_H \rangle = H$, and $\langle \Gamma_k \rangle = K$. We will define the junction $\Gamma_H \wedge \Gamma_K$ along the following paragraphs from a so-called *product automaton* of the reduced wedge automata Γ_H and Γ_K.

To begin with, we define the set of primary vertices of the product, which we will denote by $\widetilde{\Gamma}$, as the cartesian product of the primary vertices in Γ_H and the primary vertices in Γ_K,

$$V_0\widetilde{\Gamma} = V_0 \Gamma_H \times V_0 \Gamma_K \,.$$

Accordingly, the basepoint $\circledcirc$ of $\widetilde{\Gamma}$ is defined to be the pair of respective basepoints, i.e., $\circledcirc = (\circledcirc_H, \circledcirc_K)$.

Now, for $\nu = 1, 2$, we consider the subsets of primary vertices (in Γ_H and Γ_K) adjacent to some ν-secondary vertex,

$$\begin{aligned} V_{0\leftarrow\nu} \Gamma_H &= \{ p \in V_0 \Gamma_H \mid p \text{ is adj. to a } \nu\text{-secondary in } \Gamma_H \} \subseteq V_0 \Gamma_H, \\ V_{0\leftarrow\nu} \Gamma_K &= \{ p \in V_0 \Gamma_K \mid p \text{ is adj. to a } \nu\text{-secondary in } \Gamma_K \} \subseteq V_0 \Gamma_K, \end{aligned}$$

and define the following relation $\equiv_\nu$ on the set $V_{0\leftarrow\nu} \Gamma_H \times V_{0\leftarrow\nu} \Gamma_K \subseteq V_0\widetilde{\Gamma}$: $(p_1, p_1') \equiv_\nu (p_2, p_2')$, if and only if there exist two ν-elementary walks: γ from p_1 to p_2 in Γ_H (say $\gamma = p_1 e_1^{-1} q e_2 p_2$, with $q \in V_\nu \Gamma_H$), and γ' from p_1' to p_2' in Γ_K (say $\gamma' = p_1' e_1'^{-1} q' e_2' p_2'$, with $q' \in V_\nu \Gamma_K$), such that the intersection of their labels is nonempty, i.e.,

$$(p_1, p_1') \equiv_\nu (p_2, p_2') \;\Leftrightarrow\; \exists\, \nu\text{-vertices } q, q' : \begin{cases} p_1 - q - p_2 \,,\; p_1' - q' - p_2' \\ \ell_{p_1 \to q \to p_2} \cap \ell_{p_1' \to q' \to p_2'} \neq \varnothing \,. \end{cases} \tag{7.15}$$

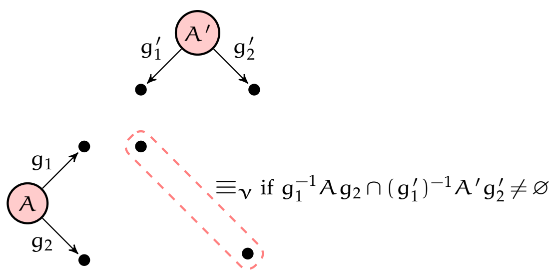

Fig. 7.22: ν-equivalence of primary vertices in $\widetilde{\Gamma}$

Lemma 7.4.32. *For $\nu = 1, 2$, the relation $\equiv_\nu$ defined in (7.15) is an equivalence relation on the subset of primary vertices $V_{0\leftarrow\nu}\Gamma_H \times V_{0\leftarrow\nu}\Gamma_K \subseteq V_0\widetilde{\Gamma}$.*

Proof. If any of the factors (Γ_H or Γ_K) does not have ν-secondary vertices, then its ν-neighborhood is empty, and the result follows immediately. So, we can assume that both factors do have ν-secondary vertices.

Fix an element $\mathbf{p} = (p, p') \in V_{0\leftarrow\nu}\Gamma_H \times V_{0\leftarrow\nu}\Gamma_K$. Let q be the ν-secondary vertex in Γ_H incident to p (there exists exactly one since $p \in V_{0\leftarrow\nu}\Gamma_H$, and Γ is reduced), and let $\gamma = pe^{-1}qep$ be the corresponding ν-elementary degenerate walk. Since $\ell_\gamma = \ell_e^{-1}\ell_q\ell_e$ is a subgroup of G_ν, it always contains the trivial element. Repeating the same on the second coordinate, we get $1 \in \ell_\gamma \cap \ell_{\gamma'} \neq \varnothing$ and so, $(p, p') \equiv_\nu (p, p')$. This shows that $\equiv_\nu$ is reflexive.

By construction, $\equiv_\nu$ is clearly symmetric.

Now, consider three elements $(p_1, p_1'), (p_2, p_2'), (p_3, p_3') \in V_0\Gamma_H \times V_0\Gamma_K$ and assume $(p_1, p_1') \equiv_\nu (p_2, p_2') \equiv_\nu (p_3, p_3')$. This means that, in Γ_H, there exist ν-elementary walks $\gamma_1 = p_1e_1^{-1}q_1e_2p_2$ and $\gamma_2 = p_2e_4^{-1}q_2e_3p_3$ with $q_1, q_2 \in V_\nu\Gamma_H$, and, in Γ_K, ν-elementary walks $\gamma_1' = p_1'e_1'^{-1}q_1'e_2'p_2'$ and $\gamma_2' = p_2'e_4'^{-1}q_2'e_3'p_3'$ with $q_1', q_2' \in V_\nu\Gamma_K$ such that

$$\varnothing \neq \ell_{e_1}^{-1}\ell_{q_1}\ell_{e_2} \cap \ell_{e_1'}^{-1}\ell_{q_1'}\ell_{e_2'} \subseteq G_\nu\,,$$
$$\varnothing \neq \ell_{e_4}^{-1}\ell_{q_2}\ell_{e_3} \cap \ell_{e_4'}^{-1}\ell_{q_2'}\ell_{e_3'} \subseteq G_\nu\,.$$

Since Γ_H and Γ_K are reduced automata, we deduce that $q_1 = q_2$ and $q_1' = q_2'$ (call them, respectively, q and q'), and $e_2 = e_4$ and $e_2' = e_4'$. Take elements in the above two nonempty sets,

$$\ell_{e_1}^{-1}a\ell_{e_2} = x = \ell_{e_1'}^{-1}a'\ell_{e_2'}\,,$$
$$\ell_{e_2}^{-1}b\ell_{e_3} = y = \ell_{e_2'}^{-1}b'\ell_{e_3'}\,,$$

for some $a, b \in \ell_q$ and $a', b' \in \ell_{q'}$. Then, $ab \in \ell_q$, $a'b' \in \ell_{q'}$ and so,

$$\ell_{e_1}^{-1}(ab)\ell_{e_3} = (\ell_{e_1}^{-1}a\ell_{e_2})(\ell_{e_2}^{-1}b\ell_{e_3}) = xy = (\ell_{e_1'}^{-1}a'\ell_{e_2'})(\ell_{e_2'}^{-1}b'\ell_{e_3'}) = \ell_{e_1'}^{-1}(a'b')\ell_{e_3'}\,.$$

Therefore, the ν-elementary walks $p_1e_1^{-1}qe_3p_3$ and $p_1'e_1'^{-1}q'e_3'p_3'$, in Γ_H and Γ_K, respectively, show that $(p_1, p_1') \equiv_\nu (p_3, p_3')$. This proves transitivity. □

Remark 7.4.33.

After Lemma 7.4.32, we can advance in the construction of the product automaton $\widetilde{\Gamma}$. For $\nu = 1, 2$, define the ν-secondary vertices of $\widetilde{\Gamma}$ to be the equivalence classes modulo $\equiv_\nu$. Namely,

$$V_\nu \widetilde{\Gamma} = (V_{0\leftarrow\nu} \Gamma_H \times V_{0\leftarrow\nu} \Gamma_K) / \equiv_\nu$$

Moreover, for each such secondary vertex $\mathbf{q} \in V_\nu \widetilde{\Gamma}$, and each primary vertex $\mathbf{p} = (p, p') \in \mathbf{q}$, add a ν-arc $\mathbf{q} \longrightarrow \mathbf{p}$. That is, the membership relation between primary and secondary vertices defines the arcs in $\widetilde{\Gamma}$:

$$E \widetilde{\Gamma} := \left\{ \mathbf{q} \longrightarrow \mathbf{p} \,:\, \mathbf{q} \in V_{1,2} \widetilde{\Gamma},\ \mathbf{p} \in V_0 \widetilde{\Gamma}, \text{ and } \mathbf{p} \in \mathbf{q} \right\}.$$

This defines the complete set of ν-arcs $E_\nu \widetilde{\Gamma}$, and finishes the definition of the underlying digraph of $\widetilde{\Gamma}$.

Observe that a primary vertex $(p, p') \in V_0 \widetilde{\Gamma}$ with one of its two coordinates not being adjacent to any ν-secondary vertex plays no role in the definition of $\equiv_\nu$ and so, becomes non-adjacent to any ν-secondary vertex in $\widetilde{\Gamma}$, either (note that, for a given (p, p') this could be the case for $\nu = 1$ and not for $\nu = 2$, or viceversa, or for both at the same time); otherwise, $(p, p') \in V_{0\leftarrow\nu} \Gamma_H \times V_{0\leftarrow\nu} \Gamma_K$ belongs to one and only one equivalence class of $\equiv_\nu$. Therefore, for $\nu = 1, 2$, every primary vertex of $\widetilde{\Gamma}$ is adjacent to at most one ν-secondary vertex of $\widetilde{\Gamma}$ through at most one arc.

On the other hand, observe that $\widetilde{\Gamma}$ may not be connected in general, even with Γ_H and Γ_K being so.

Once the underlying digraph for the product $\widetilde{\Gamma}$ of Γ_H and Γ_K is established, we define the projection (digraph) homomorphisms in the natural way.

Definition 7.4.34. Let $\widetilde{\Gamma}$ be the product of Γ_H and Γ_K. We define a digraph homomorphism $\pi \colon \widetilde{\Gamma} \to \Gamma_H$ by parts:

(a) its *restriction to primary vertices* is the projection $\pi \colon V_0 \widetilde{\Gamma} \to V_0 \Gamma_H,\ (p, p') \mapsto p$;

(b) its *restriction* $\pi \colon V_\nu \widetilde{\Gamma} \to V_\nu \Gamma_H$ *to ν-secondary vertices* assigns to every vertex $\mathbf{q} \in V_\nu \widetilde{\Gamma}$, the only ($\nu$-secondary) vertex in Γ_H adjacent to every vertex $p = (p, p')\pi \in V_0 \Gamma_H$ such that $(p, p') \in \mathbf{q}$;

(Note that this corresponds precisely to the image $\mathbf{q}\pi$ of the secondary vertex $\mathbf{q}$ thought as a class of primary vertices.)

(c) its *restriction* $\pi \colon E_\nu \widetilde{\Gamma} \to E_\nu \Gamma_H$ *to ν-arcs* is defined as follows: for every $\mathbf{q} \in V_\nu \widetilde{\Gamma}$, and every $(p, p') \in \mathbf{q}$, the (unique) ν-arc in $\widetilde{\Gamma}$ from $\mathbf{q}$ to (p, p') is assigned to the (unique) ν-arc in Γ_H from $q = \mathbf{q}\pi$ to $p = (p, p')\pi$.

Clearly, π is a well defined digraph homomorphism, called the *projection* to Γ_H. The projection to Γ_K, denoted by $\pi' \colon \widetilde{\Gamma} \to \Gamma_K$, is defined in an analogous way. Then, we write $\boldsymbol{\pi} := (\pi, \pi') \colon \widetilde{\Gamma} \to \Gamma_H \times \Gamma_K$.

Moreover, the respective restrictions of $\boldsymbol{\pi}$ to primary vertices and edges are injective. Thus, we will usually write $\mathbf{p} = (\mathbf{p}\pi, \mathbf{p}\pi') = (p, p')$, $\mathbf{q} = (\mathbf{q}\pi, \mathbf{q}\pi') = (q, q')$, and $\mathbf{e} = (\mathbf{e}\pi, \mathbf{e}\pi') = (e, e')$ (see Figure 7.23).

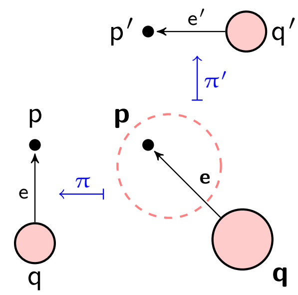

Fig. 7.23: Projection maps π and π'

In order to complete the definition of the wedge automaton $\widetilde{\Gamma}$, it remains to establish the labels for its vertices and arcs. We will do it in a way that extends the digraph homomorphisms π and π' into homomorphisms of the corresponding automata.

Definition 7.4.35. For $\nu = 1, 2$, and every ν-secondary vertex $\mathbf{q} \in V_\nu \widetilde{\Gamma}$, choose a particular primary vertex $\mathbf{p_q} = (p_\mathbf{q}, p'_\mathbf{q}) \in \mathbf{q}$, and let $\mathbf{e_q}$ be the (only) arc in $\widetilde{\Gamma}$ form $\mathbf{q}$ to the representative $\mathbf{p_q}$. This means that:

- in Γ_H there exists a ν-arc $e_\mathbf{q} := (\mathbf{e_q})\pi \in E_\nu \Gamma_H$ from $q = \mathbf{q}\pi \in V_\nu \Gamma_H$ to $p_\mathbf{q}$; and
- in Γ_K there exists a ν-arc $e'_\mathbf{q} := (\mathbf{e_q})\pi' \in E_\nu \Gamma_K$ from $q' = \mathbf{q}\pi' \in V_\nu \Gamma_K$ to $p'_\mathbf{q}$.

Then, we define the label of vertex $\mathbf{q}$ as

$$\ell_\mathbf{q} := \ell_{e_\mathbf{q}}^{-1} \, \ell_q \, \ell_{e_\mathbf{q}} \cap \ell_{e'_\mathbf{q}}^{-1} \, \ell_{q'} \, \ell_{e'_\mathbf{q}} = \ell_q^{\ell_{e_\mathbf{q}}} \cap \ell_{q'}^{\ell_{e'_\mathbf{q}}} \leqslant G_\nu \,. \tag{7.16}$$

Finally, for any ν-arc $\mathbf{e} \in E_\nu \widetilde{\Gamma}$ form $\mathbf{q}$ to a certain primary vertex $\mathbf{p} = (p, p') \in \mathbf{q}$, call $e := \mathbf{e}\pi$, $e' := \mathbf{e}\pi'$, and define the label of $\mathbf{e}$ as an arbitrary element from the corresponding coset intersection; i.e.,

$$\ell_\mathbf{e} \in \ell_{e_\mathbf{q}}^{-1} \, \ell_q \, \ell_e \cap \ell_{e'_\mathbf{q}}^{-1} \, \ell_{q'} \, \ell_{e'}, \tag{7.17}$$

which is nonempty because $(p_\mathbf{q}, p'_\mathbf{q}) \equiv_\nu (p, p')$ by construction (see Figure 7.24). Note that, in particular, we can take $\ell_{\mathbf{e_q}} = 1$.

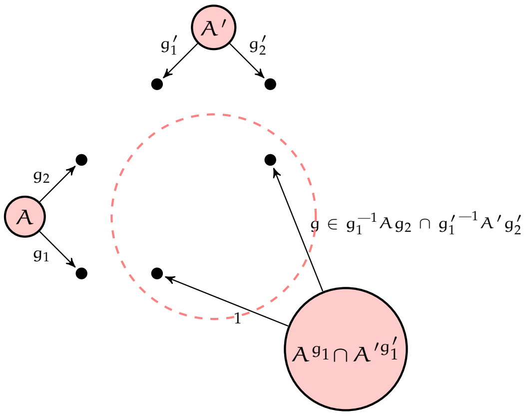

Fig. 7.24: Secondary vertex and related labels in $\widetilde{\Gamma}$

This completes the definition of a product automaton $\widetilde{\Gamma}$ of Γ_H and Γ_K (depending on some choices made on the way). Note also that $\widetilde{\Gamma}$ may not be connected, in general. The main property of the labels defined for $\widetilde{\Gamma}$ is expressed in the following lemma.

Lemma 7.4.36. *If $\widetilde{\Gamma}$ is a product of two reduced wedge automata Γ_H and Γ_K. Then, for every ν-elementary walk $\boldsymbol{\gamma}$ in $\widetilde{\Gamma}$, the projected walks $\boldsymbol{\gamma}\pi$ and $\boldsymbol{\gamma}\pi'$ are ν-elementary in Γ_H and Γ_K respectively; and we have*

$$\ell_{\boldsymbol{\gamma}} = \ell_{\boldsymbol{\gamma}\pi} \cap \ell_{\boldsymbol{\gamma}\pi'}. \tag{7.18}$$

Furthermore, $\boldsymbol{\gamma}$ is degenerate if and only if both $\boldsymbol{\gamma}\pi$ and $\boldsymbol{\gamma}\pi'$ are degenerate.

Proof. Clearly, the projections by π and π' of ν-elementary walks in $\widetilde{\Gamma}$ are also ν-elementary walks (in Γ_H and Γ_K, respectively), with the original one being degenerate if and only if both projections are degenerate. Note, however, that $\boldsymbol{\gamma}$ could be nondegenerate with one (and only one) of $\boldsymbol{\gamma}\pi$ and $\boldsymbol{\gamma}\pi'$ being degenerate; see Figure 7.24.

To see the equality in labels, let $\boldsymbol{\gamma} = \mathbf{p_1 e_1^{-1} q e_2 p_2}$ be a ν-elementary walk in $\widetilde{\Gamma}$, where $\mathbf{p_1} = (p_1, p_1')$ and $\mathbf{p_2} = (p_2, p_2')$, and let $\boldsymbol{\gamma}\pi = p_1 e_1^{-1} q e_2 p_2$ and $\boldsymbol{\gamma}\pi' = p_1'(e_1')^{-1} q' e_2' p_2'$ be the corresponding ν-elementary walks in Γ_H and Γ_K, respectively. No matter if $\boldsymbol{\gamma}$ is degenerate (i.e., if $\mathbf{e_1} = \mathbf{e_2}$) or not, we have $\ell_{\boldsymbol{\gamma}\pi} = \ell_{e_1}^{-1}\, \ell_q\, \ell_{e_2}$ and $\ell_{\boldsymbol{\gamma}\pi'} = \ell_{e_1'}^{-1}\, \ell_{q'}\, \ell_{e_2'}$.

Now, consider the chosen ν-arc $\mathbf{e_q}$ incident to $\mathbf{q}$ (possibly equal to $\mathbf{e_1}$ and/or $\mathbf{e_2}$). According to definitions (7.16) and (7.17), we have:

$$\ell_{\mathbf{q}} = \ell_q^{\ell_{e_q}} \cap \ell_{q'}^{\ell_{e_q'}} \quad \text{and} \quad \ell_{\mathbf{e_i}} \in \ell_{e_q}^{-1}\, \ell_q\, \ell_{e_i} \cap \ell_{e_q'}^{-1}\, \ell_{q'}\, \ell_{e_i'},$$

for $i = 1, 2$ (see Figure 7.24). Therefore,

$$\ell_{\boldsymbol{\gamma}} = \ell_{\mathbf{e_1}}^{-1}\, \ell_{\mathbf{q}}\, \ell_{\mathbf{e_2}} = \ell_{\mathbf{e_1}}^{-1} \cdot \left(\ell_q^{\ell_{e_q}} \cap \ell_{q'}^{\ell_{e_q'}} \right) \cdot \ell_{\mathbf{e_2}} = \ell_{e_1}^{-1}\, \ell_q\, \ell_{e_2} \cap \ell_{e_1'}^{-1}\, \ell_{q'}\, \ell_{e_2'} = \ell_{\boldsymbol{\gamma}\pi} \cap \ell_{\boldsymbol{\gamma}\pi'}.$$

This completes the proof. $\square$

We can now state the definition of junction automaton $\Gamma_H \wedge \Gamma_K$ and show that it is a reduced (G_1, G_2)-wedge automaton such that $\langle \Gamma_H \wedge \Gamma_K \rangle = H \cap K$.

Definition 7.4.37. Let G_1, G_2 be two groups, let $H, K \leqslant G_1 * G_2$ be two subgroups, and let Γ_H, Γ_K be reduced (G_1, G_2)-wedge automata recognizing H and K, respectively; and let $\widetilde{\Gamma}$ be a product of Γ_H and Γ_K (recall that there are some arbitrary choices made on the way).

Then, the *junction (automaton)* of Γ_H and Γ_K (relative to the product $\widetilde{\Gamma}$), which we will loosely denote by $\Gamma_H \wedge \Gamma_K$, is the connected component of $\widetilde{\Gamma}$ containing the basepoint.

Proposition 7.4.38. *Any junction automaton $\Gamma_H \wedge \Gamma_K$ is a (G_1, G_2)-reduced automaton recognizing $H \cap K$.*

Proof. By definition, the junction $\Gamma_H \wedge \Gamma_K$ is connected. So, it satisfies property (ii) from Definition 7.4.18.

As observed above, for $\nu = 1, 2$, every primary vertex of $\widetilde{\Gamma}$ is adjacent to at most one ν-secondary vertex of $\widetilde{\Gamma}$ through at most one arc; so $\widetilde{\Gamma}$, and hence $\Gamma_H \wedge \Gamma_K$, satisfies property (ii) from Definition 7.4.18.

To see property (iii) in Definition 7.4.18, take $\nu = 1, 2$, and let $\mathbf{q} \in V_\nu \widetilde{\Gamma}$ be a ν-secondary vertex of $\widetilde{\Gamma}$, and let $\mathbf{e_1}, \mathbf{e_2}$ be two *different* ν-arcs from $\mathbf{q}$ to $\mathbf{p_1} = (p_1, p_1')$ and $\mathbf{p_2} = (p_2, p_2')$ respectively; i.e., $\boldsymbol{\gamma} = \mathbf{p}\mathbf{e_1}^{-1}\mathbf{q}\mathbf{e_2}\mathbf{p_2}$ is a nondegenerate elementary walk in $\widetilde{\Gamma}$.

By symmetry, we can assume $p_1 \neq p_2$, i.e., that $\boldsymbol{\gamma}\pi = p_1 e_1^{-1} q e_2 p_2$ is a nondegenerate elementary walk in Γ_H. Since Γ_H is a (G_1, G_2)-reduced automaton, $1 \notin \ell_{e_1}^{-1} \ell_q \ell_{e_2} = \ell_{\gamma\pi}$. Therefore, by Lemma 7.4.36, $1 \notin \ell_\gamma = \ell_{\gamma\pi} \cap \ell_{\gamma\pi'}$. This shows that $\widetilde{\Gamma}$, and hence $\Gamma_H \wedge \Gamma_K$, satisfies property (iii) from Definition 7.4.18.

Therefore, $\Gamma_H \wedge \Gamma_K$ is a reduced (G_1, G_2)-automaton.

It remains to show that $\langle \Gamma_H \wedge \Gamma_K \rangle = H \cap K$ (or, equivalently, $\langle \widetilde{\Gamma} \rangle = H \cap K$).

Indeed, let $\boldsymbol{\gamma}$ be an arbitrary $\odot$-walk in $\widetilde{\Gamma}$, and let $\boldsymbol{\gamma} = \boldsymbol{\gamma_1} \cdots \boldsymbol{\gamma_r}$ be its elementary decomposition. Clearly, the elementary decompositions of $\boldsymbol{\gamma}\pi$ in Γ_H, and $\boldsymbol{\gamma}\pi'$ in Γ_K, are $\boldsymbol{\gamma}\pi = (\boldsymbol{\gamma_1}\pi) \cdots (\boldsymbol{\gamma_r}\pi)$ and $\boldsymbol{\gamma}\pi' = (\boldsymbol{\gamma_1}\pi') \cdots (\boldsymbol{\gamma_r}\pi')$, respectively. Then, by Lemma 7.4.36,

$$\ell_\gamma = \ell_{\gamma_1} \cdots \ell_{\gamma_r} \subseteq \ell_{\gamma_1\pi} \cdots \ell_{\gamma_r\pi} = \ell_{(\gamma_1\pi)\cdots(\gamma_r\pi)} = \ell_{\gamma\pi} \subseteq H.$$

Since this is true for every γ, we deduce $\langle \Gamma_H \wedge \Gamma_K \rangle \subseteq H$; and, by the symmetric argument, also $\langle \Gamma_H \wedge \Gamma_K \rangle \subseteq K$.

For the converse inclusion, take an element $g \in H \cap K$, and let $g = g_1 \cdots g_r$ be its syllable decomposition in $G_1 * G_2$. Since Γ_H is a (G_1, G_2)-reduced automaton and $\langle \Gamma_H \rangle = H$, Lemma 7.4.21 ensures us that $g \in \widetilde{\ell}_\gamma$ for some *alternating* walk γ, closed at the basepoint of Γ_H. In this situation, its elementary decomposition, $\gamma = \gamma_1 \cdots \gamma_r$, correspond to the syllable decomposition $g = g_1 \cdots g_r$, i.e., $g_i \in \widetilde{\ell}_{\gamma_i}$, for $i = 1, \ldots, r$. Similarly, there exists an *alternating* walk γ', closed at the basepoint of Γ_K, whose elementary decomposition $\gamma' = \gamma'_1 \cdots \gamma'_r$ again correspond to the syllable decomposition $g = g_1 \cdots g_r$, i.e., $g_i \in \widetilde{\ell}_{\gamma'_i}$, for $i = 1, \ldots, r$.

Write $\gamma_i = p_{i-1} e_i^{-1} q_i f_i p_i$ and $\gamma'_i = p'_{i-1} (e'_i)^{-1} q'_i f'_i p'_i$. Then, for each $i = 1, \ldots, r$, we have $g_i \in \ell_{\gamma_i} = \ell_{e_i}^{-1} \ell_{q_i} \ell_{f_i}$ and $g_i \in \ell_{\gamma'_i} = \ell_{e'_i}^{-1} \ell_{q'_i} \ell_{f'_i}$; so,

$$\varnothing \neq \ell_{e_i}^{-1} \ell_{q_i} \ell_{f_i} \cap \ell_{e'_i}^{-1} \ell_{q'_i} \ell_{f'_i} \subseteq G_\nu \,.$$

and this means that $\mathbf{p_{i-1}} = (p_{i-1}, p'_{i-1}) \equiv_{\nu_i} (p_i, p'_i) = \mathbf{p_i}$, where ν_i is the common type of the vertices q_i (in Γ_H) and q'_i (in Γ_K); see (7.15). Therefore, $\mathbf{p_{i-1}}$ and $\mathbf{p_i}$ are both incident to a common ν_i-secondary vertex in $\Gamma_H \wedge \Gamma_K$. In other words, there is a ν_i-elementary walk in $\Gamma_H \wedge \Gamma_K$, say $\boldsymbol{\gamma_i}$, from $\mathbf{p_{i-1}}$ to $\mathbf{p_i}$. Finally, by Lemma 7.4.36, $g_i \in \ell_{\gamma_i} \cap \ell_{\gamma'_i} = \ell_{\boldsymbol{\gamma_i}}$. Therefore,

$$g = g_1 \cdots g_r \in \ell_{\boldsymbol{\gamma_1}} \cdots \ell_{\boldsymbol{\gamma_r}} = \ell_{\boldsymbol{\gamma_1} \cdots \boldsymbol{\gamma_r}} \subseteq \langle \Gamma_H \wedge \Gamma_K \rangle \,,$$

concluding the proof. ☐

Corollary 7.4.39. *In the above situation, $H \cap K$ is finitely generated if and only if all the vertex labels of $\Gamma_H \wedge \Gamma_K$ are finitely generated.*

Proof. By construction, the underlying graph of $\Gamma_H \wedge \Gamma_K$ is finite. So, the result follows from Proposition 7.4.38 and Corollary 7.4.26. ☐

7.4.5 Understanding intersections of cosets

According to Lemma 7.4.16, given a wedge automaton Γ_H, the union of the labels of all the walks in Γ_H from the basepoint to a primary vertex $p \in V_0 \Gamma_H$, denoted by $\langle \Gamma_H \rangle_{(\circledast, p)}$, constitute a coset of the recognized subgroup $\langle \Gamma_H \rangle = H$. In general, though, this does not reflect all the cosets of H (consider, for example the cases when Γ_H has only finitely many primary vertices). However, we can slightly modify the automaton Γ_H to achieve this purpose:

Let $u = a_1b_1 \cdots a_sb_s \in G_1 * G_2$, written in normal form. Consider the (G_1, G_2)-wedge automaton (also denoted by u) consisting on a single segment line spelling the normal form for g, and having trivial vertex labels (like the petal in Figure 7.10, but without identifying the initial and terminal vertices ιu and τu); let us call it the *hair* for u. Attach this hair to Γ_H by identifying the basepoint $\circledast_H$ with ιu, and then apply the folding process until no more foldings are possible (see the proof of Theorem 7.4.30). Observe that operation (II) will not be used, and the triviality of the vertex labels in the hair implies that the vertex groups already present in Γ_H will not change along the process. So, the output is the exact same graph Γ_H with a terminal segment of the hair (maybe the whole of it) attached somewhere and sticking out; denote this by Γ_{Hu}. Clearly, Γ_{Hu} is a (G_1, G_2)-reduced automaton, like Γ_H, and furthermore, since the new secondary vertices out of Γ_H have trivial label, $\langle \Gamma_{Hu} \rangle = \langle \Gamma_H \rangle = H$.

By Lemma 7.4.16 (iii), $\langle \Gamma_{Hu} \rangle_{(\circledast_H, \tau u)} = \langle \Gamma_H \rangle \cdot u$ (the situation where this coset could already be represented by a vertex in Γ_H corresponds to the fact that the hair happens to fold completely and so, $\Gamma_{Hu} = \Gamma_H$).

Now let us go back to the graph $\Gamma_H \wedge \Gamma_K$. It is useful to understand the intersection of H and K but also, adding the corresponding hairs, it will be useful to understand the intersection of two arbitrary cosets Hu and Kv.

Given elements $u, v \in G_1 * G_2$, consider the (G_1, G_2)-reduced automata Γ_{Hu} and Γ_{Kv}, and the graph $\Gamma_{Hu} \wedge \Gamma_{Kv}$ constructed exactly like $\Gamma_H \wedge \Gamma_K$ but starting with Γ_{Hu} and Γ_{Kv} instead of Γ_H and Γ_K.

Lemma 7.4.40. *With the above notation,*

(i) $\Gamma_H \wedge \Gamma_K$ *is a* (G_1, G_2)*-reduced subgraph of* $\Gamma_{Hu} \wedge \Gamma_{Kv}$*;*

(ii) $Hu \cap Kv \neq \varnothing$ *if and only if the vertex* $(\tau u, \tau v)$ *belongs to* $\Gamma_{Hu} \wedge \Gamma_{Kv}$*;*

(iii) *for any walk* γ *in* $\Gamma_{Hu} \wedge \Gamma_{Kv}$ *from* $(\circledast_H, \circledast_K)$ *to* $(\tau u, \tau v)$*, and any* $g \in \ell_\gamma$*, we have*

$$Hu \cap Kv = \langle \Gamma_{Hu} \wedge \Gamma_{Kv} \rangle_{(\circledast_H, \circledast_K), (\tau u, \tau v)} = (H \cap K)g.$$

Proof. Note that the initial set of primary vertices for $\Gamma_{Hu} \wedge \Gamma_{Kv}$, namely $V_0 \Gamma_{Hu} \times V_0 \Gamma_{Kv}$, contains as a subset $V_0 \Gamma_H \times V_0 \Gamma_K$, the initial set of primary vertices for $\Gamma_H \wedge \Gamma_K$. And two old vertices $(p_1, p_1'), (p_2, p_2') \in \Gamma_H \wedge \Gamma_K$ are $\equiv_\gamma$-equivalent in $\Gamma_H \wedge \Gamma_K$ if and only if they are $\equiv_\gamma$-equivalent as vertices in $\Gamma_{Hu} \wedge \Gamma_{Kv}$ (since vertices of $\Gamma_{Hu} \wedge \Gamma_{Kv}$ outside $\Gamma_H \wedge \Gamma_K$ have always trivial labels). This proves (i).

Suppose first that $(\tau u, \tau v) \in V\,\Gamma_{Hu} \wedge \Gamma_{Kv}$, let γ be a walk in $\Gamma_{Hu} \wedge \Gamma_{Kv}$ from $(\circledast_H, \circledast_K)$ to $(\tau u, \tau v)$, and consider its basic label $\ell_\gamma^\bullet \in G_1 * G_2$. By the same argument as in

Proposition 7.4.38, $\ell^{\bullet}_{\gamma} \in \ell_{\gamma\pi} \cap \ell_{\gamma\pi'}$. But $\gamma\pi$ (resp., $\gamma\pi'$) is a walk in Γ_{Hu} from $\circledast_H$ to τu (resp., a walk in Γ_{Kv} from $\circledast_K$ to τv) hence, by Lemma 7.4.16 (iii),

$$\ell^{\bullet}_{\gamma} \in \langle \Gamma_{Hu} \rangle_{(\circledast_H, \tau u)} \cap \langle \Gamma_{Kv} \rangle_{(\circledast_K, \tau v)} = Hu \cap Kv \,,$$

concluding that $Hu \cap Kv \neq \varnothing$.

Conversely, suppose that $Hu \cap Kv \neq \varnothing$ and let $g \in Hu \cap Kv$. Again by Lemma 7.4.16 (iii), there exist walks γ in Γ_{Hu} from $\circledast_H$ to τu, and γ' in Γ_{Kv} from $\circledast_K$ to τv, such that $g \in \ell_{\gamma} \cap \ell_{\gamma'}$. Again, with an argument like in the proof of Proposition 7.4.38, there exists a walk $\boldsymbol{\gamma}$ in $\Gamma_{Hu} \wedge \Gamma_{Kv}$ from $(\circledast_H, \circledast_K)$ to $(\tau u, \tau v)$ such that $g \in \ell_{\boldsymbol{\gamma}}$. In particular, $(\tau u, \tau v) \in \Gamma_{Hu} \wedge \Gamma_{Kv}$. This proves (ii) and (iii). □

7.4.6 Algorithmic treatment of the junction automaton

Let us now address the algorithmic aspects of this construction. For all the present section, assume the two starting reduced (G_1, G_2)-automata Γ_H and Γ_K to be of finite type (namely, H and K are finitely generated subgroups of $G_1 * G_2$).

By construction, the underlying graph of $\Gamma_H \wedge \Gamma_K$ is finite, but not necessarily of finite type (since the labels of the vertices in $\Gamma_H \wedge \Gamma_K$ may very well be infinitely generated as a possible result of intersections of finitely generated subgroups of G_1 and G_2).

A first easy observation is that, under the assumption that both G_1 and G_2 are Howson, then $\Gamma_H \wedge \Gamma_K$ will always of finite type. This recovers a classical result originally proved by Baumslag in [Bau66].

Theorem 7.4.41 (Baumslag, 1966, [Bau66]; Ivanov, 1999, [Iva99])**.** *Any arbitrary free product of Howson groups is again Howson.*

Proof. Suppose G_1 and G_2 are Howson, and let $H, K \leqslant G_1 * G_2$ be finitely generated. By Theorem 7.4.30, there exist (G_1, G_2)-reduced automata of finite type Γ_H and Γ_K respectively recognizing H and K. Then, the (G_1, G_2)-reduced automaton $\Gamma_H \wedge \Gamma_K$ is again of finite type (this is clear since the order of the junction is bounded, and you can only obtain finitely generated labels after finitely many foldings). And, by Corollary 7.4.26, $H \cap K = \langle \Gamma_H \wedge \Gamma_K \rangle$ is finitely generated. Hence, $G_1 * G_2$ is Howson.

By induction, any finite free product of Howson groups is again Howson. Finally, an arbitrary free product $\ast_{i \in I} G_i$ of Howson groups G_i is also again Howson since any finitely generated subgroup $H \leqslant \ast_{i \in I} G_i$ is contained in $\ast_{i \in I_0} G_i$ for some big enough finite set of indices $I_0 \subseteq I$. □

In the present section — *without assuming the Howson property for the factors* — we shall give an algorithm which, on input two finitely generated subgroups $H, K \leqslant G_1 * G_2$ (given by finite sets of generators in normal form), decides whether $H \cap K$ is finitely generated or not and, in the affirmative case, computes an automaton of the form $\Gamma_H \wedge \Gamma_K$ recognizing this intersection (assuming SIP in H and K). For this purpose, we shall need to decide whether certain intersections of pairs of finitely generated subgroups of G_i are finitely generated again, or not; and in the affirmative case, whether certain cosets of them do intersect or not. This conducts us to the proofs of Theorems 7.3.1 and 7.3.2.

Theorem 7.3.2. *If two finitely presented groups* G_1 *and* G_2 *satisfy* ESIP, *then their free product* $G_1 * G_2$ *also satisfies* ESIP. □

Proof. Assume that both G_1 and G_2 satisfy ESIP; and suppose we are given two finitely generated subgroups $H, K \leqslant G_1 * G_2$ by finite sets of generators, and two extra elements $u, v \in G_1 * G_2$, all of them in normal form. By Remark 7.1.4, both G_1 and G_2 also have solvable membership problem; and by Theorem 7.4.30, we can compute reduced (G_1, G_2)-automata Γ_H and Γ_K such that $\langle \Gamma_H \rangle = H$, and $\langle \Gamma_K \rangle = K$.

Now recall Proposition 7.4.38 and Corollary 7.4.39: the intersection $H \cap K$ is recognized by the (finite) junction automaton $\Gamma_H \wedge \Gamma_K$, and is finitely generated if and only if all the vertex labels in $\Gamma_H \wedge \Gamma_K$ are finitely generated.

So, the next step is to start constructing the graph $\Gamma_H \wedge \Gamma_K$; locally from its basepoint: For every new ν-secondary vertex $\mathbf{q}$ added to the picture, we shall algorithmically detect whether its label $\ell_{\mathbf{q}}$ is finitely generated or not: if it is not, then we kill the whole process and deduce that $H \cap K$ is infinitely generated; otherwise, we shall compute a finite set of generators for $\ell_{\mathbf{q}}$, and proceed with the construction of $\Gamma_H \wedge \Gamma_K$. In this way, we shall either detect that the intersection $H \cap K$ is infinitely generated, or will complete the construction of $\Gamma_H \wedge \Gamma_K$.

To eventually compute this construction, we start looking at the basepoint $\odot = (\odot_H, \odot_K)$, with the whole set $V_0 \widetilde{\Gamma} = V_0 \Gamma_H \times V_0 \Gamma_K$ in the background. We have to keep adding ν-secondary vertices (with their labels), and ν-arcs (with their labels too) connecting them to certain primaries, until getting $\Gamma_H \wedge \Gamma_K$, the full connected component of $\widetilde{\Gamma}$ containing the basepoint $\odot$.

We start checking whether there exists $\nu \in \{1, 2\}$, such that both $\odot_H$ and $\odot_K$ have nonempty ν-neighborhoods. If not, then the basepoint $\odot$ is not adjacent to any secondary vertex in $\widetilde{\Gamma}$, and we are done (namely, the product $\widetilde{\Gamma}$ is the trivial automata, and the intersection $H \cap K = 1$). Otherwise, let $q \in V_\nu \Gamma_H$, and $e \in E_\nu \Gamma_H$ with $\iota e = q$, $\tau e = \odot_H$; and let $q' \in V_\nu \Gamma_K$ and $e' \in E_\nu \Gamma_K$ with $\iota e' = q'$, $\tau e' = \odot_K$, and enlarge our picture by drawing a new ν-secondary vertex, say $\mathbf{q}$, and a new ν-arc, say $\mathbf{e} = (e, e')$, from $\mathbf{q}$ to $\circledbullet$. According to (7.16) — and with respect to the choice $(p_{\mathbf{q}}, p'_{\mathbf{q}}) = (\circledbullet_H, \circledbullet_K)$ — we know that the label of $\mathbf{q}$ is $\ell_{\mathbf{q}} = \ell_q^{\ell_e} \cap \ell_{q'}^{\ell_{e'}} \leqslant G_\nu$.

Applying SIP for G_ν to the (finitely generated) subgroups $\ell_q^{\ell_e}$ and $\ell_{q'}^{\ell_{e'}}$, we can decide whether $\ell_{\mathbf{q}}$ is finitely generated or not. In case it is not, kill the whole process and declare $H \cap K$ infinitely generated. Otherwise, compute a finite set of generators for $\ell_{\mathbf{q}}$, assign $\ell_{\mathbf{e}} = 1$, and check which other primary vertices from $\widetilde{\Gamma}$ are adjacent to $\mathbf{q}$: $\mathbf{p} = (p, p') \in V_0 \widetilde{\Gamma}$ is adjacent to $\mathbf{q}$ if and only if $(p, p') \equiv_\nu (p_{\mathbf{q}}, p'_{\mathbf{q}})$, which happens if and only if there exists $f \in E_\nu \Gamma_H$ from q to p, and $f' \in E_\nu \Gamma_K$ from q' to p', such that $\ell_e^{-1} \ell_q \ell_f \cap \ell_{e'}^{-1} \ell_{q'} \ell_{f'} \neq \varnothing$. So, run over every $p \in V_0 \Gamma_H$ adjacent to q, and every $p' \in V_0 \Gamma_K$ adjacent to q' and, for each such pair, check whether the intersection of (right) cosets

$$\ell_q^{\ell_e} \cdot (\ell_e^{-1} \ell_f) \cap \ell_{q'}^{\ell_{e'}} \cdot (\ell_{e'}^{-1} \ell_{f'}) = \ell_e^{-1} \ell_q \ell_f \cap \ell_{e'}^{-1} \ell_{q'} \ell_{f'} \tag{7.19}$$

is empty or not; this can be done using the above call to ESIP from G_ν, since they are right cosets of $\ell_q^{\ell_e}, \ell_{q'}^{\ell_{e'}} \leqslant G_\nu$, whose intersection happens to be finitely generated.

Then, in case that intersection is not empty, add a ν-arc, say $\mathbf{f} = (f, f')$, from $\mathbf{q}$ to $\mathbf{p}$, and $\ell_{(f,f')}$ arbitrarily chosen from that nonempty intersection. After this procedure, we have a complete picture of the 1-elementary and 2-elementary walks in $\Gamma_H \wedge \Gamma_K$ starting at the basepoint $\circledbullet$.

Now, for every $\nu = 1, 2$, and every primary vertex $\mathbf{p} = (p, p')$ added to the picture and not yet explored yet, repeat the same process (with $\mathbf{p}$ in place of $\circledbullet$). Since the underlying graph of $\widetilde{\Gamma}$ is finite, this procedure will either find an infinitely generated vertex label (so detecting the infinite generated type of the intersection), or finish in finite time, with output the complete junction $\Gamma_H \wedge \Gamma_K$, from which we can obtain generators for $H \cap K$ (in fact, its Kurosh decomposition) applying Theorem 7.4.25.

Hence, so far we have proved $\mathsf{SIP}(G_1 * G_2)$. Note that the decision about the disjunctiveness of the intersection of two finitely generated cosets (when the corresponding intersection of subgroups is finitely generated) is still pending.

To carry it out, let us place ourselves in the case where $H \cap K$ finitely generated (and so, we have constructed the full junction automata $\Gamma_H \wedge \Gamma_K$). Now the inputs u, v start playing, and we have to decide whether the intersection of right cosets $Hu \cap Kv$ is empty or not. We can extend the computation of $\Gamma_H \wedge \Gamma_K$ to that of $\Gamma_{Hu} \wedge \Gamma_{Kv}$; or, if you prefer, construct directly $\Gamma_{Hu} \wedge \Gamma_{Kv}$ — since we know (see Lemma 7.4.40 (i)) that $\Gamma_H \wedge \Gamma_K$ will show up as a subgraph and so, the process will not be killed for the presence of infinitely generated vertex labels.

Tt only remains to check whether the vertex $(\tau u, \tau v)$ appears in $\Gamma_{Hu} \wedge \Gamma_{Kv}$, or not. Namely, using Lemma 7.4.40 (ii): $(\tau u, \tau v)$ is connected to $(\circledbullet_H, \circledbullet_K)$ if and only if the intersection $Hu \cap Kv$ is nonempty; and, if so, any element g from the label of any walk from $(\circledbullet_H, \circledbullet_K)$ to $(\tau u, \tau v)$ belongs to such intersection, $g \in Hu \cap Kv = (H \cap K)g$.

This concludes the proof. □

Finally, we complement the arguments in the last proof to prove that TIP also passes through free products.

Theorem 7.3.1. *If two finitely presented groups* G_1 *and* G_2 *satisfy* TIP, *then their free product* $G_1 * G_2$ *also satisfies* TIP. □

Proof. Since TIP implies ESIP, Theorem 7.3.2 already gives us ESIP for $G_1 * G_2$. It remains to solve CIP in the case where the given finitely generated subgroups H, K have infinitely generated intersection.

Given $H, K \leqslant G_1 * G_2$ finitely generated, and $w, w' \in G_1 * G_2$, run the same algorithm as in the proof of Theorem 7.3.2: construct Γ_{Hu} and Γ_{Kv} and start building the graph $\Gamma_{Hu} \wedge \Gamma_{Kv}$. When we encounter a secondary vertex $\mathbf{q}$ whose label $\ell_{\mathbf{q}} = \ell_q^{\ell_e} \cap \ell_{q'}^{\ell_{e'}} \leqslant G_\nu$ is infinitely generated, instead of computing a set of generators for it (which is not possible), we just put the trivial subgroup as a label in place of $\ell_{\mathbf{q}}$. Then, when analyzing which other primary vertices are adjacent to $\mathbf{q}$, we need to decide if the intersection of cosets from equation (7.19) are empty or not: even though $\ell_q^{\ell_e} \cap \ell_{q'}^{\ell_{e'}} \leqslant G_\nu$ is infinitely generated, the decision can be made effective using CIP from G_ν. This way, we can algorithmically complete the description of $\Gamma_{Hu} \wedge \Gamma_{Kv}$ *except that*, for some secondary vertices $\mathbf{q}$, instead of having generators for $\ell_{\mathbf{q}}$, we just have the trivial element in them.

Of course, this is not enough information for computing a set of generators for $H \cap K$ (which, according to Corollary 7.4.39, is infinitely generated). But it suffices for deciding whether the vertices $(\circledbullet_H, \circledbullet_K)$ and $(\tau w, \tau w')$ belong to the same connected component of $\Gamma_{Hu} \wedge \Gamma_{Kv}$. By Lemma 7.4.40, this allows us to decide whether the intersection of cosets $Hu \cap Kv$ is empty or not, and in case it is not, we can compute an element from it, just choosing a walk γ from $(\circledbullet_H, \circledbullet_K)$ to $(\tau w, \tau w')$, and then picking an element from ℓ_γ (if γ traverses some secondary vertex with infinitely generated label, we just recorded an element from it for this purpose).

This completes the proof. □

7.5 The direct product case

This section is devoted to proving Theorems 7.3.3 and 7.3.4 (concerning preservability of intersection problems through direct products with free-abelian groups). To this end, we analyze the Droms groups presented by connected Droms graphs (i.e., by Droms cones).

Fix an arbitrary Droms graph Γ_0 on $n \geqslant 1$ vertices, say $V\Gamma_0 = X = \{x_1, \ldots, x_n\}$, and a complete graph K_m on m vertices, say $VK_m = T = \{t_1, \ldots, t_m\}$, and consider the join $\Gamma = K_m \vee \Gamma_0$. Recall that every Droms group $\mathbb{G}_\Gamma$ admits such a *primary decomposition* (Corollary 7.2.8), with Γ_0 disconnected (or empty, if Γ is complete), and in that case, the starting group $\mathbb{G}_\Gamma$ is connected if and only if $m \geqslant 1$. Of course, this situation, corresponds to K_m agglutinating the center of Γ, and gives rise to the short exact sequence:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \mathbb{Z}^m & \longrightarrow & \mathbb{G}_\Gamma & \xrightarrow{\pi_0} & \mathbb{G}_{\Gamma_0} & \longrightarrow & 1\,, \\ & & & & t^{\mathbf{a}}\,u & \longmapsto & u & & \end{array} \tag{7.20}$$

where $\pi_0 : \mathbb{G}_\Gamma \to \mathbb{G}_{\Gamma_0}$ is the natural map killing the center (i.e., at the level of words, the map π_0 just erases the occurrences of letters in $T^{\pm} = \{t_1, \ldots, t_m\}^{\pm}$).

We have to show that if $\mathbb{G}_{\Gamma_0}$ satisfies SIP (resp., ESIP), then so does $\mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$.

Remark 7.5.1. First of all, observe that if any of this results were true for disconnected graphs Γ_0, it would automatically be true in general: since any Droms group Γ_0 is of the form $\Gamma_0 = K_r \vee \Gamma_0'$, with Γ_0' disconnected; and if $\mathbb{G}_{\Gamma_0}$ satisfies SIP (resp ESIP), its subgroup $\mathbb{G}_{\Gamma_0'}$ satisfies it too, then

$$\mathbb{Z}^{m+r} \times \mathbb{G}_{\Gamma_0'} = \mathbb{Z}^m \times (\mathbb{Z}^r \times \mathbb{G}_{\Gamma_0'}) = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0} \tag{7.21}$$

will satisfy SIP (resp ESIP) as well.

Therefore, without lost of generality, we can assume for the rest of the proof that Γ_0 is disconnected, say, $\Gamma_0 = \Gamma_{0,1} \sqcup \Gamma_{0,2}$, with $\Gamma_{0,1}, \Gamma_{0,2} \neq \varnothing$ (i.e., $\mathbb{G}_{\Gamma_0} = \mathbb{G}_{\Gamma_{0,1}} * \mathbb{G}_{\Gamma_{0,2}}$ with $\mathbb{G}_{\Gamma_{0,1}}, \mathbb{G}_{\Gamma_{0,2}} \neq 1$); in particular, $Z(\mathbb{G}_\Gamma) = \langle t_1, \ldots, t_m \rangle \simeq \mathbb{Z}^m$.

Every element in $\mathbb{G}_\Gamma$ can be written as a word on $\{t_1, \ldots, t_m, x_1, \ldots, x_n\}$, where the t_i's are free to move at any position, because they commute with any other letter. We will write all these t_i's systematically on the left, and will abbreviate them as a vectorial power of a formal symbol 't'. This way, every element in $\mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$ can be written in the form

$$t_1^{a_1} \cdots t_m^{a_m}\, u(x_1, \ldots, x_n) = t^{\mathbf{a}}\, u(x_1, \ldots, x_n)\,,$$

where $\mathbf{a} = (a_1, \ldots, a_m) \in \mathbb{Z}^m$, and $u = u(x_1, \ldots, x_n)$ is a word on the x_i's. Clearly, the product of elements under this form is given by the rule

$$(t^{\mathbf{a}} u) \cdot (t^{\mathbf{b}} v) \;=\; t^{\mathbf{a}+\mathbf{b}} uv\,.$$

Now, we introduce a notion that will simplify the forthcoming discussion in this setting.

Definition 7.5.2. Let $\mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$ be the primary decomposition of a Droms group. Then, for a given subgroup $H \leqslant \mathbb{G}_\Gamma$, and an element $u \in \mathbb{G}_{\Gamma_0}$, we define the *(abelian) completion* of u in H (a.k.a. the H-*completion* of u) to be the set

$$\mathcal{C}_H(u) \;=\; \{\mathbf{a} \in \mathbb{Z}^m \mid t^{\mathbf{a}} u \in H\}\,.$$

Lemma 7.5.3. *The completion $\mathcal{C}_H(u)$ is either empty (when $u \notin H\pi_0$), or a coset of $\mathbb{Z}^m \cap H$. Moreover, if $u_1, \ldots, u_n \in H\pi_0$, and $\omega(u_1, \ldots, u_n)$ is an arbitrary word on them, then*

$$\mathcal{C}_H(\omega(u_1, \ldots, u_n)) \;=\; \sum_{i=1}^{n} \omega_i\, \mathcal{C}_H(u_i)\,,$$

where $\omega_i = |\omega|_i$ is the total exponent of the variable u_i in ω.

Proof. Obviously, $\mathcal{C}_H(u) \neq \varnothing \Leftrightarrow u \in H\pi_0$. Also, $\mathcal{C}_H(u)$ happens to be a coset of $\mathbb{Z}^m \cap H$ since, given $t^{\mathbf{a}} \in \mathcal{C}_H(u)$,

$$\mathbf{b} \in \mathcal{C}_H(u) \;\Leftrightarrow\; t^{\mathbf{b}} u \in H \;\Leftrightarrow\; (t^{\mathbf{b}} u)(t^{\mathbf{a}} u)^{-1} \in H \;\Leftrightarrow\; t^{\mathbf{b}-\mathbf{a}} \in \mathbb{Z}^m \cap H\,.$$

With the regular addition of cosets, $(\mathbf{a} + L) + (\mathbf{b} + L) = (\mathbf{a} + \mathbf{b}) + L$, and $\lambda(\mathbf{a} + L) = \lambda\mathbf{a} + L$, the second claim is straightforward to see. ☐

The next lemma shows how the subgroups of a connected Droms group (which we know that are again PC-groups) are related to its center, and allows us to derive useful consequences.

Lemma 7.5.4. *Let $\mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$ be the primary decomposition of a connected Droms PC-group. Then, any subgroup $H \leqslant \mathbb{G}_\Gamma$ splits as:*

$$H \;=\; (\mathbb{Z}^m \cap H) \times H\pi_0\sigma, \tag{7.22}$$

where $\pi_0 \colon \mathbb{G}_\Gamma \to \mathbb{G}_{\Gamma_0}$ is the natural projection killing the center of $\mathbb{G}_\Gamma$, and $\sigma \colon H\pi_0 \to H$ is a section of $\pi_{0|H}$.

Proof. Note that the full subgraph Γ_0 must be Droms as well. Let $X = \{x_1, \ldots, x_n\}$ be the (finite) set of vertices of Γ_0 (i.e., let $\mathbb{G}_{\Gamma_0} = \langle X \mid R \rangle$, where $R \subseteq [X, X]$); and let $\mathbb{Z}^m = \langle t_1, \ldots, t_m \mid [t_i, t_j] \; \forall i, j \rangle$. Now, consider the restriction to a subgroup $H \leqslant \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$ of the natural short exact sequence associated to $\mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \mathbb{Z}^m & \longrightarrow & \mathbb{G}_\Gamma & \xrightarrow{\pi_0} & \mathbb{G}_{\Gamma_0} & \longrightarrow & 1 \\ & & \rotatebox{90}{$\leqslant$} & & \rotatebox{90}{$\leqslant$} & & \rotatebox{90}{$\leqslant$} & & \\ 1 & \longrightarrow & \mathbb{Z}^m \cap H & \longrightarrow & H & \longrightarrow & H\pi_0 & \longrightarrow & 1\,, \end{array} \qquad \begin{array}{r} (7.23) \\ \\ (7.24) \end{array}$$

Since $\mathbb{G}_{\Gamma_0}$ is Droms, we know that $H\pi_0 \leqslant \mathbb{G}_{\Gamma_0}$ is again a PC-group. Thus, there exists a (not necessarily finite) subset $Y = \{y_j\}_j \subseteq \mathbb{G}_{\Gamma_0}$ such that $H\pi_0 \simeq \langle Y \mid S \rangle$, where S is some subset of the commutator set $[Y, Y]$.

Now, observe that any map $\sigma : Y \to H$ sending each $y_j \in Y$ back to any of its π_0-preimages in H will respect the relations in S: indeed, for each commutator $[y_i, y_j] \in S$, we have $[(y_i)\sigma, (y_j)\sigma] = [t^{\mathbf{a}_i} y_i, t^{\mathbf{a}_j} y_j] = [y_i, y_j]$ (for certain abelian completions $\mathbf{a}_i, \mathbf{a}_j \in \mathbb{Z}^m$). Therefore, any such map σ defines a (injective) section of the restriction $\pi_{0|H}$ (which we will denote by σ as well). Thus, the short exact sequence (7.24) splits; and $H\pi_0 \simeq H\pi_0\sigma$, for any such section σ.

Moreover, since the kernel of the subextension (7.24) lies in the center of $\mathbb{G}_\Gamma$, the conjugation action is trivial, and the claimed result follows. □

Remark 7.5.5. Note that, for an arbitrary finitely generated subgroup

$$H = \langle t^{\mathbf{b}_1}, \ldots, t^{\mathbf{b}_r}, t^{\mathbf{a}_1} u_1, \ldots, t^{\mathbf{a}_s} u_s \rangle \leqslant \mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}\,, \tag{7.25}$$

where $u_i \neq 1$, for all $i = 1, \ldots, s$ (we have placed the generators of H belonging to the center $\mathbb{Z}^m$ of $\mathbb{G}_\Gamma$ at the beginning of the tuple), we always have:

$$\langle t^{\mathbf{b}_1}, \ldots, t^{\mathbf{b}_r} \rangle \overset{(1)}{\leqslant} \mathbb{Z}^m \cap H = Z(\mathbb{G}_\Gamma) \cap H \overset{(2)}{\leqslant} Z(H)\,, \tag{7.26}$$

but these inclusions are — in general — not strict, since:

(1) a nontrivial product of the last s generators in (7.25) could, in principle, be equal to $t^{\mathbf{c}}$ for some element $\mathbf{c} \notin \langle \mathbf{b}_1, \ldots, \mathbf{b}_r \rangle$;

(2) a generator $t^{\mathbf{a}_i} u_i$ in (7.25) can commute with every other generator in (7.25), but it does not commute with every generator in $\mathbb{G}_\Gamma$ (since u_i contains noncentral generators of $\mathbb{G}_\Gamma$). Consider, for example, the element x in the subgroup $\langle t, x \rangle \leqslant \langle t \mid - \rangle \times \langle x, y \mid - \rangle \simeq \mathbb{Z} \times \mathbb{F}_2$.

Since subgroups of finitely generated free-abelian groups are always of finite rank, the next result follows immediately from Lemma 7.5.4.

Corollary 7.5.6. *Let $\mathbb{G}_\Gamma$ be a Droms group, and H a subgroup of $\mathbb{G}_\Gamma$. Then,*

$$H \textit{ is finitely generated} \;\Leftrightarrow\; H\pi_0 \textit{ is finitely generated}\,.$$

Proof. The claim is a tautology for centerless Droms groups. Otherwise, the direct implication is obviously true; whereas the converse follows from Equation (7.22), and the fact that subgroups of finitely generated free-abelian groups are always finitely generated. □

Therefore, in order to decide whether the intersection of two subgroups H_1, H_2 of a Droms group is finitely generated, it will be enough to find out what happens with the projection $(H_1 \cap H_2)\pi_0$. We will see that the behaviour of the embedding $(H_1 \cap H_2)\pi_0 \leqslant (H_1)\pi_0 \cap (H_2)\pi_0$ is crucial to this end; we describe it below.

Lemma 7.5.7. *Let H_1, H_2 be subgroups of a Droms PC-group $G = \mathbb{G}_\Gamma$, and let π_0 be the projection killing the center of G. Then,*

(i) $(H_1 \cap H_2)\pi_0 \leqslant H_1\pi_0 \cap H_2\pi_0$, *sometimes with strict inclusion;*

(ii) $(H_1 \cap H_2)\pi_0 \trianglelefteq H_1\pi_0 \cap H_2\pi_0$;

(iii) $[\![(H_1)\pi_0 \cap (H_2)\pi_0]\!] \leqslant (H_1 \cap H_2)\pi_0$.

That is, we have the following chain of inclusions:

$$[\![(H_1)\pi_0 \cap (H_2)\pi_0]\!] \leqslant (H_1 \cap H_2)\pi_0 \trianglelefteq (H_1)\pi_0 \cap (H_2)\pi_0 \leqslant \mathop{\ast}_{i=1}^{k} \mathbb{G}_{\Lambda_i}, \tag{7.27}$$

where $\{\Lambda_i\}_{i=1}^{k}$ $(k \geqslant 2)$ are the connected components of $\Gamma_0 = \Gamma \setminus Z(\Gamma)$.

(We denote the commutator subgroup of a group G by $[\![G]\!]$, in order to distinguish it from the set *of commutators $[G, G]$; i.e., $[\![G]\!] = \langle [G, G] \rangle$.)*

Proof. If the Droms group $\mathbb{G}_\Gamma$ is centerless, then π_0 is the identity, and all three claims are obvious. We prove them for connected Droms groups:

(i) This is clear because $(H_1 \cap H_2)\pi_0$ consists of those elements $u \in \mathbb{G}_{\Gamma_0}$ that have a common completion in H_1 and H_2 — i.e., such that $t^{\mathbf{a}} u \in H_1 \cap H_2$, for some $\mathbf{a} \in \mathbb{Z}^m$; whereas $H_1\pi_0 \cap H_2\pi_0$ contains every element $u \in \mathbb{G}_{\Gamma_0}$ with (not necessarily common) completions in both H_1 and H_2 — i.e., such that $t^{\mathbf{a}} u \in H_1$ and $t^{\mathbf{b}} u \in H_2$ for some (not necessarily equal) vectors $\mathbf{a}, \mathbf{b} \in \mathbb{Z}^m$.

(ii) For normality, consider $u \in (H_1 \cap H_2)\pi_0$, and $v \in (H_1)\pi_0 \cap (H_2)\pi_0$. Then, there must exist elements $t^{\mathbf{a}}u \in H_1 \cap H_2$, and $t^{\mathbf{b}_i}v \in H_i$ (for $i = 1, 2$). Now observe that

$$t^{\mathbf{a}}(v^{-1}uv) = v^{-1}(t^{\mathbf{a}} u)v = (t^{\mathbf{b}_i} v)^{-1}(t^{\mathbf{a}} u)(t^{\mathbf{b}_i} v) \in H_i \quad (\text{for } i = 1, 2).$$

Thus, $t^{\mathbf{a}} v^{-1}uv \in H_1 \cap H_2$; and so, $v^{-1}uv \in (H_1 \cap H_2)\pi_0$, as we wanted to prove.

(iii) Finally, take $u, v \in (H_1)\pi_0 \cap (H_2)\pi_0$. Then, there exist elements $t^{a_i}u \in H_i$, $t^{b_i}v \in H_i$ $(i = 1, 2)$. Now, observe that

$$[u, v] = u^{-1}v^{-1}uv = (t^{a_i}u)^{-1}(t^{b_i}v)^{-1}(t^{a_i}u)(t^{b_i}v) \in H_i \quad (\text{for } i = 1, 2).$$

Thus, $[u, v]$ belongs to $H_1 \cap H_2$, and to $(H_1 \cap H_2)\pi_0$, as claimed.

The final chain (7.4) is nothing more that the two previous inclusions, followed by that coming from the fact that centerless Droms graphs are disconnected. □

Lemma 7.5.8. *Let H_1, H_2 be subgroups of a Droms PC-group, and let π_0 be the natural projection killing the center. Then:*

$$(H_1 \cap H_2)\pi_0 \textit{ is abelian} \Leftrightarrow (H_1)\pi_0 \cap (H_2)\pi_0 \textit{ is abelian}.$$

Otherwise,

$$\mathbb{F}_\infty \leqslant [\![(H_1)\pi_0 \cap (H_2)\pi_0]\!] \leqslant (H_1 \cap H_2)\pi_0 .$$

Proof. The converse implication is clear since $(H_1 \cap H_2)\pi_0 \leqslant (H_1)\pi_0 \cap (H_2)\pi_0$. We prove the direct implication by transposition. Recall that $(H_1)\pi_0 \cap (H_2)\pi_0$ is a subgroup of a Droms group, and thus a PC-group as well. So, if $(H_1)\pi_0 \cap (H_2)\pi_0$ is not abelian, then there exist at least one missing edge — say $\{u, v\}$ — in its commutation graph.

Therefore, the (infinitely-generated free) derived subgroup $[\![\mathbb{F}_{\{u,v\}}]\!] \simeq \mathbb{F}_\infty$ of the free subgroup generated by $\{u, v\}$ in $(H_1)\pi_0 \cap (H_2)\pi_0$, must be included in the derived subgroup of $(H_1)\pi_0 \cap (H_2)\pi_0$, and thus in $(H_1 \cap H_2)\pi_0$ which, therefore, is not abelian. This concludes the proof. □

Lemma 7.5.9. *Let H_1, H_2 be subgroups of a Droms PC-group $\mathbb{G}_\Gamma$. Then, if $(H_1)\pi_0 \cap (H_2)\pi_0$ is infinitely generated, then $(H_1 \cap H_2)\pi_0$ (and so $H_1 \cap H_2$) is also infinitely generated.*

Proof. Assume that $H_1\pi_0 \cap H_2\pi_0$ is infinitely generated, but $(H_1 \cap H_2)\pi_0$ is finitely generated, and let us find a contradiction.

Since both subgroups lie within a Droms group, $H_1\pi_0 \cap H_2\pi_0$ is again a PC-group with *infinite* $\{P_4, C_4\}$-free commutation graph, say Δ, and $(H_1 \cap H_2)\pi_0 \leqslant \mathbb{G}_{\Delta_0} \leqslant \mathbb{G}_\Delta$, where Δ_0 is the full subgraph of Δ determined by the vertices appearing in the reduced expressions of elements in $(H_1 \cap H_2)\pi_0$. Note that the assumption of finite generation for $(H_1 \cap H_2)\pi_0$ implies that Δ_0 is finite. Note also that, by construction, Δ_0 is minimal (i.e., for any $x \in V\Delta_0$, there exists and element $g \in (H_1 \cap H_2)\pi_0$ such that $g \notin \mathbb{G}_{\Delta_0 \setminus \{x\}}$).

Recall that then, $(H_1)\pi_0 \cap (H_2)\pi_0$ can not be abelian since, if so, we would have an infinitely generated free-abelian group embedded in the finitely generated PC-group $\mathbb{G}_\Gamma$, which is not possible (see Lemma 6.2.4). So, $(H_1)\pi_0 \cap (H_2)\pi_0$ is not abelian (and infinitely generated by hypothesis); and, from Lemma 7.5.8, $(H_1 \cap H_2)\pi_0$ is not abelian as well. Accordingly, neither the (infinite) graph Δ, nor the (finite) graph Δ_0 can be complete.

Suppose now there is a missing edge between some element $x \in V\Delta_0$ and some vertex $y \in V\Delta \setminus V\Delta_0$. Take an element $g \in (H_1 \cap H_2)\pi_0$ with $g \notin G(\Delta_0 \setminus \{x\})$ and Lemma 7.5.7(ii) would tell us that $y^{-1}gy \in \mathbb{G}_{\Delta_0}$, which is a contradiction.

Hence, in Δ, every vertex from Δ_0 is connected to every vertex outside Δ_0. But then, two non-adjacent vertices x_1, x_2 from Δ_0 (there is at least one pair because Δ_0 is not complete) together with two non-adjacent vertices y_1, y_2 from $\Delta \setminus \Delta_0$ (there must be many because $\Delta \setminus \Delta_0$ is infinite and $\mathbb{Z}^\infty$ does not embed into $\mathbb{G}_\Gamma$) form a copy of C_4, a square, as a full subgraph of Δ; this is again a contradiction with Δ being Droms.

So, the starting assumption must be false, and the proof is concluded. □

7.5.1 Inductive theorems for connected Droms groups

Let us now prove Theorem 7.3.3; and afterwards, with an extension of the same arguments, we shall put cosets into the picture and prove Theorem 7.3.4.

Theorem 7.3.3. *Let* G *be a Droms PC-group. If* G *satisfies* SIP, *then* $\mathbb{Z}^m \times G$ *also satisfies* SIP. □

Proof. Recall that it is enough to prove the theorem when G is a Droms group with disconnected commuting graph, say Γ_0 (Remark 7.5.1). So, we will assume $G = \mathbb{G}_{\Gamma_0} = \langle X \mid R \rangle$ (a graphic presentation) for the rest of the proof.

We are given finite sets of generators for two subgroups H_1, H_2 of a Droms group with primary decomposition $\mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$.

If we project the given generators to $\mathbb{G}_{\Gamma_0}$, and then apply Proposition 7.2.10 on the projected generating sets, we can compute bases for $H_1\pi_0$ and $H_2\pi_0$ respectively. Now, for each such basis, say $u_1, \ldots, u_{n_i}$ of $H_i\pi_0$, compute its completion, say $t^{a_1} u_1, \ldots, t^{a_{n_i}} u_{n_i}$, in H_i. This can be easily done by taking the words expressing the u_i's in terms of the generators of $H_i\pi$ projection of the initial generators, and recomputing them on the original generators given for H_i.

Now, for each of the original generators of H_i, say $t^c v$, we can write $v \in H_i\pi$ in terms of the basis $u_1, \ldots, u_{n_i}$, say $v = v(u_1, \ldots, u_{n_i})$ and compute

$$v(t^{a_1} u_1, \ldots, t^{a_{n_i}} u_{n_i}) \;=\; t^{d}\, v(u_1, \ldots, u_{n_i}) \;=\; t^{d}\, v\,.$$

We obtain $t^c v, t^d v \in H_i$, and so $t^{c-d} \in H_i \cap \mathbb{Z}^m$. Repeating this operation for each generator of H_i, we get a collection of vectors generating $H_i \cap \mathbb{Z}^m$. Cleaning them with the use of standard linear algebra, we obtain a free abelian basis of $H_i \cap \mathbb{Z}^m$, say $\{t^{b_1}, \ldots, t^{b_{m_1}}\}$. It is then immediate to see that $\{t^{b_1}, \ldots, t^{b_{m_1}}, t^{a_1} u_1, \ldots, t^{a_{n_i}} u_{n_i}\}$ is a basis of H_i, with the extra properties that $\{t^{b_1}, \ldots, t^{b_{m_1}}\}$ is a free-abelian basis of $H_i \cap \mathbb{Z}^m$, and $\{u_1, \ldots, u_{n_i}\}$ is a basis of $H_i \pi$.

So, we can assume the given generators to be of the form:

$$\begin{aligned} H_1 &= \langle t^{b_1}, \ldots, t^{b_{m_1}}, t^{a_1} u_1, \ldots, t^{a_{n_1}} u_{n_1} \rangle \leqslant \mathbb{G}_\Gamma \,, \\ H_2 &= \langle t^{b'_2}, \ldots, t^{b'_{m_2}}, t^{a'_2} u'_2, \ldots, t^{a'_{n_2}} u'_{n_2} \rangle \leqslant \mathbb{G}_\Gamma \,, \end{aligned} \tag{7.28}$$

where $\{t^{b_1}, \ldots, t^{b_{m_1}}\}$ is a free-abelian basis for $L_1 = H_1 \cap \mathbb{Z}^m$, and $\{u_1, \ldots, u_{n_1}\}$ is a basis of $H_1 \pi$ (resp., $\{t^{b'_1}, \ldots, t^{b'_{m_2}}\}$ is a free-abelian basis for $L_2 = H_2 \cap \mathbb{Z}^m$, and $\{u'_1, \ldots, u'_{n_2}\}$ is a basis of $H_2 \pi$)

Let us first make the decision on whether $H_1 \cap H_2$ is finitely generated or not (equivalently, whether $(H_1 \cap H_2)\pi_0 \leqslant \mathbb{G}_{\Gamma_0}$ is finitely generated or not). In the affirmative case, we shall then compute a set of generators for this intersection.

We have generators (indeed bases) $U = \{u_1, \ldots, u_{n_1}\}$, and $U' = \{u'_1, \ldots, u'_{n_2}\}$ — written as words on $X = \{x_1, \ldots, x_n\}$ — for the respective subgroups $H_1\pi_0$, and $H_2\pi_0$ of the Droms group $\mathbb{G}_{\Gamma_0}$. Let Δ_1 and Δ_2 be the respective commutation graphs for them; i.e.,

$$H_1\pi_0 \simeq \mathbb{G}_{\Delta_1} \quad \text{and} \quad H_2\pi_0 \simeq \mathbb{G}_{\Delta_2} \,,$$

with U and U' corresponding to the vertices of Δ_1 and Δ_2 respectively.

Applying property SIP from the hypothesis to these two subgroups of $\mathbb{G}_{\Gamma_0}$, we can decide whether $H_1\pi_0 \cap H_2\pi_0$ is finitely generated or not. If it is infinitely generated then (by Lemma 7.5.9) so is $(H_1 \cap H_2)\pi_0$; hence (by Corollary 7.5.6) $H_1 \cap H_2$ is also infinitely generated, and we are done.

Thus, for the rest of the proof, assume that $H_1\pi_0 \cap H_2\pi_0 \leqslant \mathbb{G}_{\Gamma_0}$ is finitely generated, in which case the hypothesis provides us with a finite set of generators for it. Applying again Proposition 7.2.10, we can compute its commutation graph — say Δ_3, with basis $W = \{w_1, \ldots, w_{n_3}\}$ — for $H_1\pi_0 \cap H_2\pi_0$; i.e.,

$$H_1\pi_0 \cap H_2\pi_0 = \mathbb{G}_{\Delta_3} = \langle w_1, \ldots, w_{n_3} \rangle \leqslant \mathbb{G}_{\Gamma_0} \,,$$

where the w_i's are words on $X = \{x_1, \ldots, x_n\}$.

Recall that in order to decide about finite generation of $H_1 \cap H_2$, it will be enough to decide about finite generation of $(H_1 \cap H_2)\pi_0$, which is a normal subgroup of the subgroup $H_1\pi_0 \cap H_2\pi_0$ described above.

Also note that the w_i's in W can be algorithmically written as words $w_i = \omega_i(u_1, \ldots, u_{n_1})$ on U (resp., $w_i = \omega_i(u'_1, \ldots, u'_{n_2})$ on U') just enumerating all products of the elements in U (resp., U'), and waiting to hit $w_1, \ldots, w_{n_3}$ (we can algorithmically recognize them with the help of the word problem for $\mathbb{G}_{\Gamma_0}$, but recall that only the YES part is needed). These words constitute the formal expressions of the inclusions $\iota_1\colon H_1\pi_0 \cap H_2\pi_0 \rightarrowtail H_1\pi_0$, and $\iota_2\colon H_1\pi_0 \cap H_2\pi_0 \rightarrowtail H_2\pi_0$ in terms of the corresponding bases (i.e., $\iota_1\colon \mathbb{G}_{\Delta_3} \rightarrowtail \mathbb{G}_{\Delta_1}$, and $\iota_2\colon \mathbb{G}_{\Delta_3} \rightarrowtail \mathbb{G}_{\Delta_2}$).

Abelianizing these two morphisms, we obtain the integral matrices $\mathbf{P_1}$ (of size $n_3 \times n_1$), and $\mathbf{P_2}$ (of size $n_3 \times n_2$) and complete the upper half of the diagram in Figure 7.25, where the ρ_i's are the corresponding abelianization maps. Note that, even though ι_1 and ι_2 are injective, their abelianizations $\mathbf{P_1}$ and $\mathbf{P_2}$ need not be (n_3 could very well be bigger than n_1 or n_2).

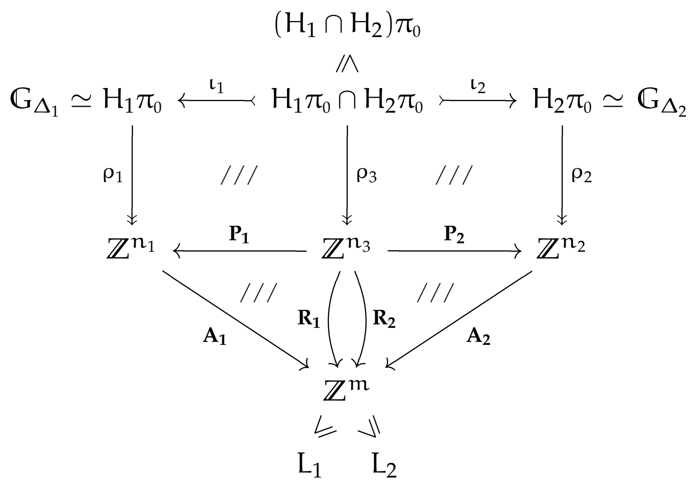

Fig. 7.25: Intersection diagram for subgroups of Droms groups

Now, we can recompute the words ω_i (resp., ω'_i) as words on the $(t^{\mathbf{a}_i} u_i)$'s (resp., on the $(t^{\mathbf{a}'_i} u'_i)$'s) to get particular preimages of the w_i's in H_1 (resp., H_2). Namely,

$$\begin{aligned}
\omega_i(t^{\mathbf{a}_1} u_1, \ldots, t^{\mathbf{a}_{n_1}} u_{n_1}) &= t^{\boldsymbol{\omega}_i \mathbf{A_1}}\, \omega_i(u_1, \ldots, u_{n_1}) = t^{\boldsymbol{\omega}_i \mathbf{A_1}}\, w_i \in H_1\,, \\
\omega'_i(t^{\mathbf{a}'_1} u'_1, \ldots, t^{\mathbf{a}'_{n_2}} u'_{n_2}) &= t^{\boldsymbol{\omega}'_i \mathbf{A_2}}\, \omega'_i(u'_1, \ldots, u'_{n_2}) = t^{\boldsymbol{\omega}'_i \mathbf{A_2}}\, w_i \in H_2\,,
\end{aligned}$$

where $\boldsymbol{\omega}_i = (\omega_i)^{\mathrm{ab}}$, $\boldsymbol{\omega}'_i = (\omega'_i)^{\mathrm{ab}}$; and the integral matrices

$$\mathbf{A_1} = \begin{pmatrix} \mathbf{a}_1 \\ \vdots \\ \mathbf{a}_{n_1} \end{pmatrix} \quad \text{and} \quad \mathbf{A_2} = \begin{pmatrix} \mathbf{a}'_1 \\ \vdots \\ \mathbf{a}'_{n_2} \end{pmatrix}$$

have sizes $n_1 \times m$ and $n_2 \times m$, respectively.

Hence, the abelian completions of $w_i \in H_1\pi_0 \cap H_2\pi_0$ in H_1 and H_2 are the linear varieties:

$$\begin{aligned} \mathcal{C}_{H_1}(w_i) &= \boldsymbol{\omega_i A_1} + L_1 = w_i \iota_1 \rho_1 \mathbf{A_1} + L_1 = w_i \rho_3 \mathbf{R_1} + L_1 \,, \\ \mathcal{C}_{H_2}(w_i) &= \boldsymbol{\omega'_i A_2} + L_2 = w_i \iota_2 \rho_2 \mathbf{A_2} + L_2 = w_i \rho_3 \mathbf{R_2} + L_2 \,, \end{aligned}$$

where $L_i = \mathbb{Z}^m \cap H_i$, and we have used the commutation $\iota_i \rho_i = \rho_3 \mathbf{P_i}$ together with the definition $\mathbf{R_i} \coloneqq \mathbf{P_i A_i}$, for $i = 1, 2$ (see diagram in Figure 7.25).

To finish our argument, it is crucial to understand which elements of $H_1\pi_0 \cap H_2\pi_0$ do belong to (the normal subgroup) $(H_1 \cap H_2)\pi_0$. They are, precisely, those whose H_1-completion and H_2-completion (which are linear varieties in $\mathbb{Z}^m$ of directions L_1 and L_2, respectively) do intersect:

$$\begin{aligned} (H_1 \cap H_2)\pi_0 &= \{ w \in H_1\pi_0 \cap H_2\pi_0 \mid \mathcal{C}_{H_1}(w) \cap \mathcal{C}_{H_2}(w) \neq \varnothing \} \\ &= \{ w \in H_1\pi_0 \cap H_2\pi_0 \mid (w\rho_3 \mathbf{P_1 A_1} + L_1) \cap (w\rho_3 \mathbf{P_2 A_2} + L_2) \neq \varnothing \} \\ &= (\{ \mathbf{d} \in \mathbb{Z}^{n_3} \mid (\mathbf{d}\mathbf{R_1} + L_1) \cap (\mathbf{d}\mathbf{R_2} + L_2) \neq \varnothing \})\, \rho_3^{\leftarrow} \qquad (7.29) \\ &= (\{ \mathbf{d} \in \mathbb{Z}^{n_3} \mid \mathbf{d}(\mathbf{R_1} - \mathbf{R_2}) \in L_1 + L_2 \})\rho_3^{\leftarrow} \\ &= (L_1 + L_2)(\mathbf{R_1} - \mathbf{R_2})^{\leftarrow} \rho_3^{\leftarrow} = M\rho_3^{\leftarrow} \,, \end{aligned}$$

where $M \coloneqq (L_1 + L_2)(\mathbf{R_1} - \mathbf{R_2})^{\leftarrow}$ (denoting the preimage of $L_1 + L_2$ by $\mathbf{R_1} - \mathbf{R_2}$) is a subgroup of $\mathbb{Z}^{n_3}$ for which we can easily compute a free-abelian basis using linear algebra (note that the data of the problem allows us to compute L_1, L_2, $\mathbf{R_1}$ and $\mathbf{R_2}$).

From this computations one can see clearly the inclusions stated in Lemma 7.5.7:

$$[\![H_1\pi_0 \cap H_2\pi_0]\!] \leqslant (H_1 \cap H_2)\pi_0 \trianglelefteqslant H_1\pi_0 \cap H_2\pi_0 \,.$$

At this point, we can decide whether $(H_1 \cap H_2)\pi_0$ is finitely generated or not, by distinguishing two cases.

If Δ_3 is complete (this includes the case where Δ_3 is empty and $n_3 = 0$), then $H_1\pi_0 \cap H_2\pi_0 \simeq \mathbb{Z}^{n_3}$ is abelian, ρ_3 is the identity, and $(H_1 \cap H_2)\pi_0 = (L_1 + L_2)(\mathbf{R_1} - \mathbf{R_2})^{\leftarrow} = M$ is always finitely generated, and a basis for it is easily computable with basic linear algebra techniques.

So, assume Δ_3 is not complete. Since it is a Droms graph, it will have a primary decomposition, say $\Delta_3 = K_{n_4} \vee \Delta_5$, where $n_4 \geqslant 0$, and Δ_5 is Droms again, disconnected, with $|V\Delta_5| = n_5 = n_3 - n_4 \geqslant 2$.

Let us denote $VK_{n_4} \eqqcolon \{z_1, \ldots, z_{n_4}\}$, and $V\Delta_5 \eqqcolon \{y_1, \ldots, y_{n_5}\}$, the vertices of Δ_3.

Algebraically, we have that $H_1\pi_0 \cap H_2\pi_0 \simeq \mathbb{G}_{\Delta_3} = \mathbb{Z}^{n_4} \times \mathbb{G}_{\Delta_5}$, where $n_4 \geqslant 0$, and $\mathbb{G}_{\Delta_5} \neq 1$ decomposing as a non-trivial free product. Furthermore, note that the normal subgroup $(H_1 \cap H_2)\pi_0 \trianglelefteqslant \mathbb{G}_{\Delta_3}$ is *not* contained in $\mathbb{Z}^{n_4}$ (taking two vertices, say y_i, y_j, in different components of Δ_5, Lemma 7.5.7(iii) tells us that $1 \neq [y_i, y_j] \in (H_1 \cap H_2)\pi_0$).

In this situation, the abelianization map $\rho_3\colon \mathbb{G}_{\Delta_3} \twoheadrightarrow \mathbb{Z}^{n_3}$ is the identity on the center $\mathbb{Z}^{n_4}$ of $\mathbb{G}_{\Delta_3}$ and so, can be written in the form

$$\begin{array}{rcl} \rho_3 = \mathrm{id} \times \rho_5\colon \mathbb{G}_{\Delta_3} = \mathbb{Z}^{n_4} \times \mathbb{G}_{\Delta_5} & \twoheadrightarrow & \mathbb{Z}^{n_4} \times \mathbb{Z}^{n_5} = \mathbb{Z}^{n_3} \\ (\mathbf{c}, \nu) & \mapsto & (\mathbf{c}, \mathbf{v})\ , \end{array} \tag{7.30}$$

where, as usual, $\mathbf{v}$ denotes the abelianization, $\mathbf{v} = \nu^{\mathsf{ab}} \in \mathbb{Z}^{n_5}$. Of course, if $n_4 = 0$ then $\rho_5 = \rho_3$ and the decomposition in (7.30) is vacuous.
Now consider the image of $(H_1 \cap H_2)\pi_0$ under the projection $\pi_1\colon \mathbb{Z}^{n_4} \times \mathbb{G}_{\Delta_5} \twoheadrightarrow \mathbb{G}_{\Delta_5}$ (which is non-trivial since $(H_1 \cap H_2)\pi_0 \not\leqslant \mathbb{Z}^{n_4}$). We have $1 \neq (H_1 \cap H_2)\pi_0\pi_1 \trianglelefteq \mathbb{G}_{\Delta_5}$, a non-trivial normal subgroup in a group which decomposes as a non-trivial free product. Therefore,

$$\begin{aligned} H_1 \cap H_2 \text{ is finitely generated} &\Leftrightarrow (H_1 \cap H_2)\pi_0 \text{ is finitely generated} \\ &\Leftrightarrow (H_1 \cap H_2)\pi_0\pi_1 \text{ is finitely generated} \\ &\Leftrightarrow (H_1 \cap H_2)\pi_0\pi_1 \trianglelefteq_{\mathsf{fi}} \mathbb{G}_{\Delta_5} \\ &\Leftrightarrow M\rho_3^{\leftarrow}\pi_1 \trianglelefteq_{\mathsf{fi}} \mathbb{G}_{\Delta_5} \\ &\Leftrightarrow M\pi_1^{\mathsf{ab}}\rho_5^{\leftarrow} \trianglelefteq_{\mathsf{fi}} \mathbb{G}_{\Delta_5} \\ &\Leftrightarrow M\pi_1^{\mathsf{ab}} \trianglelefteq_{\mathsf{fi}} \mathbb{Z}^{n_5} \\ &\Leftrightarrow \mathrm{rk}\left(M\pi_1^{\mathsf{ab}}\right) = n_5\,. \end{aligned} \tag{7.31}$$

(Note that:

- the first and second equivalences are applications of Corollary 7.5.6.
- The third equivalence is an application of the following theorem by Baumslag in [Bau66, Section 6].

  Theorem 7.5.10. *Let* G *be the free product of two non-trivial groups. Let H be a finitely generated subgroup containing a non-trivial normal subgroup of* G. *Then* H *is of finite index in* G. □

- The fifth equivalence is correct because $\pi_1\rho_5 = \rho_3\pi_1^{\mathsf{ab}}$, and all of them are surjective maps.
- The sixth equivalence is correct because backwards the epimorphism $\rho_5\colon \mathbb{G}_{\Delta_5} \twoheadrightarrow \mathbb{Z}^{n_5}$, a subgroup $M\pi_1^{\mathsf{ab}} \leqslant \mathbb{Z}^{n_5}$ is of finite index if and only if its full preimage $M\pi_1^{\mathsf{ab}}\rho_5^{\leftarrow}$ is of finite index in $\mathbb{G}_{\Delta_5}$ — in which case, furthermore, the two indices do coincide; namely, $[\mathbb{Z}^{n_5} : M\pi_1^{\mathsf{ab}}] = [\mathbb{G}_{\Delta_5} : M\pi_1^{\mathsf{ab}}\rho_5^{\leftarrow}]$.)

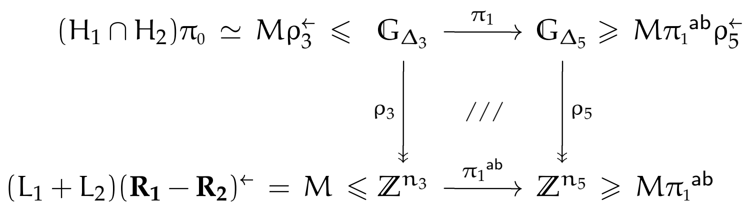

Fig. 7.26: The map π_1 and its abelianization

It only remains to realize that, since we know L_1, L_2, $\mathbf{R_1}$, and $\mathbf{R_2}$ (they are all clearly computable from the input in Equation (7.28)), the last condition in (7.31) — namely, whether $\operatorname{rk}(M\pi_1^{ab}) = n_5 = |V\Delta_5|$ — can effectively be checked with linear algebra (recall that, from Proposition 7.2.10, a basis, and thus the rank n_5 of $\mathbb{G}_{\Delta_5}$, can be computed from the finite generating set obtained from the induction hypothesis on the subgroups $H_1\pi_0$, and $H_2\pi_0$).
Hence, we can algorithmically decide whether $H_1 \cap H_2$ is finitely generated or not (ultimately, in terms of some integral matrix having the correct rank). This solves the decision part of SIP for Droms groups.

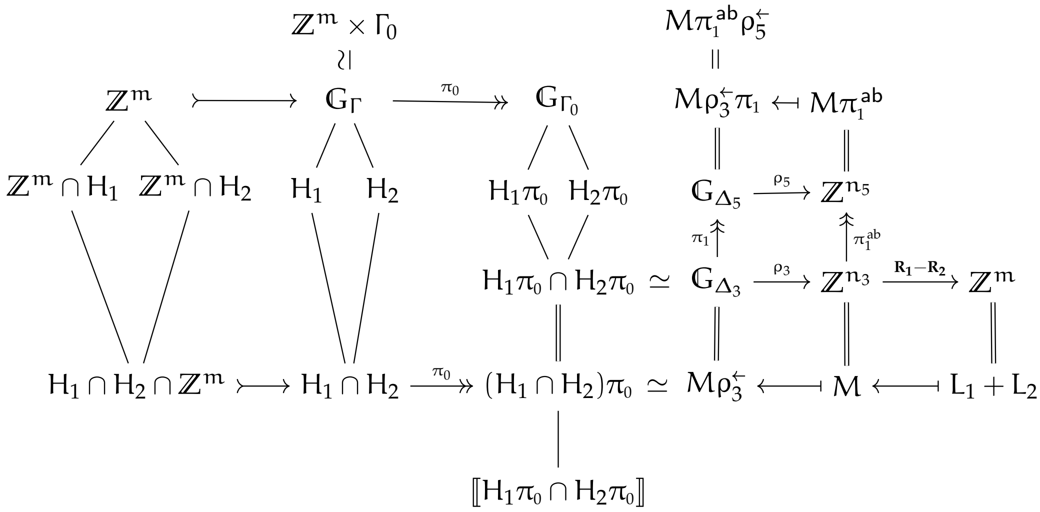

Fig. 7.27: Roadmap to SIP induction for connected Droms groups

The second part of the proof consists in computing a (finite) set of generators for $H_1 \cap H_2$ assuming it is finitely generated; i.e., assuming we are in the situation described by any of the conditions in (7.31); e.g. $\operatorname{rk} M\pi_1^{ab} = n_5$, where $M = (L_1 + L_2)(\mathbf{R_1} - \mathbf{R_2})^{\leftarrow} \leqslant \mathbb{Z}^{n_3}$. We will take advantage of Theorem 7.5.10 to deduce finite generating sets from the corresponding Schreier graphs.
Note however, that $(H_1 \cap H_2)\pi \simeq M\rho_3^{\leftarrow}$ can be finitely generated, even when M is of infinite index in $\mathbb{Z}^{n_3}$ (recall that the actual condition is $M\pi_1^{ab}$ being of finite index in $\mathbb{Z}^{n_5}$). So, we will start computing a (finite) family C of coset representatives of $\mathbb{Z}^{n_5}$ modulo $M\pi_1^{ab}$ (recall that $M\pi_1^{ab}$ is known, and thus such a

family is computable), and then choose ρ_5-preimages in $\mathbb{G}_{\Delta_5}$, say $\{v_1,\ldots,v_r\}$, where $r = [\mathbb{G}_{\Delta_5} : M\pi_1^{ab}\rho_5^{\leftarrow}] = [\mathbb{Z}^{n_5} : M\pi_1^{ab}]$. This can be done, for example, choosing for each vector $\mathbf{a} = (a_1,\ldots,a_{n_5}) \in \mathbb{Z}^{n_5}$, the element $y_1^{a_1}\cdots y_{n_5}^{a_{n_5}} \in \mathbb{G}_{\Delta_5}$.
Now, we shall construct the Schreier graph of the subgroup

$$(H_1 \cap H_2)\pi_0\pi_1 = M\rho_3^{\leftarrow}\pi_1 = M\pi_1^{ab}\rho_5^{\leftarrow} \leqslant_{fi} \mathbb{G}_{\Delta_5}$$

with respect to $V\Delta_5$, in the following way:

- Draw as vertices the cosets $[v_1],\ldots,[v_r]$.
- For every $[v_i]$ ($i = 1,\ldots,r$), and every y_j ($j = 1,\ldots,n_5$), Draw an edge labelled y_j from $[v_i]$ to $[v_i y_j]$.

  Here, we have to algorithmically recognize which is the coset $[v_i y_j]$ from our list of vertices, but this is easy since:

$$[v_i y_j] = [v_k] \;\Leftrightarrow\; v_i y_j v_k^{-1} \in M\pi_1^{ab}\rho_5^{\leftarrow} \;\Leftrightarrow\; (v_i y_j v_k^{-1})\rho_5 \in M\pi_1^{ab}.$$

After having the full picture of the Schreier graph for $(H_1 \cap H_2)\pi_0\pi_1 \leqslant_{fi} \mathbb{G}_{\Delta_5}$, we can obtain a finite set of generators for $(H_1 \cap H_2)\pi_0\pi_1$ just reading the basic labels of the petals $\mathbf{T}[e]$ corresponding to the arcs outside a chosen maximal tree $\mathbf{T}$ (this is a general fact, see Figure 5.49). These will be words on $V\Delta_5 = \{y_1,\ldots,y_{n_5}\}$, i.e., elements of $\mathbb{G}_{\Delta_3}$ not using the central vertices $\{z_1,\ldots,z_{n_4}\}$.
The next step is to elevate the obtained generators of $(H_1 \cap H_2)\pi_0\pi_1$ to generators of $(H_1 \cap H_2)\pi_0$, pulling them back through π_1. For each one of them, say $g(y_1,\ldots,y_{n_5})$, we look for its preimages in $(H_1 \cap H_2)\pi_0$; they all are of the form

$$z_1^{\lambda_1}\cdots z_{n_4}^{\lambda_{n_4}} g(y_1,\ldots,y_{n_5}), \tag{7.32}$$

where the unknowns $\lambda_1,\ldots,\lambda_{n_4} \in \mathbb{Z}$ can be found by solving the system of linear equations coming from the imposition of the condition:

$$z_1^{\lambda_1}\cdots z_{n_4}^{\lambda_{n_4}} g(y_1,\ldots,y_{n_5}) \in (H_1 \cap H_2)\pi_0 = M\rho_3^{\leftarrow}.$$

That is, $(\lambda_1,\ldots,\lambda_{n_4},|g|_1,\ldots,|g|_{n_5}) \in M = (L_1 + L_2)(\mathbf{R_1} - \mathbf{R_2})^{\leftarrow}$, or equivalently:

$$(\lambda_1,\ldots,\lambda_{n_4},|g|_1,\ldots,|g|_{n_5})(\mathbf{R_1} - \mathbf{R_2}) \in L_1 + L_2.$$

For each such $g(y_1,\ldots,y_{n_5})$, we compute a particular preimage of the form (7.32) and put them all to constitute a set of generators for $(H_1 \cap H_2)\pi_0$, adding on top of them a free basis for

$$\ker \pi_1 \cap (H_1 \cap H_2)\pi_0 = \left\{ z_1^{\lambda_1}\cdots z_{n_4}^{\lambda_{n_4}} \;:\; (\lambda_1,\ldots,\lambda_{n_4},0,\ldots,0)(\mathbf{R_1} - \mathbf{R_2}) \in L_1 + L_2 \right\}.$$

Finally, we have to lift the just computed set of generators for $(H_1 \cap H_2)\pi_0$, to $H_1 \cap H_2$:

- Write each generator $h_j \in (H_1 \cap H_2)\pi_0 \leqslant (H_\nu)\pi_0$ $(\nu = 1,2)$ as a word $\omega_j(u_1, \ldots, u_{n_1})$ (resp., $\omega'_j(u'_1, \ldots, u'_{n_2})$) in the original basis U for $H_1\pi_0$ (resp., U' for $H_2\pi_0$).

- Evaluate each ω_j and ω'_j in the noncentral part of original bases (7.28) for H_1 and H_2 respectively, to obtain vectors $\mathbf{c}_j, \mathbf{c}'_j \in \mathbb{Z}^m$ such that:

$$\omega_j(t^{\mathbf{a}_1} u_1, \ldots, t^{\mathbf{a}_{n_1}} u_{n_1}) = t^{\mathbf{c}_j} \omega_j(u_1, \ldots, u_{n_1}) \in H_1 ,$$
$$\omega'_j(t^{\mathbf{a}'_1} u'_1, \ldots, t^{\mathbf{a}'_{n_2}} u'_{n_2}) = t^{\mathbf{c}'_j} \omega'_j(u'_1, \ldots, u'_{n_2}) \in H_2 .$$

Now, for each j, compute a vector $\mathbf{d}_j \in (\mathbf{c}_j + L_1) \cap (\mathbf{c}'_j + L_2)$ (note that all these intersections must be nonempty because $h_j \in (H_1 \cap H_2)\pi_0$), and consider the element $t^{\mathbf{d}_j} h \in H_1 \cap H_2$. All these elements $t^{\mathbf{d}_j} h_j$, together with a free-abelian basis for $H_1 \cap H_2 \cap \mathbb{Z}^m = (H_1 \cap \mathbb{Z}^m) \cap (H_2 \cap \mathbb{Z}^m) = L_1 \cap L_2$ constitute the desired set of generators for $H_1 \cap H_2$.
Therefore, $\mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$ satisfies both the decision, and the search SIP properties, and the proof is completed. □

Below, we extend the previous arguments to prove Theorem 7.3.4.

Theorem 7.3.4. *Let* G *be a Droms PC-group. If* G *satisfies* ESIP, *then* $\mathbb{Z}^m \times G$ *also satisfies* ESIP. □

Proof. Let us use the same notation as above: $G = \mathbb{G}_{\Gamma_0}$, and also assume Γ_0 disconnected (using the argument in Remark 7.5.1). We want to prove that if $\mathbb{G}_{\Gamma_0}$ satisfies ESIP, then $\mathbb{G}_\Gamma = \mathbb{Z}^m \times \mathbb{G}_{\Gamma_0}$ also satisfies ESIP.

We are given finite sets of generators for two subgroups $H_1, H_2 \leqslant \mathbb{G}_\Gamma$, and two elements $t^{\mathbf{a}} u$, $t^{\mathbf{a}'} u' \in \mathbb{G}_\Gamma$. By Proposition 7.2.10, we can assume the generators to be bases of the respective subgroups, like in (7.28). Note that we can immediatly apply Theorem 7.3.3 to solve the SIP part of ESIP; see Definition (7.1).

So, we can assume that $H_1 \cap H_2$ is finitely generated, and that we have already computed a set of generators, say $\{v_1, \ldots, v_p\}$, following the proof of Theorem 7.3.3. Here is where $t^{\mathbf{a}} u$ and $t^{\mathbf{a}'} u'$ start playing its role: we have to decide whether the coset intersection $(t^{\mathbf{a}} u)H_1 \cap (t^{\mathbf{a}'} u')H_2$ is empty or not. Note that

$$(t^{\mathbf{a}} u)H_1 \cap (t^{\mathbf{a}'} u')H_2 = \varnothing \;\Leftrightarrow\; \Big((t^{\mathbf{a}} u)H_1 \cap (t^{\mathbf{a}'} u')H_2\Big)\,\pi_0 = \varnothing,$$

and that

$$\Big((t^{\mathbf{a}} u)H_1 \cap (t^{\mathbf{a}'} u')H_2\Big)\,\pi_0 \subseteq ((t^{\mathbf{a}} u)H_1)\pi_0 \cap ((t^{\mathbf{a}'} u')H_2)\pi_0 = u(H_1\pi_0) \cap u'(H_2\pi_0) .$$

Then, since $H_1 \cap H_2$ is finitely generated, we know from Lemma 7.5.9 that $H_1\pi_0 \cap H_2\pi_0$ is finitely generated as well. So, an application of ESIP inductive hypothesis to $H_1\pi_0, H_2\pi_0 \leqslant \mathbb{G}_{\Gamma_0}$ and $u, u' \in \mathbb{G}_{\Gamma_0}$, tells us (algorithmically) whether the coset intersection $u(H_1\pi_0) \cap u'(H_2\pi_0)$ is empty or not. If it is empty, then $((t^{\mathbf{a}}\, u)H_1 \cap (t^{\mathbf{a}'} u')H_2)\pi_0$ is empty as well, and we are done.

Otherwise, $u(H_1\pi_0) \cap u'(H_2\pi_0) \neq \varnothing$, and we can assume — by induction hypothesis — that we have computed an element $v_0 \in u(H_1\pi_0) \cap u'(H_2\pi_0)$ as a word on $V\Gamma_0 = \{x_1, \ldots, x_n\}$; and thus $u(H_1\pi_0) \cap u'(H_2\pi_0) = v_0(H_1\pi_0 \cap H_2\pi_0)$.

Observe that $((t^{\mathbf{a}}\, u)H_1 \cap (t^{\mathbf{a}'} u')H_2)\pi_0$ consists precisely of those elements $v_0 w$, with $w \in H_1\pi_0 \cap H_2\pi_0$, for which there exists a vector $\mathbf{c} \in \mathbb{Z}^m$ such that $t^{\mathbf{c}} v_0 w \in (t^{\mathbf{a}}\, u)H_1 \cap (t^{\mathbf{a}'} u')H_2$; that is, such that $t^{\mathbf{c}-\mathbf{a}}\, u^{-1} v_0 w \in H_1$, and $t^{\mathbf{c}-\mathbf{a}'}\, (u')^{-1} v_0 w \in H_2$. In the language of completions, this is the same as saying that $\mathbf{c} - \mathbf{a} \in \mathcal{C}_{H_1}(u^{-1} v_0 w)$, and $\mathbf{c} - \mathbf{a}' \in \mathcal{C}_{H_2}((u')^{-1} v_0 w)$; i.e., that these linear varieties do intersect.

Therefore, $((t^{\mathbf{a}}\, u)H_1 \cap (t^{\mathbf{a}'} u')H_2)\pi_0 = \varnothing$ if and only if

$$\forall w \in H_1\pi_0 \cap H_2\pi_0\,,\ \big(\mathbf{a} + \mathcal{C}_{H_1}(u^{-1} v_0 w)\big) \cap \big(\mathbf{a}' + \mathcal{C}_{H_2}((u')^{-1} v_0 w)\big) = \varnothing\,.$$

But then, for an arbitrary word $w = \omega(w_1, \ldots, w_{n_3})$, with $|\omega|_i = \lambda_i$, and choosing $\mathbf{c} \in \mathcal{C}_{H_1}(u^{-1} v_0)$, $\mathbf{c}' \in \mathcal{C}_{H_2}((u')^{-1} v_0)$, and $\mathbf{d_i} \in \mathcal{C}_{H_1}(w_i)$, $\mathbf{d}'_\mathbf{i} \in \mathcal{C}_{H_2}(w_i)$, for $i = 1, \ldots, n_3$; we have:

$$\begin{aligned}
\big(\mathbf{a} + \mathcal{C}_{H_1}(u^{-1} v_0 w)\big) \cap \big(\mathbf{a}' + \mathcal{C}_{H_2}((u')^{-1} v_0 w)\big) &= \\
&= \big(\mathbf{a} + \mathbf{c} + \mathcal{C}_{H_1}(w)\big) \cap \big(\mathbf{a}' + \mathbf{c}' + \mathcal{C}_{H_2}(w)\big) \\
&= \left(\mathbf{a} + \mathbf{c} + \sum_{i=1}^{n_3} \lambda_i \mathcal{C}_{H_1}(w_i)\right) \cap \left(\mathbf{a}' + \mathbf{c}' + \sum_{i=1}^{n_3} \lambda_i \mathcal{C}_{H_2}(w_i)\right) \\
&= \left(\mathbf{a} + \mathbf{c} + \sum_{i=1}^{n_3} \lambda_i \mathbf{d_i} + L_1\right) \cap \left(\mathbf{a}' + \mathbf{c}' + \sum_{i=1}^{n_3} \lambda_i \mathbf{d}'_\mathbf{i} + L_2\right)\,.
\end{aligned}$$

Hence, the coset intersection $\big((t^{\mathbf{a}}\, u)H_1 \cap (t^{\mathbf{a}'} u')H_2\big)\,\pi_0$ is empty if and only if:

$$\forall \lambda_1, \ldots, \lambda_{n_3} \in \mathbb{Z}\,,\ \ (\mathbf{a} - \mathbf{a}' + \mathbf{c} - \mathbf{c}') + \sum_{i=1}^{n_3} \lambda_i (\mathbf{d_i} - \mathbf{d}'_\mathbf{i}) \notin L_1 + L_2\,,$$

or, equivalently:

$$(\mathbf{a} - \mathbf{a}' + \mathbf{c} - \mathbf{c}') + \langle\, \mathbf{d_1} - \mathbf{d}'_\mathbf{1}, \ldots, \mathbf{d}_{\mathbf{n_3}} - \mathbf{d}'_{\mathbf{n_3}} \,\rangle \cap (L_1 + L_2) = \varnothing\,.$$

This can be effectively decided, so the proof is concluded. □

Bibliography

[Adi57a] S. I. Adian. "Finitely Presented Groups and Algorithms". *Akad. Nauk Armyan. SSR Dokl.* 117 (1957), pp. 9–12 (cit. on p. 62).

[Adi57b] S. I. Adian. "The Unsolvability of Certain Algorithmic Problems in the Theory of Groups". *Trudy Moskov. Mat. Obsc.* 6 (1957), pp. 231–298 (cit. on p. 62).

[AH89] I. J. Aalbersberg and H. J. Hoogeboom. "Characterizations of the Decidability of Some Problems for Regular Trace Languages". *Mathematical systems theory* 22.1 (Dec. 1989), pp. 1–19 (cit. on p. 210).

[AMS11] Y. Antolín, A. Martino, and I. Schwabrow. "Kurosh Rank of Intersections of Subgroups of Free Products of Orderable Groups" (Sept. 1, 2011). arXiv: 1109.0233 (cit. on p. 187).

[Art10] M. Artin. *Algebra*. 2nd ed. Addison Wesley, Aug. 13, 2010. 560 pp. (cit. on pp. 19, 42).

[Arz00] G. Arzhantseva. "A Property of Subgroups of Infinite Index in a Free Group". *Proceedings of the American Mathematical Society* 128.11 (2000), pp. 3205–3210 (cit. on p. 124).

[ASS14] V. Araujo, P. V. Silva, and M. Sykiotis. "Finiteness Results for Subgroups of Finite Extensions" (Feb. 3, 2014). arXiv: 1402.0401 [math] (cit. on p. 187).

[Bau66] B. Baumslag. "Intersections of Finitely Generated Subgroups in Free Products". *Journal of the London Mathematical Society* s1-41 (Jan. 1, 1966), pp. 673–679 (cit. on pp. 25, 237, 246, 259).

[Bau77] A. Baudisch. "Kommutationsgleichungen in Semifreien Gruppen". *Acta Mathematica Academiae Scientiarum Hungaricae* 29 (1977), pp. 235–249 (cit. on p. 199).

[Bau81] A. Baudisch. "Subgroups of Semifree Groups". *Acta Mathematica Academiae Scientiarum Hungaricae* 38 (1-4 Mar. 1981), pp. 19–28 (cit. on p. 199).

[BC12] J. Behrstock and R. Charney. "Divergence and Quasimorphisms of Right-Angled Artin Groups". *Mathematische Annalen* 352.2 (2012), pp. 339–356 (cit. on p. 196).

[BH92] M. Bestvina and M. Handel. "Train Tracks and Automorphisms of Free Groups". *Annals of Mathematics*. Second Series 135.1 (Jan. 1, 1992), pp. 1–51. JSTOR: 2946562 (cit. on p. 36).

[BK98] R. Burns and S.-M. Kam. "On the Intersection of Double Cosets in Free Groups, with an Application to Amalgamated Products". *Journal of Algebra* 210.1 (Dec. 1, 1998), pp. 165–193 (cit. on pp. 3, 25).

[BM15] O. Bogopolski and O. Maslakova. "An Algorithm for Finding a Basis of the Fixed Point Subgroup of an Automorphism of a Free Group". *International Journal of Algebra and Computation* 26 (01 Nov. 30, 2015), pp. 29–67 (cit. on pp. 36, 48).

[BMV07] O. Bogopolski, A. Martino, and E. Ventura. "The Automorphism Group of a Free-by-Cyclic Group in Rank 2". *Communications in Algebra* 35.5 (2007), pp. 1675–1690 (cit. on pp. 83, 84).

[BMV10] O. Bogopolski, A. Martino, and E. Ventura. "Orbit Decidability and the Conjugacy Problem for Some Extensions of Groups". *Transactions of the American Mathematical Society* 362.4 (2010), pp. 2003–2036 (cit. on pp. 25, 47, 50, 153).

[BNS87] R. Bieri, W. D. Neumann, and R. Strebel. "A Geometric Invariant of Discrete Groups". *Inventiones mathematicae* 90.3 (Oct. 1, 1987), pp. 451–477 (cit. on pp. 77–79).

[Bog+06] O. Bogopolski, A. Martino, O. Maslakova, and E. Ventura. "The Conjugacy Problem Is Solvable in Free-by-Cyclic Groups". *The Bulletin of the London Mathematical Society* 38.5 (2006), pp. 787–794 (cit. on p. 47).

[Bog00] O. Bogopolski. "Classification of Automorphisms of the Free Group of Rank 2 by Ranks of Fixed-Point Subgroups". *Journal of Group Theory* 3.3 (Mar. 2000), pp. 339–351 (cit. on pp. 36, 83).

[Bog08] O. Bogopolski. *Introduction to Group Theory*. Zurich, Switzerland: European Mathematical Society Publishing House, Feb. 29, 2008 (cit. on p. 25).

[BR84] G. Baumslag and J. E. Roseblade. "Subgroups of Direct Products of Free Groups". *Journal of the London Mathematical Society* s2-30.1 (Aug. 1, 1984), pp. 44–52 (cit. on p. 198).

[Bri00] P. Brinkmann. "Hyperbolic Automorphisms of Free Groups". *Geometric and Functional Analysis* 10 (Dec. 2000), pp. 1071–1089 (cit. on p. 91).

[Bro87] K. S. Brown. "Trees, Valuations, and the Bieri-Neumann-Strebel Invariant". *Inventiones mathematicae* 90.3 (Oct. 1, 1987), pp. 479–504 (cit. on p. 80).

[BS10] L. Bartholdi and P. V. Silva. "Rational Subsets of Groups" (Dec. 7, 2010). arXiv: 1012.1532 (cit. on p. 106).

[BS80] R. Bieri and R. Strebel. "Valuations and Finitely Presented Metabelian Groups". *Proceedings of the London Mathematical Society* s3-41.3 (Nov. 1, 1980), pp. 439–464 (cit. on pp. 69, 70, 77).

[Bur71] R. G. Burns. "On the Intersection of Finitely Generated Subgroups of a Free Group". *Mathematische Zeitschrift* 119.2 (June 1, 1971), pp. 121–130 (cit. on p. 124).

[But05] J. O. Button. "Fibred and Virtually Fibred Hyperbolic 3-Manifolds in the Censuses". *Experimental Mathematics* 14.2 (Jan. 1, 2005), pp. 231–255 (cit. on p. 80).

[Cav+17] B. Cavallo, J. Delgado, D. Kahrobaei, and E. Ventura. "Algorithmic Recognition of Infinite Cyclic Extensions". *Journal of Pure and Applied Algebra* 221.9 (Sept. 2017), pp. 2157–2179 (cit. on p. vii).

[CDK13] M. Casals-Ruiz, A. Duncan, and I. Kazachkov. "Embedddings between Partially Commutative Groups: Two Counterexamples". *Journal of Algebra* 390 (Sept. 15, 2013), pp. 87–99 (cit. on p. 204).

[CG05] C. Champetier and V. Guirardel. "Limit Groups as Limits of Free Groups". *Israel Journal of Mathematics* 146 (2005), pp. 1–75 (cit. on p. 202).

[CGW09] J. Crisp, E. Godelle, and B. Wiest. "The Conjugacy Problem in Subgroups of Right-Angled Artin Groups". *Journal of Topology* 2.3 (Jan. 1, 2009), pp. 442–460 (cit. on p. 199).

[CH10] L. Ciobanu and A. Houcine. "The Monomorphism Problem in Free Groups". *Archiv der Mathematik* 94.5 (2010), pp. 423–434 (cit. on p. 44).

[Cha07] R. Charney. "An Introduction to Right-Angled Artin Groups". *Geometriae Dedicata* 125.1 (2007), pp. 141–158 (cit. on p. 192).

[CL83] D. J. Collins and F. Levin. "Automorphisms and Hopficity of Certain Baumslag-Solitar Groups". *Archiv der Mathematik* 40.1 (Dec. 1983), pp. 385–400 (cit. on p. 87).

[CL89] M. M. Cohen and M. Lustig. "On the Dynamics and the Fixed Subgroup of a Free Group Automorphism". *Inventiones Mathematicae* 96.3 (1989), pp. 613–638 (cit. on p. 36).

[CL99] M. M. Cohen and M. Lustig. "The Conjugacy Problem for Dehn Twist Automorphisms of Free Groups". *Commentarii Mathematici Helvetici* 74.2 (June 1999), pp. 179–200 (cit. on p. 90).

[Dah03] F. Dahmani. "Combination of Convergence Groups". *Geometry & Topology* 7 (Nov. 12, 2003), pp. 933–963 (cit. on p. 203).

[Day09] M. B. Day. "Peak Reduction and Finite Presentations for Automorphism Groups of Right-Angled {A}rtin Groups". *Geometry & Topology* 13 (Jan. 8, 2009), pp. 817–855 (cit. on pp. 16, 42).

[Day14] M. B. Day. "Full-Featured Peak Reduction in Right-Angled Artin Groups". *Algebraic & Geometric Topology* 14.3 (May 29, 2014), pp. 1677–1743 (cit. on p. 42).

[Del14a] J. Delgado. "Some Characterizations of Howson PC-Groups". *Reports@SCM* 1.1 (Oct. 1, 2014), pp. 33–38 (cit. on p. viii).

[Del14b] J. Delgado. "{Whitehead Problems for Words in Fn x Zm". In: *Extended Abstracts Fall 2012*. Ed. by J. González-Meneses, M. Lustig, and E. Ventura. Trends in Mathematics. Springer International Publishing, Jan. 1, 2014, pp. 35–38 (cit. on p. vii).

[DF01] W. Dicks and E. Formanek. "The Rank Three Case of the Hanna Neumann Conjecture". *Journal of Group Theory* 4.2 (2001), pp. 113–151 (cit. on p. 124).

[DG11] F. Dahmani and V. Guirardel. "The Isomorphism Problem for All Hyperbolic Groups". *Geometric and Functional Analysis* 21.2 (Apr. 23, 2011), pp. 223–300 (cit. on p. 91).

[DG81] J. L. Dyer and E. K. Grossman. "The Automorphism Groups of the Braid Groups". *American Journal of Mathematics* 103.6 (Dec. 1, 1981), pp. 1151–1169. JSTOR: 2374228 (cit. on p. 90).

[DI08] W. Dicks and S. V. Ivanov. "On the Intersection of Free Subgroups in Free Products of Groups". *Mathematical Proceedings of the Cambridge Philosophical Society* 144.3 (2008), pp. 511–534 (cit. on p. 124).

[DI10] W. Dicks and S. V. Ivanov. "On the Intersection of Free Subgroups in Free Products of Groups with No 2-Torsion". *Illinois Journal of Mathematics* 54.1 (2010), pp. 223–248 (cit. on p. 187).

[Dic83] W. Dicks. "Automorphisms of the Polynomial Ring in Two Variables". *Publicacions de la Secció de Matemàtiques. Universitat Autònoma de Barcelona* 27.1 (1983), pp. 155–162 (cit. on p. 84).

[Dic94] W. Dicks. "Equivalence of the Strengthened Hanna Neumann Conjecture and the Amalgamated Graph Conjecture". *Inventiones mathematicae* 117.1 (Dec. 1, 1994), pp. 373–389 (cit. on p. 124).

[Die87] V. Diekert. *On the Knuth-Bendix Completion for Concurrent Processes*. 1987 (cit. on p. 199).

[DK93] G. Duchamp and D. Krob. "Free Partially Commutative Structures". *Journal of Algebra* 156.2 (Apr. 15, 1993), pp. 318–361 (cit. on p. 199).

[Dro87a] C. Droms. "Graph Groups, Coherence, and Three-Manifolds". *Journal of Algebra* 106.2 (Apr. 1, 1987), pp. 484–489 (cit. on p. 198).

[Dro87b] C. Droms. "Isomorphisms of Graph Groups". *Proceedings of the American Mathematical Society* 100.3 (Mar. 1987), pp. 407–407 (cit. on p. 193).

[Dro87c] C. Droms. "Subgroups of Graph Groups". *Journal of Algebra* 110.2 (Oct. 15, 1987), pp. 519–522 (cit. on p. 197).

[DSS92] C. Droms, B. Servatius, and H. Servatius. "Groups Assembled from Free and Direct Products". *Discrete Mathematics* 109 (1–3 Nov. 12, 1992), pp. 69–75 (cit. on p. 198).

[DT06] N. M. Dunfield and D. P. Thurston. "A Random Tunnel Number One 3–manifold Does Not Fiber over the Circle". *Geometry & Topology* 10.4 (Dec. 15, 2006), pp. 2431–2499 (cit. on p. 80).

[Dun01] N. M. Dunfield. "Alexander and Thurston Norms of Fibered 3-Manifolds". *Pacific Journal of Mathematics* 200.1 (Sept. 1, 2001), pp. 43–58 (cit. on p. 80).

[DV13] J. Delgado and E. Ventura. "Algorithmic Problems for Free-Abelian Times Free Groups". *Journal of Algebra* 391 (Oct. 1, 2013), pp. 256–283 (cit. on pp. vii, 25, 128, 208).

[DV17a] J. Delgado and E. Ventura. "Stallings Graphs for Free-Abelian by Free Groups". *(preprint)* (2017) (cit. on p. vii).

[DV17b] J. Delgado and E. Ventura. "Twisted-Conjugacy Problem for Free-Abelian Times Free Groups". *(preprint)* (2017) (cit. on p. vii).

[DVZ17] J. Delgado, E. Ventura, and A. Zakharov. "Intersection Problems for Droms Groups". *(preprint)* (2017) (cit. on p. viii).

[EKR05] E. S. Esyp, I. V. Kazachkov, and V. N. Remeslennikov. "Divisibility Theory and Complexity of Algorithms in Free Partially Commutative Groups" (Dec. 16, 2005). arXiv: `math/0512401` (cit. on pp. 192, 199).

[FH14] M. Feighn and M. Handel. "Algorithmic Constructions of Relative Train Track Maps and CTs" (Nov. 23, 2014). arXiv: `1411.6302 [math]` (cit. on p. 36).

[Fri15] J. Friedman. “Sheaves on Graphs, Their Homological Invariants, and a Proof of the Hanna Neumann Conjecture: With an Appendix by Warren Dicks”. *Memoirs of the American Mathematical Society* 233.1100 (Jan. 2015), pp. 0–0 (cit. on p. 124).

[Ger87] S. Gersten. “Fixed Points of Automorphisms of Free Groups”. *Advances in Mathematics* 64.1 (Apr. 1987), pp. 51–85 (cit. on p. 36).

[Gol78] M. C. Golumbic. “Trivially Perfect Graphs”. *Discrete Mathematics* 24.1 (Jan. 1, 1978), pp. 105–107 (cit. on p. 211).

[Gre90] E. R. Green. “Graph Products of Groups”. 1990 (cit. on pp. 13, 192).

[GT86] R. Z. Goldstein and E. C. Turner. “Fixed Subgroups of Homomorphisms of Free Groups”. *Bulletin of the London Mathematical Society* 18.5 (Jan. 9, 1986), pp. 468–470 (cit. on p. 36).

[HM95] S. Hermiller and J. Meier. “Algorithms and Geometry for Graph Products of Groups”. *Journal of Algebra* 171.1 (Jan. 1, 1995), pp. 230–257 (cit. on p. 199).

[How54] A. G. Howson. “On the Intersection of Finitely Generated Free Groups”. *Journal of the London Mathematical Society* s1-29.4 (Oct. 1, 1954), pp. 428–434 (cit. on pp. 24, 123, 201).

[Hum94] S. P. Humphries. “On Representations of Artin Groups and the Tits Conjecture”. *Journal of Algebra* 169.3 (1994), pp. 847–862 (cit. on p. 13).

[Iva01] S. V. Ivanov. “Intersecting Free Subgroups in Free Products of Groups”. *International Journal of Algebra and Computation* 11 (03 June 1, 2001), pp. 281–290 (cit. on pp. 124, 191, 218).

[Iva08] S. V. Ivanov. “On the Kurosh Rank of the Intersection of Subgroups in Free Products of Groups”. *Advances in Mathematics* 218.2 (2008), pp. 465–484 (cit. on pp. 124, 218, 228).

[Iva99] S. V. Ivanov. “On the Intersection of Finitely Generated Subgroups in Free Products of Groups”. *International Journal of Algebra and Computation* 09 (05 Oct. 1, 1999), pp. 521–528 (cit. on pp. 218, 219, 226, 228, 235, 237, 246).

[Jec73] T. J. Jech. *The Axiom of Choice*. Amsterdam; New York: North-Holland Pub. Co. ; American Elsevier Pub. Co., 1973 (cit. on p. 151).

[JJC96] Y. Jing-Ho, C. Jer-Jeong, and G. J. Chang. “Quasi-Threshold Graphs”. *Discrete Applied Mathematics* 69.3 (Aug. 27, 1996), pp. 247–255 (cit. on p. 211).

[Kam09] M. Kambites. “On Commuting Elements and Embeddings of Graph Groups and Monoids”. *Proceedings of the Edinburgh Mathematical Society (Series 2)* 52 (01 2009), pp. 155–170 (cit. on p. 204).

[Kha02] B. Khan. “Positively Generated Subgroups of Free Groups and the Hanna Neumann Conjecture”. *Contemporary Mathematics* 296 (2002), pp. 155–170 (cit. on p. 124).

[KK13] S.-H. Kim and T. Koberda. “Embedability between Right-Angled Artin Groups”. *Geometry & Topology* 17.1 (Apr. 2, 2013), pp. 493–530 (cit. on p. 204).

[KLV01] S. Krstić, M. Lustig, and K. Vogtmann. “An Equivariant Whitehead Algorithm and Conjugacy for Roots of Dehn Twist Automorphisms”. *Proceedings of the Edinburgh Mathematical Society* 44 (01 Feb. 2001), p. 117 (cit. on p. 91).

[KM02] I. Kapovich and A. Myasnikov. "Stallings Foldings and Subgroups of Free Groups". *Journal of Algebra* 248.2 (Feb. 15, 2002), pp. 608–668 (cit. on pp. 20, 106).

[KMM15] N. Koban, J. McCammond, and J. Meier. "The BNS-Invariant for the Pure Braid Groups". *Groups, Geometry, and Dynamics* 9.3 (2015), pp. 665–682 (cit. on p. 61).

[Kob13] T. Koberda. "Right Angled Artin Groups and Their Subgroups". *Lecture notes Yale University* (Mar. 7, 2013), pp. 1–50 (cit. on pp. 192, 193).

[KP14] N. Koban and A. Piggott. "The Bieri–Neumann–Strebel Invariant of the Pure Symmetric Automorphisms of a Right-Angled Artin Group". *Illinois Journal of Mathematics* 58.1 (2014), pp. 27–41 (cit. on p. 61).

[Krs89] S. Krstić. "Actions of Finite Groups on Graphs and Related Automorphisms of Free Groups". *Journal of Algebra* 124.1 (July 1989), pp. 119–138 (cit. on p. 90).

[KWM05] I. Kapovich, R. Weidmann, and A. Myasnikov. "Foldings, Graphs of Groups and the Membership Problem". *International Journal of Algebra and Computation* 15 (01 Feb. 2005), pp. 95–128 (cit. on pp. 200, 218, 219, 226, 228).

[Lau95] M. R. Laurence. "A Generating Set for the Automorphism Group of a Graph Group". *Journal of the London Mathematical Society* 52.2 (Jan. 10, 1995), pp. 318–334 (cit. on p. 16).

[Lev08] G. Levitt. "Unsolvability of the Isomorphism Problem for [Free Abelian]-by-Free Groups" (Oct. 6, 2008). arXiv: 0810.0935 [math] (cit. on p. 153).

[Los96] J. E. Los. "On the Conjugacy Problem for Automorphisms of Free Groups". *Topology* 35.3 (July 1996), pp. 779–806 (cit. on p. 90).

[LS01] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*. Reprint. Springer, Mar. 1, 2001. 353 pp. (cit. on pp. 9, 18, 31, 42, 51, 95).

[LS08] M. Lohrey and B. Steinberg. "The Submonoid and Rational Subset Membership Problems for Graph Groups". *Journal of Algebra*. Computational Algebra 320.2 (July 15, 2008), pp. 728–755 (cit. on pp. 200, 210).

[Lus00] M. Lustig. *Structure and Conjugacy for Automorphisms of Free Groups I*. Bonn: Max-Planck Inst. Math., 2000 (cit. on p. 90).

[Lus01] M. Lustig. *Structure and Conjugacy for Automorphisms of Free Groups II*. Bonn: Max-Planck Inst. Math., 2001 (cit. on p. 90).

[Lus07] M. Lustig. "Conjugacy and Centralizers for Iwip Automorphisms of Free Groups". In: *Geometric Group Theory*. Ed. by G. N. Arzhantseva, J. Burillo, L. Bartholdi, and E. Ventura. Trends in Mathematics. Birkhäuser Basel, 2007, pp. 197–224 (cit. on p. 90).

[LWZ90] H.-N. Liu, C. Wrathall, and K. Zeger. "Efficient Solution of Some Problems in Free Partially Commutative Monoids". *Information and Computation* 89.2 (Dec. 1990), pp. 180–198 (cit. on p. 199).

[Mak82] G. Makanin. "Equations in Free Groups (Russian)". *Izv. Akad. Nauk SSSR Ser. Mat.* 46 (1982), pp. 1190–1273 (cit. on p. 44).

[Mar07] L. Markus-Epstein. "Stallings Foldings and Subgroups of Amalgams of Finite Groups". *International Journal of Algebra and Computation* 17.8 (2007), pp. 1493–1535 (cit. on p. 219).

[Mei90] G.-N. Meigniez. “Bouts d’un groupe opérant sur la droite, I : théorie algébrique”. *Annales de l’institut Fourier* 40.2 (1990), pp. 271–312 (cit. on p. 78).

[Mik58] K. A. Mikhailova. “The Occurrence Problem for Direct Products of Groups”. *Doklady Akademii Nauk SSSR* 119 (1958), pp. 1103–1105 (cit. on pp. 199, 200, 217).

[Mik68] K. A. Mikhailova. “The Occurrence Problem for Free Products of Groups”. *Mathematics of the USSR-Sbornik* 4.2 (Feb. 28, 1968), pp. 181–190 (cit. on p. 217).

[Mil04] C. F. Miller III. “Combinatorial Group Theory (Notes)”. Apr. 7, 2004 (cit. on p. 230).

[Mil92] C. F. Miller III. “Decision Problems for Groups — Survey and Reflections”. In: *Algorithms and Classification in Combinatorial Group Theory*. Ed. by G. Baumslag and C. F. Miller. Mathematical Sciences Research Institute Publications 23. Springer New York, 1992, pp. 1–59 (cit. on pp. 62, 76, 77).

[Min12] I. Mineyev. “Submultiplicativity and the Hanna Neumann Conjecture”. *Annals of Mathematics* 175.1 (Jan. 1, 2012), pp. 393–414 (cit. on p. 124).

[MKS04] W. Magnus, A. Karrass, and D. Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. 2 Revised. Dover Publications, Nov. 12, 2004. 464 pp. (cit. on p. 95).

[MV95] J. Meier and L. Vanwyk. “The Bieri-Neumann-Strebel Invariants for Graph Groups”. *Proceedings of the London Mathematical Society* s3-71.2 (Sept. 1, 1995), pp. 263–280 (cit. on p. 61).

[MW02] J. Meakin and P. Weil. “Subgroups of Free Groups: A Contribution to the Hanna Neumann Conjecture”. *Geometriae Dedicata* 94.1 (Oct. 1, 2002), pp. 33–43 (cit. on p. 124).

[Neu56] H. Neumann. “On the Intersection of Finitely Generated Free Groups.” *Publicationes Mathematicae* 4 (1956), pp. 186–189 (cit. on p. 124).

[Neu90] W. D. Neumann. “On Intersections of Finitely Generated Subgroups of Free Groups”. In: *Groups—Canberra 1989*. Ed. by L. G. Kovács. Vol. 1456. Springer Berlin Heidelberg, 1990, pp. 161–170 (cit. on p. 124).

[Nos16] G. A. Noskov. “Mineyev–Dicks Proof of the HN-Conjecture and the Euler–Poincaré Characteristic”. *Mathematical Notes* 99 (3-4 Mar. 1, 2016), pp. 390–396 (cit. on p. 124).

[PS10] S. Papadima and A. I. Suciu. “Bieri–Neumann–Strebel–Renz Invariants and Homology Jumping Loci”. *Proceedings of the London Mathematical Society* 100.3 (Jan. 5, 2010), pp. 795–834 (cit. on p. 61).

[Rab58] M. O. Rabin. “Recursive Unsolvability of Group Theoretic Problems”. *Annals of Mathematics*. Second Series 67.1 (Jan. 1, 1958), pp. 172–194. JSTOR: 1969933 (cit. on p. 62).

[Rob96] D. J. S. Robinson. *A Course in the Theory of Groups*. Vol. 80. Graduate Texts in Mathematics. New York, NY: Springer New York, 1996 (cit. on p. 56).

[RSS13] E. Rodaro, P. V. Silva, and M. Sykiotis. “Fixed Points of Endomorphisms of Graph Groups”. *Journal of Group Theory* 16.4 (Jan. 5, 2013), pp. 573–583 (cit. on pp. 41, 202, 204, 210).

[Rub15] C. Rubio-Montiel. “A New Characterization of Trivially Perfect Graphs”. *Electronic Journal of Graph Theory and Applications (EJGTA)* 3.1 (Mar. 22, 2015), pp. 22–26 (cit. on p. 211).

[Sah15] J. Sahattchieve. “On Convex Hulls and the Quasiconvex Subgroups of Fm x Zn”. *Groups, Complexity, Cryptology* 7.1 (2015), pp. 69–80 (cit. on p. 34).

[Seg90] D. Segal. “Decidable Properties of Polycyclic Groups”. *Proc. London Math. Soc* 3 (1990), pp. 61–497 (cit. on p. 89).

[Sel95] Z. Sela. “The Isomorphism Problem for Hyperbolic Groups I”. *Annals of Mathematics*. Second Series 141.2 (Mar. 1, 1995), pp. 217–283. JSTOR: 2118520 (cit. on pp. 90, 91).

[Ser89] H. Servatius. “Automorphisms of Graph Groups”. *Journal of Algebra* 126.1 (Oct. 1989), pp. 34–60 (cit. on p. 196).

[Sou08] L. Soukup. “Infinite Combinatorics: From Finite to Infinite”. In: *Horizons of Combinatorics*. Ed. by E. Győri, G. O. H. Katona, L. Lovász, and G. Sági. Bolyai Society Mathematical Studies 17. Springer Berlin Heidelberg, 2008, pp. 189–213 (cit. on p. 102).

[SSV16] P. V. Silva, X. Soler-Escrivà, and E. Ventura. “Finite Automata for Schreier Graphs of Virtually Free Groups”. *Journal of Group Theory* 19.1 (2016), pp. 25–54 (cit. on p. 219).

[Sta83] J. R. Stallings. “Topology of Finite Graphs”. *Inventiones Mathematicae* 71 (Mar. 1983), pp. 551–565 (cit. on pp. 20, 106, 115, 119, 201, 208).

[Sti93] J. Stillwell. *Classical Topology and Combinatorial Group Theory*. Vol. 72. Graduate Texts in Mathematics. New York, NY: Springer New York, 1993 (cit. on p. 102).

[Str12] R. Strebel. *Notes on the Sigma Invariants*. Apr. 1, 2012. arXiv: 1204.0214 (cit. on pp. 77–79).

[Tar92] G. Tardos. “On the Intersection of Subgroups of a Free Group”. *Inventiones mathematicae* 108.1 (Dec. 1, 1992), pp. 29–36 (cit. on p. 124).

[Tar96] G. Tardos. “Towards the Hanna Neumann Conjecture Using Dicks’ Method”. *Inventiones Mathematicae* 123.1 (1996), pp. 95–104 (cit. on p. 124).

[Tou06] N. W. M. Touikan. “A Fast Algorithm for Stallings’ Folding Process”. *International Journal of Algebra and Computation* 16 (06 Dec. 1, 2006), pp. 1031–1045 (cit. on p. 119).

[Tur95] E. C. Turner. “Finding Invisible Nielsen Paths for a Train Tracks Map”. *Proc. of a workshop held at Heriot-Watt Univ., Edinburg, 1993 (Lond. Math. Soc. Lect. Note Ser., 204), Cambridge, Cambridge Univ. Press.* (1995), pp. 300–313 (cit. on p. 36).

[Tyr71] J. Tyrer. “On Direct Products and the Hopf Property”. University of Oxford, 1971 (cit. on p. 14).

[Van94] L. Van Wyk. “Graph Groups Are Biautomatic”. *Journal of Pure and Applied Algebra* 94.3 (July 8, 1994), pp. 341–352 (cit. on p. 199).

[Ven02] E. Ventura. “Fixed Subgroups in Free Groups: A Survey”. In: *Combinatorial and Geometric Group Theory (New York, 2000/Hoboken, NJ, 2001)*. Vol. 296. Contemp. Math. Providence, RI: Amer. Math. Soc., 2002, pp. 231–255 (cit. on p. 36).

[Whi36] J. H. C. Whitehead. “On Equivalent Sets of Elements in a Free Group”. *The Annals of Mathematics*. Second Series 37.4 (Oct. 1, 1936), pp. 782–800. JSTOR: 1968618 (cit. on pp. 41, 44).

[Wil05] H. J. R. Wilton. “An Introduction to Limit Groups”. *Series for Telgiggy, Imperial College* (3-3-05) (cit. on p. 202).

[Wis05] D. T. Wise. “The Coherence of One-Relator Groups with Torsion and the Hanna Neumann Conjecture”. *Bulletin of the London Mathematical Society* 37.5 (Jan. 10, 2005), pp. 697–705 (cit. on p. 124).

[Wol62] E. S. Wolk. “The Comparability Graph of a Tree”. *Proceedings of the American Mathematical Society* 13.5 (Oct. 1, 1962), pp. 789–795. JSTOR: 2034179 (cit. on p. 211).

[Wol65] E. S. Wolk. “A Note on "The Comparability Graph of a Tree"”. *Proceedings of the American Mathematical Society* 16.1 (Feb. 1, 1965), pp. 17–20. JSTOR: 2033992 (cit. on p. 211).

[Wra88] C. Wrathall. “The Word Problem for Free Partially Commutative Groups”. *J. Symb. Comput.* 6.1 (Aug. 1988), pp. 99–104 (cit. on p. 199).

[Wra89] C. Wrathall. *Free Partially Commutative Groups*. 1989 (cit. on p. 199).

[Zak14] A. Zakharov. “On the Rank of the Intersection of Free Subgroups in Virtually Free Groups”. *Journal of Algebra* 418 (2014), pp. 29–43 (cit. on p. 187).

[Zim85] B. Zimmermann. *Zur Klassifikation höherdimensionaler Seifertscher Faserräume. (On the classification of higher dimensional Seifert fibre spaces)*. 1985 (cit. on pp. 153, 154).

Websites

[Dic12] W. Dicks. *Simplified Mineyev's Proof of Hanna Neumann Conjecture*. May 2012. URL: http://mat.uab.cat/~dicks/SimplifiedMineyev.pdf (cit. on p. 124).

Index